Cruz Operations Center (CruzOC)

User Guide

dorado
SOFTWARE

January 15, 2021

Dorado Software
4805 Golden Foothill Parkway
El Dorado Hills, CA 95762

888-939-9959

Mail corrections or
support questions to support@doradosoftware.com

www.doradosoftware.com

Produced in the United States of America.

Because Dorado Software is constantly striving to improve its products, the features and screenshots that appear in this manual may only approximately match those in the actual application described here. This confidential and proprietary document is the property of Dorado Software. All trademarks in this manual are the property of their respective holders.

Contents

Preface

Cruz Operations Center (CruzOC), often referred to as simply Cruz in user documentation, from Dorado Software offers automated, consolidated configuration and control of your converged infrastructure resources, including servers, storage and network devices. You can customize CruzOC and unify multiple systems for a single view of infrastructure assets, monitoring, compliance auditing, troubleshooting and much more.

Why Cruz Operations Center?

Productive

Discovery and wizard-driven configuration features are available within minutes of installation, which let you quickly discover, monitor and manage your multi-technology, multi-vendor infrastructure.

Easy

CruzOC provides the infrastructure information you need from a single pane-of-glass and offers advanced capabilities with minimal configuration overhead.

Value

CruzOC often costs less to use and maintain than most other solutions.

Scalable

You can scale CruzOC to almost any size.

NOTE:

CruzOC is a custom software product. The underlying software code, debug statements, installation files, Java classes, license entries, Logs and so on, may refer to names other than CruzOC. Such names only have meaning for troubleshooting or support. Most users can safely ignore these. Examples include Redcell, Synergy, Oware, and Liferay.

Application Overview

The application provides a flexible system to manage networks big and small. It distributes processing between the following elements:

- **Application Server** is the application’s “central processor.”
- **Web Portal Server** providesclients (browsers) with information based on the Application Server’s processes.
- **Mediation Server** managesmessageprocessingbetweenthe Application Serverand managed devices.
- **Database Server** stores and retrieves information about managed devices.

The installation wizard displays these options because you can install each of them to separate servers. Single-serverinstallationsarepossible. However, to manage largernetworksand provide failover High Availability (HA) installations, distributed and redundant installations are also available.

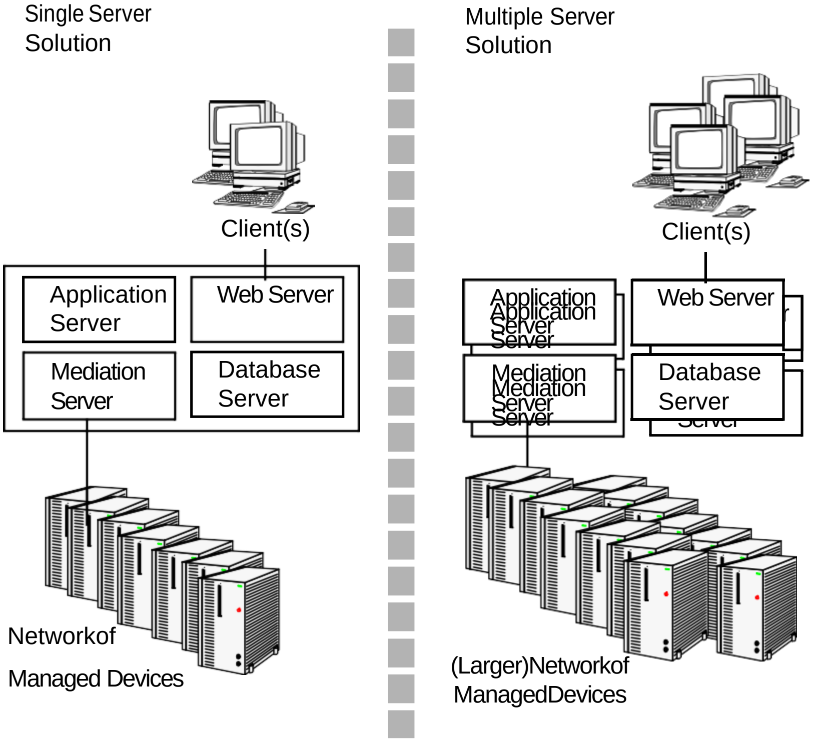

Refer to the *Installation Guide* for more information about distributed and HA installations.

Key Features

The following are some key features of CruzOC:

Feature	Description
Customizable and Flexible Web Portal	You can customize the web portal, even providing custom designed views of your data assigned to individual users. You can even create web portal accounts for departments, geographic areas, or other criteria.
Automate and Schedule Device Discovery	Device discovery populates the database and begins network analysis. You can also create schedules to automatically run Discovery whenever you need to update the initially discovered conditions.
Cruz Administration	You can administer your network—adding devices, user accounts, and web portal displays—from a secure console on your network.
Open Integration	Supports industry standards. It comes with an open-source MySQL database, and supports using Oracle databases. It also uses industry-standard MIBs and protocols, and even lets you install open-source screen elements like Google gadgets to its web portal.
Network Topology	The topology screen lets you create multi-layered, customizable, web-based top ology displays of your network to help track the state of network devices.
Alarms	Alarms respond to hundreds of possible network scenarios by default, and you can configure them to including multiple condition checks. Alarms help you recognize issues before your network's users experience productivity losses. Alarms can also trigger actions like e-mailing technical support staff, executing Perl scripts, paging, emitting SNMP traps to other systems, Syslog messaging, and executing external application.
Traps and Syslog	Let's you investigate network issues by examining traps and Syslog messages. You can set up events/alarms and then receive, process, forward, and send syslog and trap messages.
Reports and Graphs	This application comes with many pre-configured reports to di splay data from its database. You can archive and compare reports or automate creating them with the scheduler. If your package includes them, you can also configure graphs to present performance and traffic flow data.
Modularity	Modules analyze network traffic, manage services and IP address and subnet allocations. Save time adding modules to existing deployments to add feature functionality without requiring additional standalone software.

Networks with CruzOC

The beginning ofmanagement with CruzOC is executing Discovery Profiles to discoverresources on a network. After that occurs, you can configure System Topology, Resource Monitors, and Performance Dashboards.

After these initial steps, CruzOC helps you understand and troubleshoot you r network's conditions. For example: Suppose a CruzOC Performance Dashboard displays something you want to troubleshoot. You can right-click the impacted device in the Managed Resources portlet to access its configuration and initiate potential actions on it. The Network Status icon in the view indicates the device status. Its Connected Devices panel al so displays the highest severity alarm on the device or its sub-components. F or example, red indicates a *Critical* al arm.

In screens like Connected Devices on page 189 you can examine each section of device information and right-click components to see further applicable actions. Forexample, right-click to Show Performance, and edit and/or save that view of performance as another Performance Dashboard. Performance can also appear in portlets that Show Top Talkers (the busiest devices) or Show Key Metrics.

Fromlooking at Performance Dashboards or Top N [Assets] you may conclude some configuration changes made memory consumption spike. Right- click to access resource actions under File Management Menu that let you see the current configuration files on devices and compare current to previous. Youcan also back up devices (see Backing Up Configurations on page 462) and restore previously backed up files (see Restoring Configurations on page 463). Finally, you may simply want to Resync (another right-click menu item) to ensure the device and your management system are up to date.

NOTE:

Alternatively, the Alarms portlet also lets you right-click to expose Alarm Actions.

You can right-click for Direct Access – Telnet or Direct Access – MIB Browser to display a command line telnetting to the device, or an SNMP MIB browser to examine SNMP possibilities for it.

Click the plus in the upper right corner of any portlet to see its expanded version, for example: The Managed Resources expanded view. This displays detail or “Snap-in” panels at its bottom, with additional information about a selected resource.

Generating a Report let you take snapshots of network conditions to aid in analysis of trends, and Audit Trail portlets track message traffic between CruzOC and devices.

Additional Products

The following describes how to increase the power of your installation. While the documents describe everything available with the product, your installation may provide only a limited set of those features.

Updating Your License

If you have a limited license — for example CruzOC may limit discovery to a certain number of devices — then it does not function outside those licensedlimits.

You can purchase additional capabilities, and can update your license by putting the updated license file in a convenient directory. Click *License Management* from the Settings > Application Configuration Settings portlet to register your license Y our updated license should be visible in the License Viewer (See License Viewer Window on page 56 for details.)

NOTE:

If you update your installation from a previous one where you upgraded license, you must also install any new licenses.

If you import a license that, for example, changes the application's expiration date, it often does not immediately take effect. Log out and log in again to have newly licensed capabilities immediately (with a few exceptions that ma y make you wait).

Licenses support three expiration formats: *Never*, *Date certain*,and a format that indicates the license will be valid for *a number of days after registration* .

Online Help/Filter

Access general online help by clicking *Help* from the portal header. Help appropriate to each portlet appears when you click the help tool (question m ark) on the portlet title b ar.

By default, this opens a separate browser window, which is not necessarily always in front of the screen that calls it. Because it is separate, you can arrange the display, so the help screen does not conceal the portlet it describes. Click the *Show* button to display the contents, index and search tabs (*Hide* conceals them again), and the *Prev/Next* button s, or clicking table of contents to pics moves to different topics within the help set.

NOTE:

A browser's cache may interfere with help's correct appearance. If you see a table of contents node without contents, refresh (Reload) the help window or the application page.

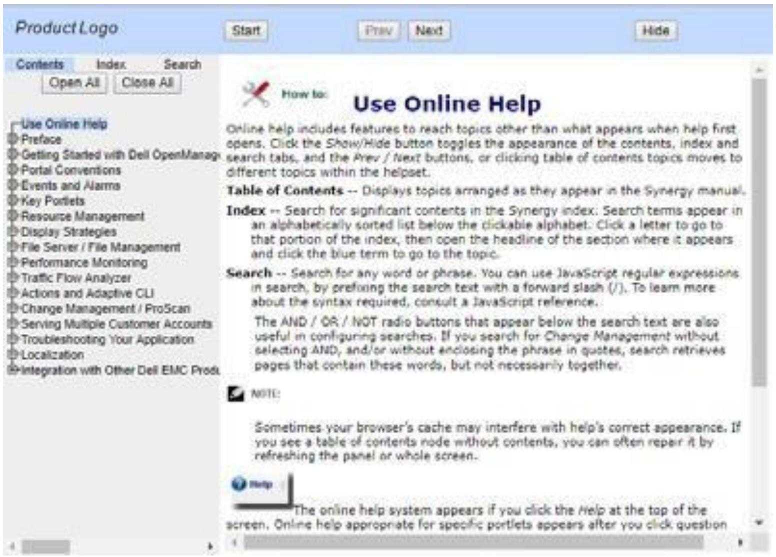

Use Online Help

Online help includes features to reach topics other than what appears when help first opens. Click the *Show/Hide* button toggles the appearance of the contents, index and search tabs, and the *Prev / Next* buttons, or clicking table of contents topics moves to different topics within the helpset.

Table of Contents -- Displays topics arranged as they appear in the Synergy manual.

Index -- Search for significant contents in the Synergy index. Search terms appear in an alphabetically sorted list below the clickable alphabet. Click a letter to go to that portion of the index, then open the headline of the section where it appears and click the blue term to go to the topic.

Search -- Search for any word or phrase. You can use JavaScript regular expressions in search, by prefixing the search text with a forward slash (/). To learn more about the syntax required, consult a JavaScript reference.

The AND / OR / NOT radio buttons that appear below the search text are also useful in configuring searches. If you search for *Change Management* without selecting AND, and/or without enclosing the phrase in quotes, search retrieves pages that contain these words, but not necessarily together.

NOTE:

Sometimes your browser's cache may interfere with help's correct appearance. If you see a table of contents node without contents, you can often repair it by refreshing the panel or whole screen.

The online help system appears if you click the *Help* at the top of the screen. Online help appropriate for specific portlets appears after you click question

A Note About Performance

CruzOC is designed to help you manage your network with alacrity. Unfortunately, the devices managed or the networks that communicate with those devices are not always as fast as this software. If discovery takes a long time (it can), often network and device latency is the culprit. You can also optimize installations to be faster (see the recommendations in the *Installation Guide* and the first chapter of this guide), and limit device queries with filters, but device and network latency limit how quickly your system can respond.

NOTE:

If you use management systems other than this one, you must perform a device level resync before performing configuration actions. Best practice is to use a single management tool whenever possible.

1

Getting Started

The section outlines the steps in a typical installation. Because the software described here is both flexible and powerful, this section does not exhaustively describe all the details of available installations. Instead, references are made to those descriptions elsewhere in this guide, the *CruzOC Installation Guide*, or online help. This guide assumes that the product is installed and operational.

This section focuses on the tasks a user performs. F or a detailed description of the CruzOC portal, windows, portlets, options, features, and so on.

In addition to the general information and the portlets and editors descript ion, this section provides administrative tasks. If you already have a good understanding of the general information and the Control Panel, portlets, and editors, go directly to the tasks you want to perform.

Administrative Tools

In addition to the Control Panel – 32, this section describes the following portlets, editors, and windows used by a system administrator or anyone else responsible for the tasks described in this section:

- Aging Policies Editor
- Quick Navigation Portlet
- Network Tools
- License Viewer Window
- Common Setup Tasks Portlet

Control Panel

Use the Control panel to configure access to the application. You must be signed in as a user with administrative permissions. (The default *admin* user has such permissions.) The *Go to > Control Panel* menu item opens a screen with the following tabs of interest:

- Administrator
- [Domains]
- Portal > Users and Organizations
- Public/Private Page Behavior
- Portal > Roles
- Portal > Portal Settings
- Portal > [Other]
- Redcell > Permission Manager
- Redcell > Database Aging Policies (DAP)
- Redcell > Data Configuration
- Redcell > Mediation
- Redcell > Filter Management
- Redcell > Application Settings
- Server Administration

Each screen begins with its description. Click the question mark button to hide or show this description. Tool tips appear when you hover the cursor over fields or click the question mark button.

Users and Organizations

Users are individuals who perform tasks using the portal. Administrators can permissions in roles, or be collected into user groups. User groups can cross may also be used to customize personal site templates for members.

Users without Administrator permissions may not see all of the features described in this guide.

See Control Panel on page 32 for the Control Panel capabilities.

Depending on your installation and configuration, some options may not be available to you, such as the Contacts portlet is not related to the Control panel Contacts Center.

Sometimes Cruz maydisplay Control Panelobjectslikeusers, roles, andorganizationsinaccurately. This occurs because **search Indexes** need to be re-indexed every so often, especially when changes to roles, users and organizations are frequent.

Tore-index go to Control Panel > Server Administration and then click on the *Reindex allsearch indexes*. This takes little time.

Administrator

Toconfigure information for your login, look for the bar titled with your account login's name. It has the following lines beneath it:

My Account —Thisconfigures your information as a user, including your e-mail address, password, and so on.

Contacts Center—This configures contacts, in other words, people within your system that you are following. This is *not* the same as customer contacts as in the Contacts portlet (see Contacts Portlet on page 215).

Click the *Find People* link to see a list of potential contacts within your system. Youmust click *Action* > *Follow* to see them listed in the *Contacts Home.* Use the *Action* button to explore other possibilities.

The contact must approve you in their requests. To *Follow* means you want to receive the followed person's activity stream, blog postings, and so on. *Friend*ing means your friendscan see your activity and you can see theirs. They must accept any *Friend*request.

✎ NOTE:

You can export vCards for all contacts in the system to use with other software that uses contacts. For example: e-mail clients.

[Domains]

A default domain name Cruz appears in *Control Panel*. *Global* and *Administrator's Personal Site,* or *[Multitenant Site Names]* site configurations may appear as additional items to configure when you expand the domains list. The *Global* option is unrelated to Cruz functionality. Refer to the *Installation Guide* or online help for more about Multitenancy, also referred to as MSP (Multitenant Service Provider) capabilities.

✎ NOTE:

Seewhether Multitenancy is installed fromthe Manage > Show Versions > Installed Extensions list. It shows as the Synergy MSP Extension.

The domain options configure the site settings and the portal's overall look and feel, reference information, and so on. See the tooltips for more complete descriptions. This also configures pages, documents, calendars, blogs, wikis, polls and so on.

Social Activity lets you alter measurements for user participation in organizations. Equity values determine the reward value of an action; equity lifespans determine when to age the reward of action.

Portal > Users and Organizations

In these screens you can create Users you later assign to Roles and Locations with the appropriate permissions (Roles for operators, administrators, and so on). The limit for User Names is 70 characters. Define the default password policy in the Control Panel under Portal > Password Policies.

Users perform tasks using the portal. Administrators can create new users or deactivate existing users. You can organize users in a hierarchy of organizations and delegate administrative rights.

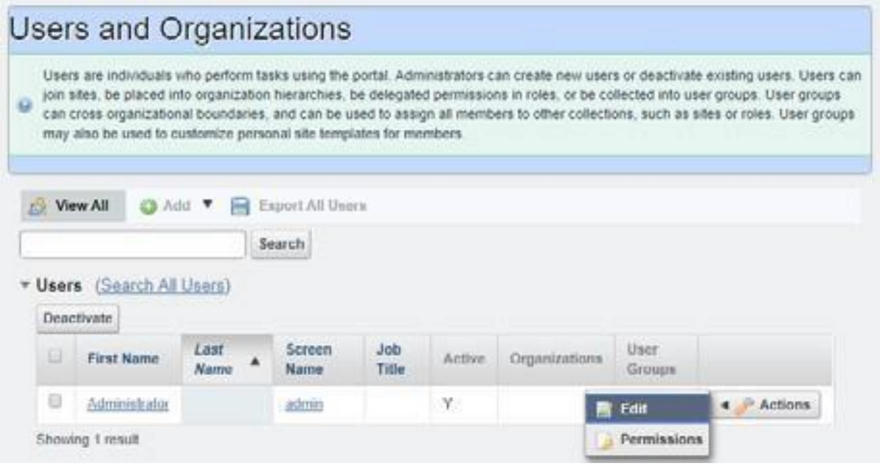

Users and Organizations

Users are individuals who perform tasks using the portal. Administrators can create new users or deactivate existing users. Users can join sites, be placed into organization hierarchies, be delegated permissions in roles, or be collected into user groups. User groups can cross organizational boundaries, and can be used to assign all members to other collections, such as sites or roles. User groups may also be used to customize personal site templates for members.

□	First Name	Last Name ▲	Screen Name	Job Title	Active	Organizations	User Groups	
□	Administrator		admin		Y			Actions

Showing 1 result

After creating them, add Usersto roleswhich configure their permissionsforaccessand action with the *Actions* menu to the right of a listed user, or during user creation.

NOTE:

Bestpractice is tospend sometimedesigning yoursystem'ssecurity before creating users, organizations and roles.

By default, every new user has the *Power User* and *User* roles. To assign a new user to specific permissions only, remove all rights on these roles, or confine their permissions to those that are universal first. You can remove users from *Power User*, but not from *User.*

When signed in, you can edit your user information by clicking the link with your username in the top right corner of the screen.

User/Power User Roles

This role's description is *Portal Role: Portal userswith viewaccess.* To turn off most permissions from the User Role, go to Redcell > Permission manager and edit the User role. The *Advanced* button opens a screen where you can select/de-select permissions in larger groups. Power User is *Portal users with extended privileges,* and Administrator is *Portal users with system privileges.*

Default User Roles — Power User

To make new users *not* assigned as Power Users by default, go to the Portal > Portal Settings > Users > Default Associations Taband remove the rolesyou do not wantassigned bydefault. Notice that you can assign/unassign to existing users in th is tab too. The role User appears in this default list, but removal does not have an impact. Cruz automatically assigns all users to the User role, so you must modify it as a universal minimum of permissions.

Multitenancy and Roles

Fora new user that is part of Customer Turnup, Cruz always assigns the Power User role, regardless of defaults, since it is a Site Contact.

Public/Private Page Behavior

Despite the small *Public/Private* label next to the My Private/My Public pages listed in the *Go To* menu, both types of pages appear only for the user(s) who created them. Page Standard settings are *Max Items, Default Filter, Max Items per Page,* and *Column Configuration.* These persist for

Admin users on the Cruz pages, or for users who have the portlet on their Public or Private pages (which makes them the owner of that instance). Without Cruz portlets, URLs for pages labeled public are accessible even to users who do not log in.

Some portlets provide extra settings—for example Alarms portlet's charting options, or the *Top N* portlets number of Top Items. These persist too.

NOTE:

Max Items, *Max Items Per Page* and *Columns* persist for both the summary and maximized portlets independently. For example: If Max Items is 50 in minimized mode it does not affect the Max Items in the Maximized window state. This lets you configure modes independently.

Cruz remembers the default sort column and order per user, whether the user has Admin rights or not. The Sort Column/Order (Descending/Ascending) is also shared between both summary and maximized portlets. A sort on IP Address in Resources persists if you expand the summary portlet to maximized mode.

The *My Public* and *My Private* pages do not appear in sites that have Multitenancy enabled unless you access them through Portal > Sites in the Control Panel. These pages are unrelated to user-specific pages in non-Multitenant installations. They refer to the site, not the user. Refer to the *Installation Guide* for more information about Multitenancy.

In any case, the administrative user can re-arrange pages and portlets in a way that persists. Non-administrative users typically cannot do this.

Portal > Roles

Rolesdetermine the applications permissions available to usersassigned them; manage them in the *Portal > Roles* screen. Notice that these permissions are for the web portal's open source capabilities. You can click the *Actions* button to the right of listed Roles to change a role's portal capabilities. Toconfigure Cruz'sfunctional permissions, overand aboveportal capabilities, use the editor described in Redcell > Permission Manager on page 37.

Click *Add* to create a *Regular Role, Site Role,* or *Organizational Role.* A *Regular Role* assigns its permissions to its members. A *Site* or *Organizational Role* assigns portal permissions to a site or organization to which you can assign users. Other than for *Regular Role*, however, only web portal permissions (not redcell permissions) are available for *Site Roles* and *Organizational Roles.*Only *Regular Role*s restrict a user's Cruz abilities.

Click the *Action* button to the right of a role to *Edit,* view or alter *Permissions, Assign Members* (this last works to see and assign users).*You* can also assign role members in the Portal > Users and Organizations user editor.

NOTE:

Owner Roles do not have an *Action* button. Owner implies something you have added or created and so actions do notapply.

Notice also that when you *Assign Members*, a screen appears with tabs where you can assign *Users, Sites, Organizations* and *User Roles.* Typical best practice is to assign users to one of these collective designations, then assign the collection to a Role.

Notice also that you can view both *Current* and *Available* members with those sub-tabs. You can even *Search* for members.

Click *Back* (in the upper right corner) or the *View All* tab to return to the screen listing roles and their *Action* buttons.

Portal > PasswordPolicies

This panel lets you configure password policies for your installation. It includes options that configure whetherand whenchange to passwordsmustoccur, syntaxchecking, history, expiration, and lockout policies for failed logins.

NOTE:

For users of Multitenant sites, logins require theprepended siteprefix. For example, an adminuser in customer site with prefix BP, logs in as BP-admin.

You can also generate a User Login Report/Last 30 Days to view the history of logins.

Portal > PortalSettings

The *Settings* screens are where users who are administrators can configure the most basic global settings for Cruz, including names, authentication, default user associations, and mail host names. These include the following:

- Users — Among other things *Default User Associations* configures Site and Role defaults for users. Remember, you can remove the Power User default here, but removing the User role does not work. You have to change permissions for User because every user gets that Role.
- Mail Host Name(s) — These are for user account notifications, not mail hosts for Cruz event-based notifications configured in Event Processing Rules, for example (see Automation and Event Processing Rules on page 288). Configure such notifications as described in SMTP Configuration on page 63.
- Email notifications, who sends them, what the contents are for account creation notices, or password change/reset notices.
- Identification, including address, phone, email and websites.
- The default landing page, and display settings like the site logo.
- Google Apps login/password.

CAUTION:

Checking *Allow Strangers to create accounts* may produce a defective login screen. Do not do it.

Portal > [Other]

Some of the remaining portal items permit the following:

Sites— Configure sites. Sites are a set of pages that display content and provide access to specific applications. Sites can have members, which are given exclusive access to specific pages or content. Refer to the *Installation Guide* (or online help) for a more in-depth explanation of Multitenancy and the real power of sites.

Site Template — Configures pages and web content for organizations. Refer to the *Installation Guide* (or online help) for a more in-depth explanation of Multitenancy and the real power of sites.

Page Template—Configures a page and portlets, as well as permissions. Refer to the *Installation Guide* (or online help) for a more in-depth explanation of Multitenancy and the use of Page Templates.

Password Policy—Configure thesecuritypoliciesyouwant, includinguserlockoutand password expiration, and assign them to users. (See the first chapter of this guide for details).

Custom Fields—Lets youconfigure custom fields for Blogentries, Bookmarks or Bookmark Folders, Calendar Events, and so on.

Monitoring —Lets you see details like accessed URLs, number of hits, and so on, for live sessions on the portal. Click a session to see its details. This is usually turned off in production for performance reasons.

Plugins Configuration — Configure access to portlets and features like themes, layouts and so on. By default, only administrators can add portlets/plugins to their pages.

Redcell > Permission Manager

Manage Permissions in these panels to manage user access to different Cruz features. These are configured aspartof Roles, whichaggregate usersregardlessofcommunity affiliation. Create Roles with Portal > Roles.

Notice that, by default, *User*, *Power User* and *Administrator* roles exist and have *redcell* permissions. Also, bydefault, theapplication assigns *User* and *PowerUser* roles' permissions to any new users you create. Since redcell logically ANDs all permissions, this may mean you want to alter the defaults that come with these roles too. You can remove the *Power User* default in the Portal > Portal Settings > User > Default User Associations panel, but removing the *User* role from these default assignments is ineffective, so take care to configure *User* to reflect your system's requirements.

The *Users* editor screen accessible from the *Action* menu for users listed in Portal > Users and Organizations lets you manage groups to which Users are assigned.

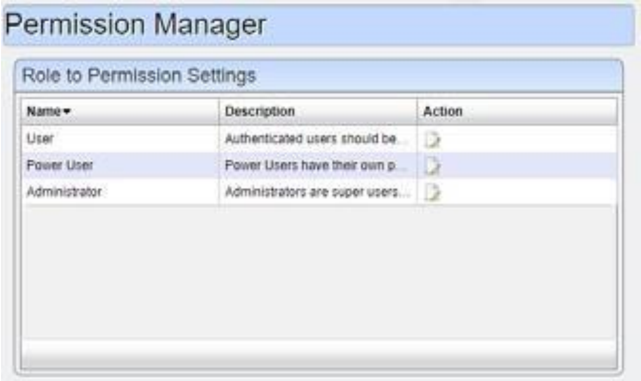

Permission Manager

Role to Permission Settings

Name ▾	Description	Action
User	Authenticated users should be...	
Power User	Power Users have their own p...	
Administrator	Administrators are super users...	

Click the Actions Edit button to see and configure its permissions.

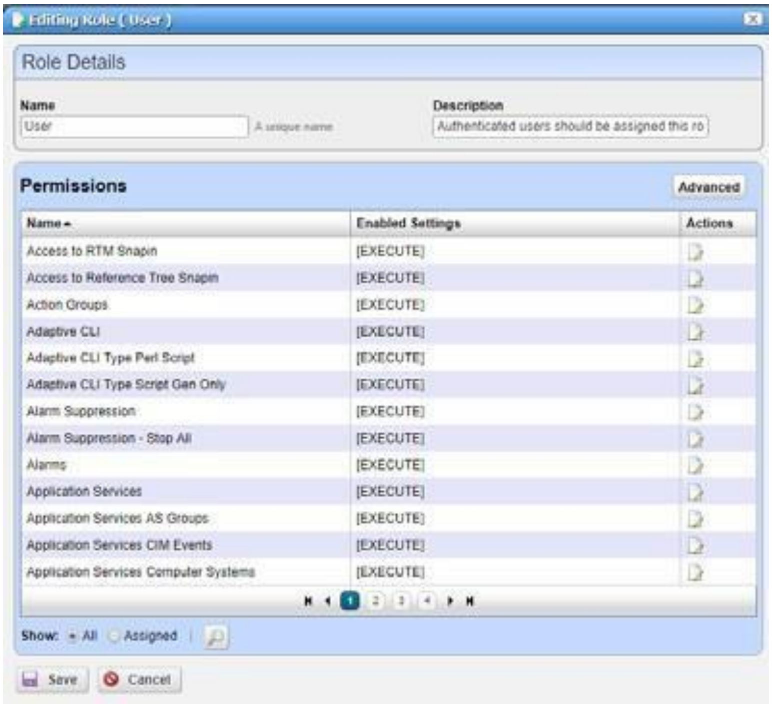

Editing Role (User)

Role Details

Name: User (A unique name)

Description: Authenticated users should be assigned this ro

Permissions

Name	Enabled Settings	Actions
Access to RTM Snapin	[EXECUTE]	
Access to Reference Tree Snapin	[EXECUTE]	
Action Groups	[EXECUTE]	
Adaptive CLI	[EXECUTE]	
Adaptive CLI Type Perl Script	[EXECUTE]	
Adaptive CLI Type Script Gen Only	[EXECUTE]	
Alarm Suppression	[EXECUTE]	
Alarm Suppression - Stop All	[EXECUTE]	
Alarms	[EXECUTE]	
Application Services	[EXECUTE]	
Application Services AS Groups	[EXECUTE]	
Application Services CIM Events	[EXECUTE]	
Application Services Computer Systems	[EXECUTE]	

Show: ● All ○ Assigned

Notice that you have the option to view *Assigned* or *All* permissions or search for the permission you want to locate.

Click Actions *Edit* button to modify the type of permissions available.

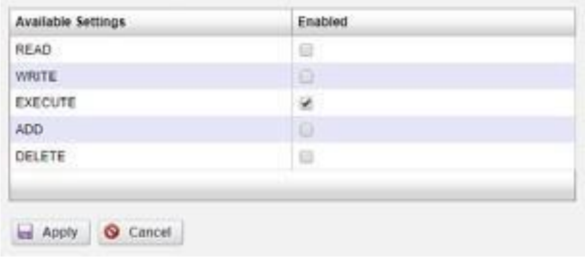

Available Settings	Enabled
READ	□
WRITE	□
EXECUTE	☑
ADD	□
DELETE	□

Thefollowingdescribes the actionsofthe permissions, when selected:

Action	Default Behavior
read	Enables *Details*, *Topology* and *View as PDF*
write	Enables the *Edit, Save,* and *Import/Export.*
execute	Let's you see the view altogether, launch from a portlet and query for elements. Alternatively, this action can control a specific application function, (typically described by the permission name) like provisioning a policy.
add	Enables the *New* menu item and *Save*. If you do not check this action, then the *New* menu item does not appear.
delete	Enables the *Delete* menu item.

The *Advanced* button lets you configure each functional permissions bytype. Forexample, expand the Read list to specify READ permissions to all or selected functions.

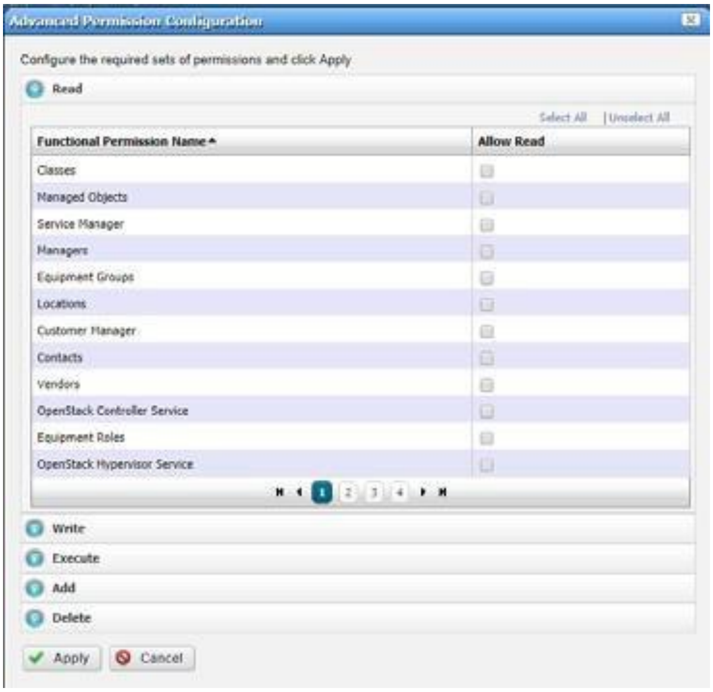

Advanced Permission Configuration

Configure the required sets of permissions and click Apply

Read

Select All | Unselect All

Functional Permission Name ▲	Allow Read
Classes	☐
Managed Objects	☐
Service Manager	☐
Managers	☐
Equipment Groups	☐
Locations	☐
Customer Manager	☐
Contacts	☐
Vendors	☐
OpenStack Controller Service	☐
Equipment Roles	☐
OpenStack Hypervisor Service	☐

When you hover the cursor over a functional permission, tooltips provide a description.

NOTE:

If you upgrade your installation and new permissions are available, edit the Administrator Role, and notice an enabled *Add* button indicates new permissions are available. Because of an upgrade, for example, theConfiguration Filesportletmightnotbevisible.Bydefault, upgrades turnoffanynew permissions, so if you want them enabled, particularly for Administrators, click *Add* and enable them for the Roles for which you want them enabled.

The portlet settings button (wrench) through which you edit portlet settings, is only accessible if you have the "Change Portlet Settings" permission. On upgrade, no users this permission. This is by design, because the intent is for the system administrator to consider which users (more specifically, which user roles) should change portlet settings and which users should not have access to this feature.

Redcell > Database Aging Policies (DAP)

Database Aging Policies prevent the Cruz database from filling up by filling up by deleting old records. You can also save designated contents to an archive file on a specified cycle. Database Aging Policies configure which contents to archive, the archive location, and the configuration of that archive file.

To view and manage such policies, right click an item with them (for example, an alarm), or click *Manage > Control Panel*, and under *Redcell* click *Database Aging Policies.*

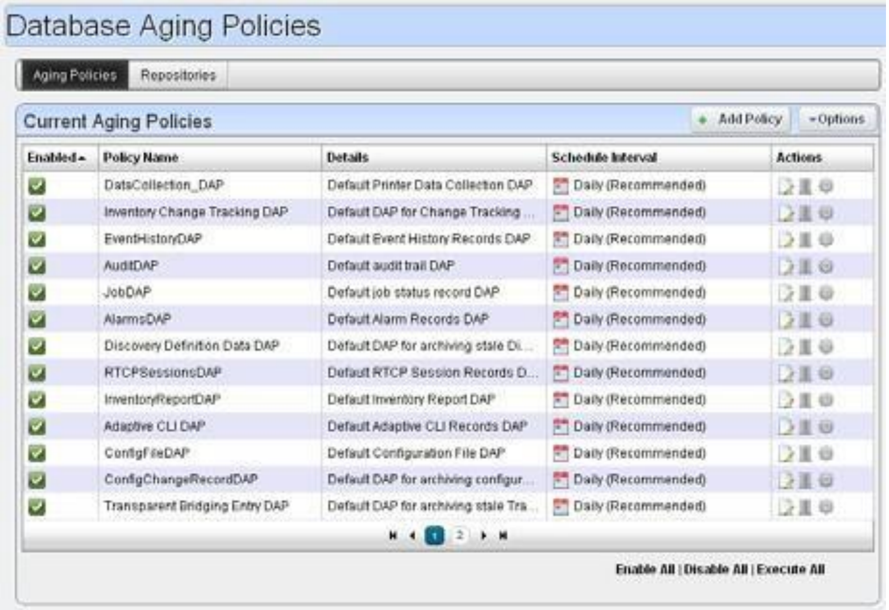

Enabled	Policy Name	Details	Schedule Interval	Actions
☑	DataCollection_DAP	Default Printer Data Collection DAP	Daily (Recommended)	
☑	Inventory Change Tracking DAP	Default DAP for Change Tracking ...	Daily (Recommended)	
☑	EventHistoryDAP	Default Event History Records DAP	Daily (Recommended)	
☑	AuditDAP	Default audit trail DAP	Daily (Recommended)	
☑	JobDAP	Default job status record DAP	Daily (Recommended)	
☑	AlarmsDAP	Default Alarm Records DAP	Daily (Recommended)	
☑	Discovery Definition Data DAP	Default DAP for archiving stale Di...	Daily (Recommended)	
☑	RTCPSessionsDAP	Default RTCP Session Records D...	Daily (Recommended)	
☑	InventoryReportDAP	Default Inventory Report DAP	Daily (Recommended)	
☑	Adaptive CLI DAP	Default Adaptive CLI Records DAP	Daily (Recommended)	
☑	ConfigFileDAP	Default Configuration File DAP	Daily (Recommended)	
☑	ConfigChangeRecordDAP	Default DAP for archiving configur...	Daily (Recommended)	
☑	Transparent Bridging Entry DAP	Default DAP for archiving stale Tra...	Daily (Recommended)	

Policies appear in the *Aging Policies* tab of this screen, with columns that indicate whether the policy is *Enabled*, the *Policy Name, Details* (description), *Scheduled Intervals* and icons triggering three *Actions* (*Edit*, *Delete* and *Execute*). Notice that the bottom right corner of this page also lets you *Enable/Disable/Execute All* policies listed.

Redcell > DataConfiguration

This panel configures custom attributes for Cruz. Click the *Edit* button next to the *Entity Type* (Managed Equipment, Port, Contact, Vendor, or Location) for which you want to create custom attributes. This opens an editor listing the available custom attributes for the entity type. Editing Custom Attributes on page 146 describes right clicking to access this directly from the portlet menu, and the details of how to edit custom attributes.

NOTE:

The custom fields configured here are for Cruz. only. The Custom Fields editor in the *Portal* portion of Control Panel manages custom fields for the rest of the portal.

Redcell > Audit Trail Definitions

This screen, accessible from Go to > Control Panel lets you manage audit trail definitions in Redcell. Audit trail entries are based on these definitions. Clicking the Edit icon to the right of a listed filter opens the editor. From this screen you can control the behavior of audit trail entries based on each respective definition.

The following fields are available for configuration:

Severity Level - Users viewing the audit trail must have at least this level of security in order to view

this type of audit trail. 0 is the most restrictive, so if a definition has a security level of 0 then only users with the highest level of security can see this type of audit trail.

Disabled - Controls whether this type of audit trail is saved to the database. If disabled is checked, then audit trail entries of this type are not saved.

Emit Event - Controls whether to emit an event when this type of audit trail is created. This can be useful, for example, if you want to forward audit log entries to a northbound system. To do this, simply check Emit Event for the type of audit trails that you want to forward and then create an Event Processing Rule to forward northbound all events of type redcellAuditTrailEntry. Note that the default behavior for the event definition redcellAuditTrailEntry is reject, meaning that these events will not be saved to the event history, but they can still nonetheless trigger the execution of event processing rules.

Redcell > Mediation

Thispanel monitors mediation serversin your system, appearingonlywhen such serversexist. Mediation servers appear listed in the *Servers* tab of this manager if mediation servers are connected to application servers.

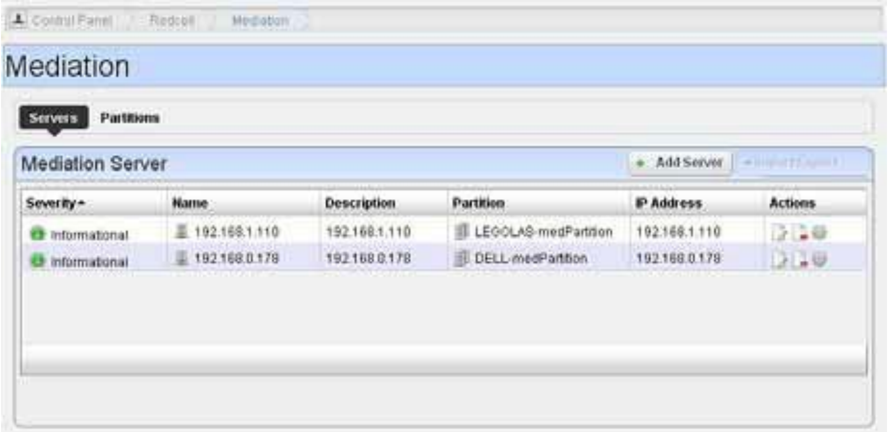

Severity	Name	Description	Partition	IP Address	Actions
Informational	192.168.1.110	192.168.1.110	LEGOLAS-medPartition	192.168.1.110	
Informational	192.168.0.178	192.168.0.178	DELL-medPartition	192.168.0.178	

Mediationserver, routingentriesand partition entriesappearautomatically when mediation server connects for the first time. You can test connectivity from appserver cluster and medserver/ partition.

You can export or import both server and partition configurations. Use the button on the right above thelisted serversorpartitions to do this. Importing Partitions/MedServersoverwritesthosein the database with the same names. Exporting a partition exports contained medservers too. Importing a partition looks for overlapping routing entries and saves the partition with only its unique entries. If no entries are unique, the partition is not saved.

NOTE:

Thispaneldoesnotappearifyouinstall Cruzinstand-alonemode, withoutaseparatemediationserver. To make it appear, add `medserver.support=true` to `portal-ext.properties` file in `\oware\synergy\tomcat-7.0.26\webapps\ROOT\WEB-INF\classes`. Remember, best practice is to override properties as described in Overriding Properties on page 112.

In addition to automatically detecting mediation servers, you can click *Add Server* to configure additional mediation servers.

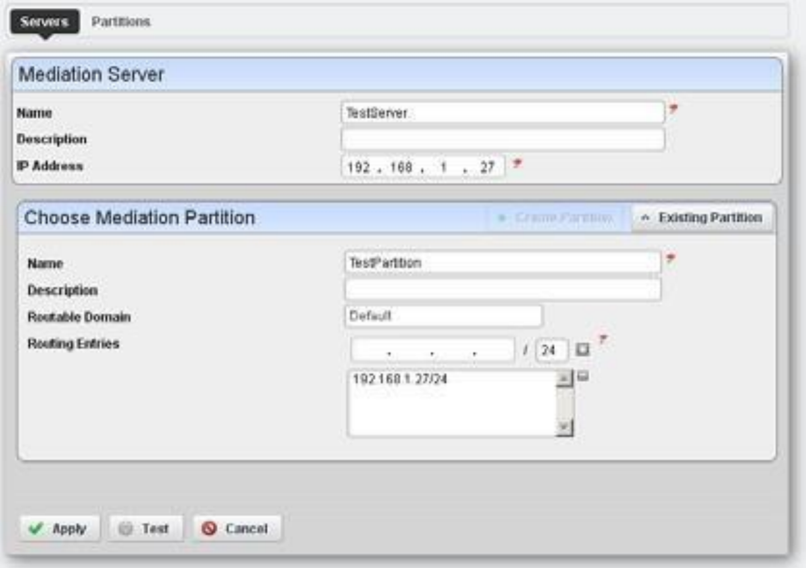

Mediation Server

Name	TestServer
Description	
IP Address	192 . 168 . 1 . 27

Choose Mediation Partition

Name	TestPartition
Description	
Routable Domain	Default
Routing Entries	. . . / 24 192.168.1.27/24

When creating a new server, enter a Name, Description and IP Address. You can also Create Partition (or select from Existing Partitions), choosing a Name, Description, Routable Domain, and Routing Entries (click the ‘ +’ to add your entries to the list).

The *Test* button scanning the ports in the proposed application server/mediation server link, validating the installed versions of Cruz in both locations are the same, and validating the connection between application server and mediation server. A window opens similar to the Audit Trail/Job Status on page 148 that tracks testing progress.

The *Partitions* tab of the Mediation monitor displays already-configured partitions, and let’s *Add Partition*, or lets you edit them with an *Edit this entry* icon. The editor screen is like the one that adds new partitions, where you enter a name, description and routable domain CIDR IP addresses.

Test listed partitions with the gear icon to the right of the partition or delete it with the *Delete this entry* icon. Notice that you can also *Import/Export* partition descriptions with that button on this screen.

Search for Mediation Server

The *Search* button in the *Partitions* tab of the Mediation monitor opens a screen where you can enter an address in *IP to Search for*.

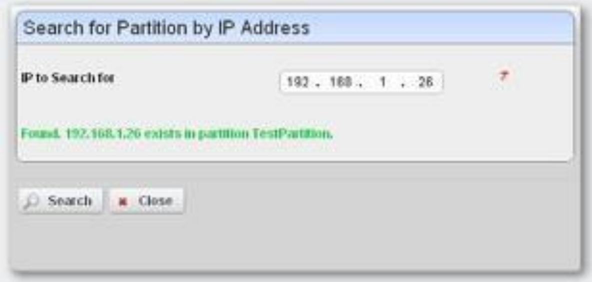

Search for Partition by IP Address

IP to Search for: 192 . 168 . 1 . 26

Found. 192.168.1.26 exists in partition TestPartition.

Clicking *Search* locates the mediation partition that serv ices the entered IP address (although it does not determine whether that partition is up and running).

Redcell > FilterManagement

This screen, accessible from *Go to > Control Panel* lets you manage the filters in Cruz.

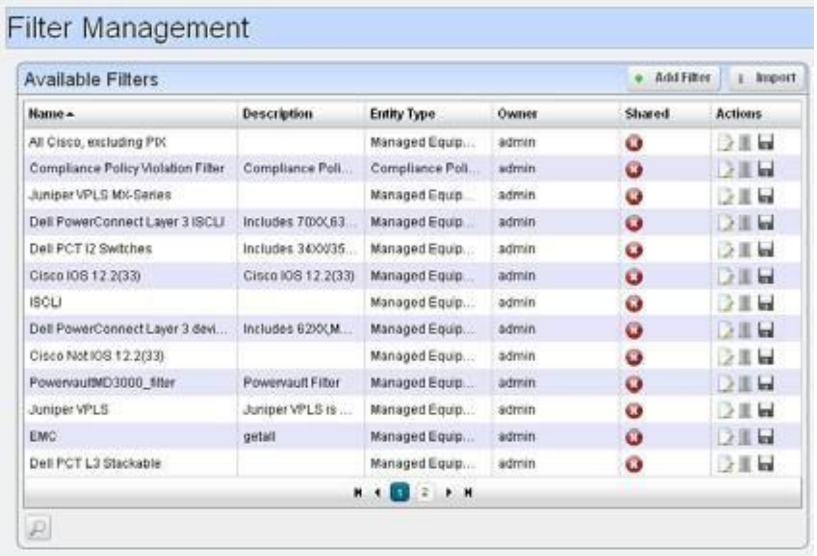

Filter Management

Available Filters

Name ▲	Description	Entity Type	Owner	Shared	Actions
All Cisco, excluding PIX		Managed Equip...	admin	✖	
Compliance Policy Violation Filter	Compliance Poli...	Compliance Poli...	admin	✖	
Juniper VPLS MX-Series		Managed Equip...	admin	✖	
Dell PowerConnect Layer 3 ISCLI	Includes 70XX,63...	Managed Equip...	admin	✖	
Dell PCT l2 Switches	Includes 34XX/35...	Managed Equip...	admin	✖	
Cisco IOS 12.2(33)	Cisco IOS 12.2(33)	Managed Equip...	admin	✖	
ISCLI		Managed Equip...	admin	✖	
Dell PowerConnect Layer 3 devi...	Includes 62XX,M...	Managed Equip...	admin	✖	
Cisco Not IOS 12.2(33)		Managed Equip...	admin	✖	
PowervaultMD3000_filter	Powervault Filter	Managed Equip...	admin	✖	
Juniper VPLS	Juniper VPLS is ...	Managed Equip...	admin	✖	
EMC	getall	Managed Equip...	admin	✖	
Dell PCT L3 Stackable		Managed Equip...	admin	✖	

Click the *Delete* icon to the right of a listed filter to remove it from the system. Click the disk icon to export the filter. Clicking the *Import* button at the top of the screen lets you import previously exported filters.

NOTE:

To find a filter, click the *Search* (magnifying glass) icon in the lower left corner of this screen.

Clicking the *Edit* icon to the right of a listed filter or clicking the *Add Filter* button opens the filter editor.

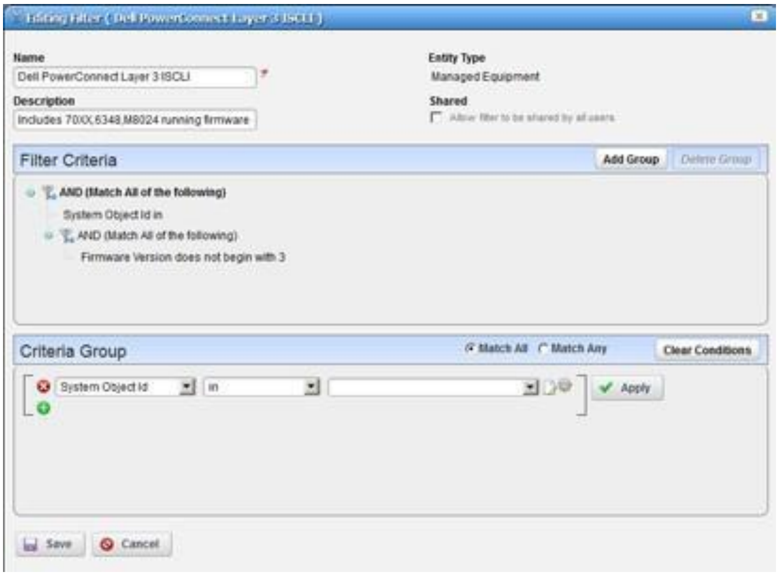

Editing Filter (Dell PowerConnect Layer 3 ISCLI)

Name: Dell PowerConnect Layer 3 ISCLI *

Description: Includes 70XX,6348,M8024 running firmware

Entity Type: Managed Equipment

Shared: □ Allow filter to be shared by all users

Filter Criteria

- AND (Match All of the following)
  - System Object Id in
  - AND (Match All of the following)
    - Firmware Version does not begin with 3

Criteria Group

● Match All ○ Match Any

System Object Id | in |

Use this editor to configure filters. Enter a *Name* and *Description*, and use the green plus (+) to select an entity type from a subsequent screen. Checking *Shared* makes the filter available for all users, not just your user. You can add groups of filter criteria (click *Add Group*) that logical AND (*Match All*) OR (*Match Any*) with each other. Click *Clear Conditions* to remove criteria. Configure the filter in the *Criteria Group* panel as described in the How to: Defining Advanced Filters on page 145. Delete filters with the *Delete this entry* icon next to the edit icon.

Redcell > Application Settings

This screen has several panels in the following tabs:

- General > Entity Change Settings
- User Interface > Map Provider
- User Interface > Job Viewer
- User Interface > Performance Chart Settings
- User Interface > Revert to Factory Defaults (self-explanatory)
- Server

General > Entity Change Settings

This panel lets you override polling/refreshing for the minimized Managed Resources, Alarms, Hierarchical View Tree, Topology and Map Context portlets. By default, these portlets poll at 40, 35, 40, 40, and 60 seconds, respectively, for changes in the data and automatically refresh. Polling times are configurable. The valid range is 10 seconds -> 3600 (1 hour) for the minimized Alarms portlet, 20 seconds - > 3600 seconds for the others.

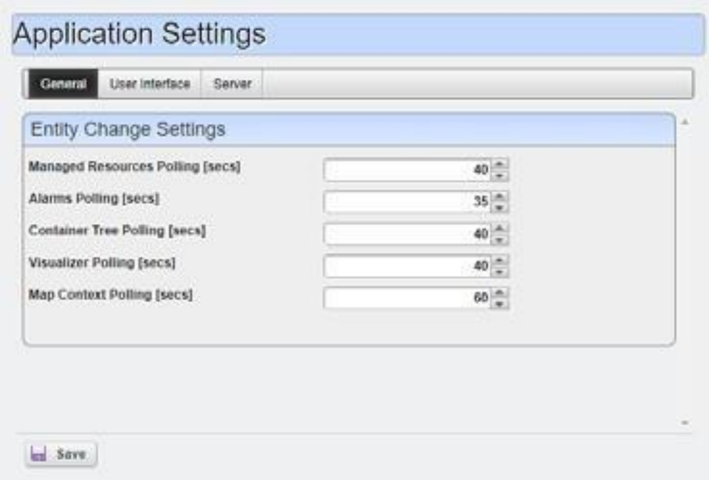

Managed Resources Polling [secs]	40
Alarms Polling [secs]	35
Container Tree Polling [secs]	40
Visualizer Polling [secs]	40
Map Context Polling [secs]	60

User Interface > Map Provider

The *Map Provider* panel lets you set whether Cruz uses Google or Nokia maps by default and sets the Initial Latitude and Longitude. Check *Use Secure API* if you want to load map APIs in secure SSL mode. Some browsers block non-secure external APIs if they are viewing a secure page, so use this if you view Cruz through an HTTPS connection.

Follow the directions in Setting Up Google Maps on page 249 to set the application to use those maps.

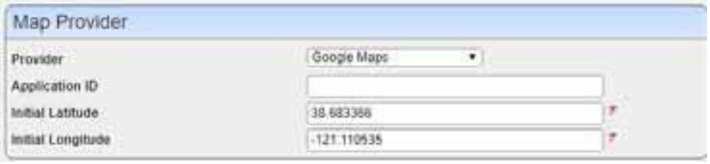

Map Provider

Provider	Google Maps ▾
Application ID	
Initial Latitude	38.683386
Initial Longitude	-121.110535

User Interface > Job Viewer

The *Job Viewer* panel lets you select the following options and set the popup width:

Show Job Viewer—Checking this displays the job viewer after Execution (most cases). Leaving it unchecked does not displayit, although you can still view jobs with *My Alerts* in the lower left portion of the screen.

Alwaysshow Job Viewerfor Actions—Whenchecked, thisdisplaysthejobviewerforexecutionof Actions or Action Groups.

Show Information Messages by Default —When checked, showsinformational message nodesby default.

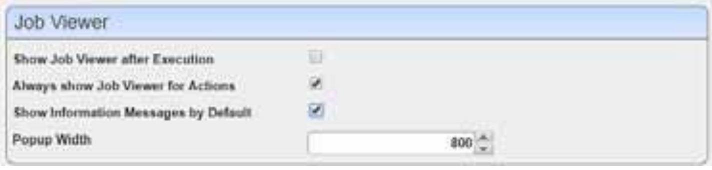

Job Viewer

Show Job Viewer after Execution	☐
Always show Job Viewer for Actions	☑
Show Information Messages by Default	☑
Popup Width	800

User Interface > Performance Chart Settings

This panel displays options for the performance dashboard and traffic charts. Day and Minute Formats depend on the locale settings in the operating system running Cruz. Select them in the pick lists that appear in this panel. The 24 Hr Clock setting makes the charts show times in 24-hour format.

The Canvas Line Charts option controls the type of line charts that are used. Earlier versions of Redcell used a Scalable Vector Graphics type line chart. Redcell now supports a Canvas based line chart which can display many more points. If you prefer to use the old style SVG line charts you can uncheck this box.

The Restrict Y-Axis Range to data range option causes the Y-axis range to be based on the range of data being graphed. If this is not checked the range will be based on the max and min values associated with the attribute definitions. For example, most % values will have a range of 0 to 100.

The Equipment Display setting controls how equipment labels are shown in the performance dashboards. The default is to display the IP address. You can also have it display the device name. The Port/Interface Display setting lets you select between showing the Port/Interface Name of the Port/Interface description.

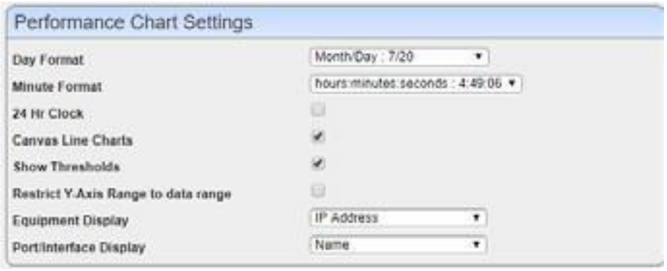

Performance Chart Settings

Day Format	Month/Day : 7/20 ▾
Minute Format	hours:minutes:seconds : 4:49:06 ▾
24 Hr Clock	□
Canvas Line Charts	☑
Show Thresholds	☑
Restrict Y-Axis Range to data range	□
Equipment Display	IP Address ▾
Port/Interface Display	Name ▾

Server

Thispanel lets you override the system defaults for web client time-outs. You must configure this if you want your web client to remain connected to the server for extended periods. The timeout extends automatically with any activity (keystrokes) on the client.

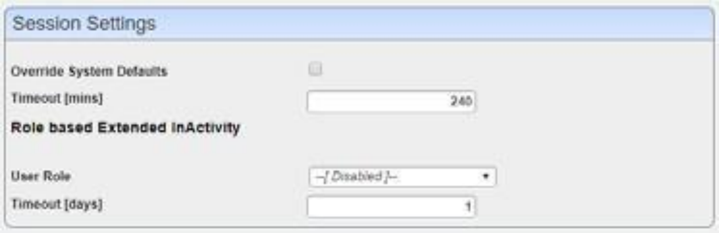

Session Settings

Override System Defaults	□
Timeout [mins]	240
Role based Extended InActivity	
User Role	--[Disabled]-- ▾
Timeout [days]	1

Set the timeout with the following fields:

Override System Defaults—Check this to start overriding the system default time-outs for all users.

Timeout [mins] — Enter the minutes of inactivity before Timeout occurs (5 - 2880 [two days])

Role based Extended Inactivity

If you want to confine extended web client session s to a particular role, then use these fields to configure that.

User Role — Select the role for which this override works from the pick list.

Timeout [days] — Select the number of days for role-based timeout override, up to 365.

Server Administration

The Server Administration option lets you manage the portal's webserver and maintain its smooth operation in a variety of ways. Click the Execute buttons to run actions, such as re-indexing the search indexes. These options are visible to administrators only and they contain helpful settings and resource information related to the server.

Access this option by selecting Go to > Control Panel > Server > Server Administration.

Aging Policies Editor

When you click *Add Policy*, a selector appears where select the kind ofpolicyyou want to create and the then the editor appears. If you click the *Edit* icon to the right of a listed policy, the Aging Policies Editor appears with that policy's in formation already filled out, ready to modify.

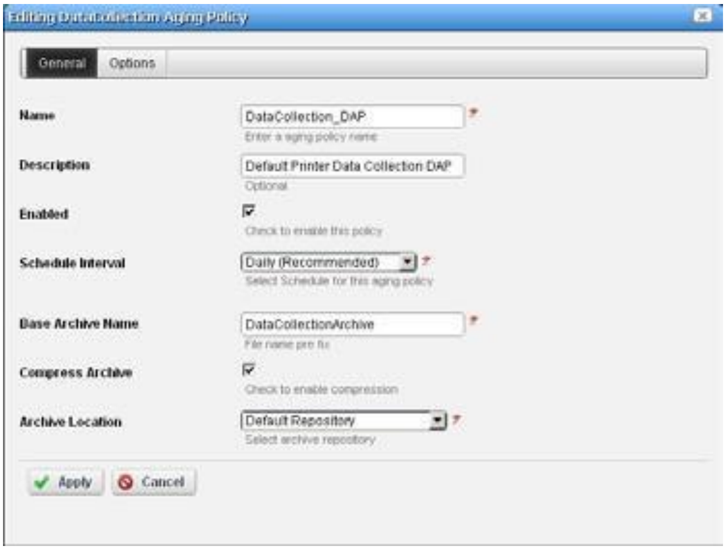

Editing DataCollection Aging Policy

Name: DataCollection_DAP *
Enter a aging policy name

Description: Default Printer Data Collection DAP
Optional

Enabled: ☑
Check to enable this policy

Schedule Interval: Daily (Recommended) *
Select Schedule for this aging policy

Base Archive Name: DataCollectionArchive *
File name pre fix

Compress Archive: ☑
Check to enable compression

Archive Location: Default Repository *
Select archive repository

The *General* screen has the following fields:

Name — An identifier for the policy

Description—A text description of the policy

Enabled — Check to enable the policy.

Schedule Interval—Use the pick list to select an interval. Once you have configured an interval here, you can re-configure it in the Schedules portlet.

Base Archive Name — The prefix for the archived file.

Compress Archive — Check to compress the archive file.

Archive Location — Select from the available Repositories in the pick list.

The contents of the *Options* tab depend on the type of DAP you are configuring. Typically, this tab is where you set the retention thresholds.

DAP SubPolicies

Some Options tabs include sub-policies for individual attribute retention.

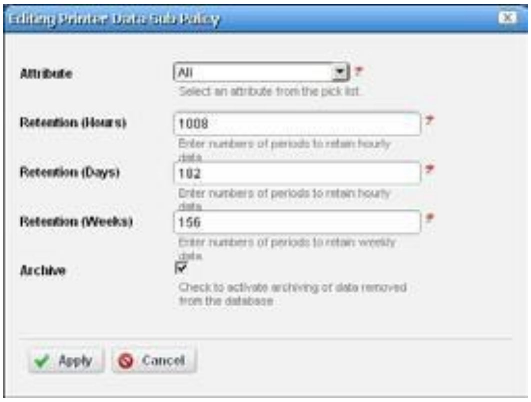

Attribute: All
Select an attribute from the pick list.

Retention (Hours): 1008
Enter numbers of periods to retain hourly data.

Retention (Days): 182
Enter numbers of periods to retain hourly data.

Retention (Weeks): 156
Enter numbers of periods to retain weekly data.

Archive: ☑
Check to activate archiving of data removed from the database

Click *Add SubPolicy* or click the *Edit* button to the right of listed policies to access the editor. The fields vary depending on the policy selected.

Editing Tips

Archiving options that appear in the Aging Policies Editor vary, based on type of policy selected. Inventory Change Tracking DAPs ask how long you would like to keep Config reports, Inventory Report DAPs ask howlong you would like to keep your Historical Reports based on number of instances, days, and weeks, monthsor years.

Set these thresholds in the *Options* tab. All DAPs require a Name and a record threshold. Check the *Enabled* checkbox to enable the policy.

DAPs run on a schedule. If the record threshold number is greater than or equal to the configured threshold then the DAP runs at the scheduled time. You may also manually click the gear icon to the right of a listed policy and execute a DAP at any time to check that threshold figure. In either case, if the threshold is not crossed Cruz creates no archives.

To verify when current DAPs are scheduled to run, open the Schedules portlet, and select the schedule on which it runs. For most DAPs, this is the Daily (recommended) DAP. Right-click to edit it. The Scheduled Aging Policies list should include all DAPs that have selected that schedule.

Aging Policies Options

The *Options* tab in this editor can vary, depending on the type of policy.

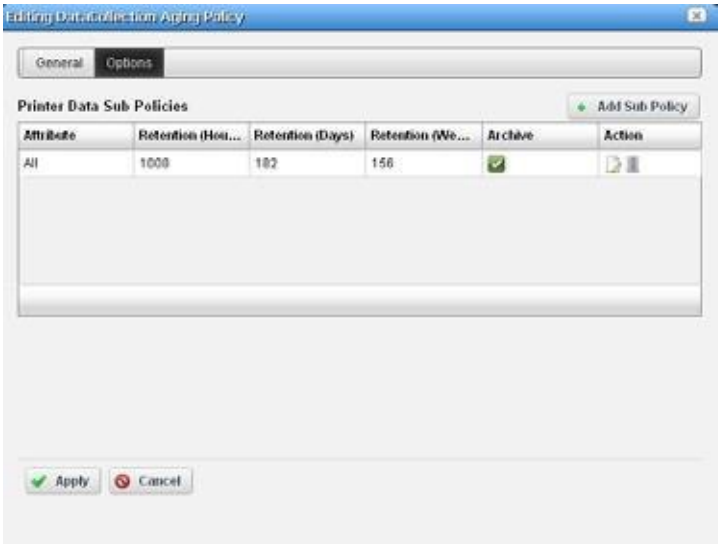

Printer Data Sub Policies

Attribute	Retention (Hou...	Retention (Days)	Retention (We...	Archive	Action
All	1000	182	156	☑	

Fields can include the following:

Keep [Aged Item] for these many days—The number of days to keep the aged item before archiving it.

Archive [Aged Item] — Check this to activated archiving according to this policy.

Sub-Policies

Some types of Database Aging Policies can have su b-policies that further refine the aging for their type of contents.

These appear listed in the Aging Policies Options tab. Click *Add Sub Policy* to create them. Notice that you can *Edit* or *Delete* listed policies with the icons in the far-right *Action* column in this list.

Such sub-policies can contain the following types of fields:

Component — Select the component for the sub-policy from the pick list.

Action Type — This further sub-classifies the *Component*.

Retention (Days) — The number of days to keep the aged item before archiving it.

Archive — Check this to activated archiving according to this policy.

Repositories

When you select a repository in the Aging Policies Editor, the available policies come from what is configured in this tab of the editor.

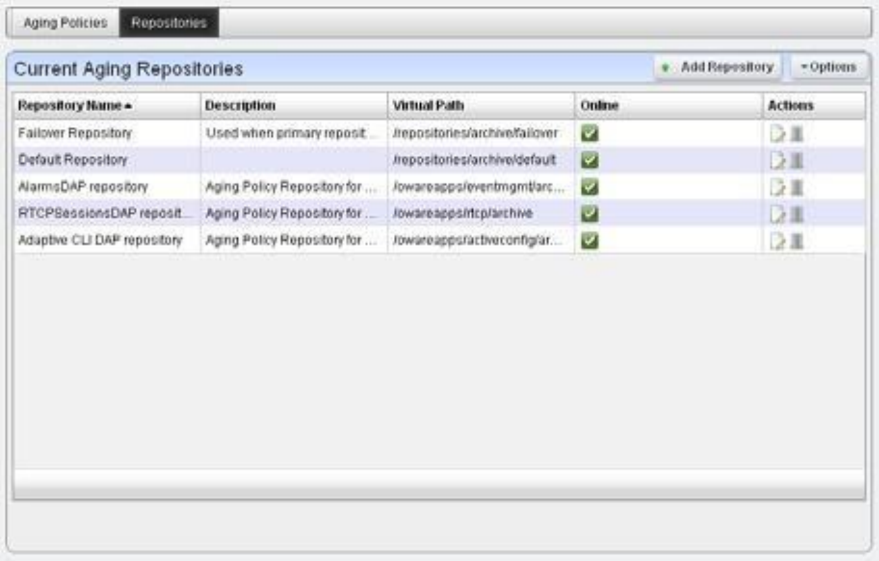

Current Aging Repositories

Repository Name ▲	Description	Virtual Path	Online	Actions
Failover Repository	Used when primary reposit...	/repositories/archive/failover	☑	
Default Repository		/repositories/archive/default	☑	
AlarmsDAP repository	Aging Policy Repository for ...	/owareapps/eventmgmt/arc...	☑	
RTCPSessionsDAP reposit...	Aging Policy Repository for ...	/owareapps/rtcp/archive	☑	
Adaptive CLI DAP repository	Aging Policy Repository for ...	/owareapps/activeconfig/ar...	☑	

Available repositories appear listed in the initial screen. Like the Aging Policies Editor, you can click *Add Repository* to create a new repository, and *Edit* or *Delete* selected, listed policies with the icons in the *Action* column. Notice the listed policies indicated whether the archiving destination is *Online* with a green icon (this is red, when the destination is offline).

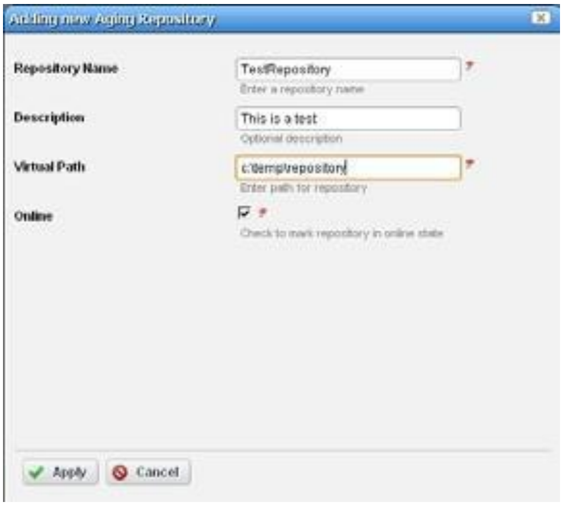

Adding new Aging Repository

Repository Name: TestRepository
Enter a repository name

Description: This is a test
Optional description

Virtual Path: c:\temp\repository
Enter path for repository

Online: ☑
Check to mark repository in online state

When you *Add Repository* or *Edit* an existing one, the following fields appear in the editor:

Repository Name — An identifier for the archiving destination.

Description — A text comment.

Virtual Path—This is the path relative to the installation root directory. Any user with administrator permissions can specify or change the default archive path here.

Online — Check this to put this repository online.

Cruz automatically writes to any configured failover repository if the primary repository is full or not writable.

NOTE:

To view any archived DAP file, use `dapviewer`. Type `oware` in a command shell, then, after pressing [Enter], type `dapviewer` to use this utility.

Quick Navigation Portlet

By default, the Quick Navigation portlet is available from the Home page. Admin users and Power users can see all tasks listed. The User role sees only options to which the user has write permissions.

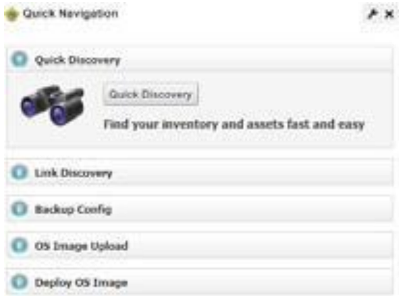

Find your inventory and assets fast and easy

The Quick Navigation portlet lets you quickly perform these basic tasks:

Task	Description
Quick Discovery	Discovers devices in you network with the Quick Discovery defaults or construct a Quick Discovery profile if not exists. See Discovering Resources on page 211 for details.
Link Discovery	Discovers their connections to resources after you discovered them. See Link Discovery on page 188. This option is visible only if you have write permission.
Backup Config Files	Backs up discovered devices’ configuration files. Before you can use this feature, you must have servers configured as described in File Servers on page 458. See also Configuration File Compare Window on page 449.
OS Image Upload	Uploads firmware updates for devices. See Firmware Image Editor on page 454 for more about these capabilities. This option is visible only if you have write permission.
Deploy OS Image	Deploys firmware updates. To deploy images, you must have File Servers configured, as described above for Backup. See Deploying Firmware on page 465.

Network Tools

The Network Toolsportlet lets youinvokeavariety ofexisting functions on a device without having the device currently discovered. It typically appears listed as an available Application to install as a portlet.

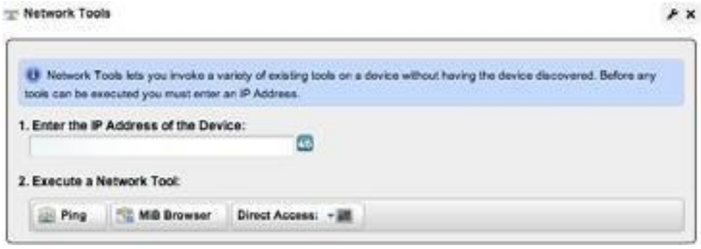

Network Tools lets you invoke a variety of existing tools on a device without having the device discovered. Before any tools can be executed you must enter an IP Address.

1. Enter the IP Address of the Device:

2. Execute a Network Tool:

Before you can use the tools, you must enter an IP address in the appropriate field. Once you have entered that address, you can use the following:

- Ping Tool
- MIB Browser Tool
- Direct Access Tool

NOTE:

If you want to restrict access for some users so they do not automatically log in with direct access, then remove direct access permissions for users, and use Network Tools for direct access.

Ping Tool

The second button is the Ping tool, which pings the selected device, and lists the time for ping response.

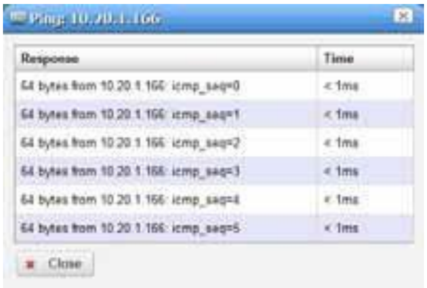

Response	Time
64 bytes from 10.20.1.166: icmp_seq=0	< 1ms
64 bytes from 10.20.1.166: icmp_seq=1	< 1ms
64 bytes from 10.20.1.166: icmp_seq=2	< 1ms
64 bytes from 10.20.1.166: icmp_seq=3	< 1ms
64 bytes from 10.20.1.166: icmp_seq=4	< 1ms
64 bytes from 10.20.1.166: icmp_seq=5	< 1ms

MIB Browser Tool

The first button displays the MIB browser with default SNMP settings. Youcan *Edit* the settings to match the device's SNMP community settings and then save them. The next time Network Tools invokes the MIB browser, it defaults to your previoussettings.

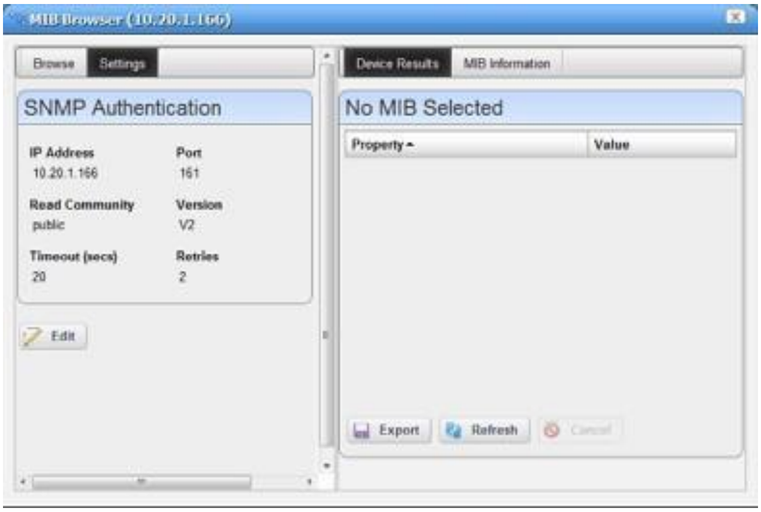

IP Address: 10.20.1.166
Port: 161
Read Community: public
Version: V2
Timeout (secs): 20
Retries: 2

Once you are done editing the SNMP settings, click *Save*. Click the *Browse* tab to look through available MIBs as you would ordinarily do in MIB browser. See MIB Browser Tool on page 53 for more about using the MIB browser. Youcan also browse MIBs in the attribute selection panel for the SNMP monitor. See SNMP on page 405.

NOTE:

Locations for MIB file included with your package are subject to change without notice, but generally are under the owareapps/[application name]/mibs directory for different application modules. Refer to the *Installation Guide* for additional information.

Direct Access Tool

The third button on the Network Tools portlet toolbar opens the Direct Access tool. It provides a command line interface terminal for Telnet, SSH, and SSH V2 access to the device.

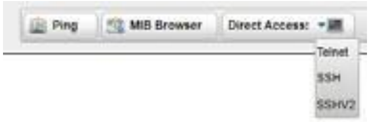

Click and select the type of direct access you want.

- Direct Access - Telnet
- Direct Access - SSH/SSH V2

Direct Access - Telnet

Telnet direct access connects to the device with telnet and displays the terminal session. You must login to the device manually, unlike the method described for discovered devices in Network Tools on page 52.

Direct Access - SSH/SSH V2

Direct Access for SSH or SSH V2 first prompts for a user name and password.

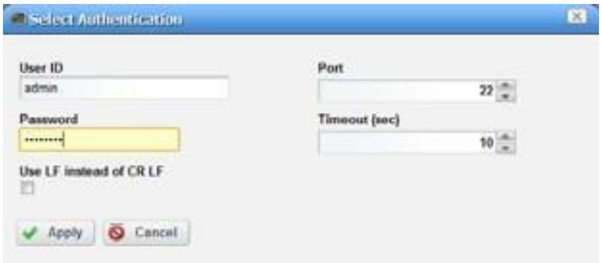

User ID: admin

Password: ••••••••

Port: 22

Timeout (sec): 10

Use LF instead of CR LF: □

The *Use LF instead of CR LF* checkbox suppresses carriage returns when you click Enter key. This is necessary for some devices (for example: some Dell Power Connect devices).

Once you log in, Cruz attempts to connect with SSH or SSH V2 using the user id and password provided. Some Dell Power Connect devices do not log when connected and prompt you to enter the user and password again.

Firefox Browser Configuration for Direct Access

To make Direct Access - Terminal tool for CLI cut-through to work on the Firefox browser; you must install or update the latest version of Java that is compatible with the browser from the following page.

https://www.java.com/en/download/manual.jsp

Once installed, go to Firefox plugins page (about:addons) and enable the Java plugin.

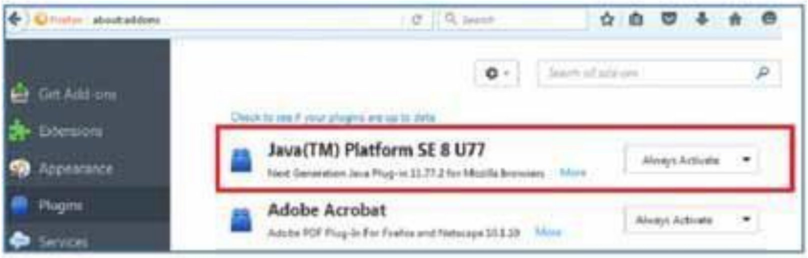

Java(TM) Platform SE 8 U77 — Next Generation Java Plug-in 11.77.2 for Mozilla browsers — Always Activate

Adobe Acrobat — Adobe PDF Plug-In For Firefox and Netscape 10.1.16 — Always Activate

You must enable Java content in the browse r from your Java Security configuration.

Control Panel > Java > Security tab > select "Enable Java content in the browser" option.

You might also have to add the Cruz server IP Address to the Java Security Exception Site list.

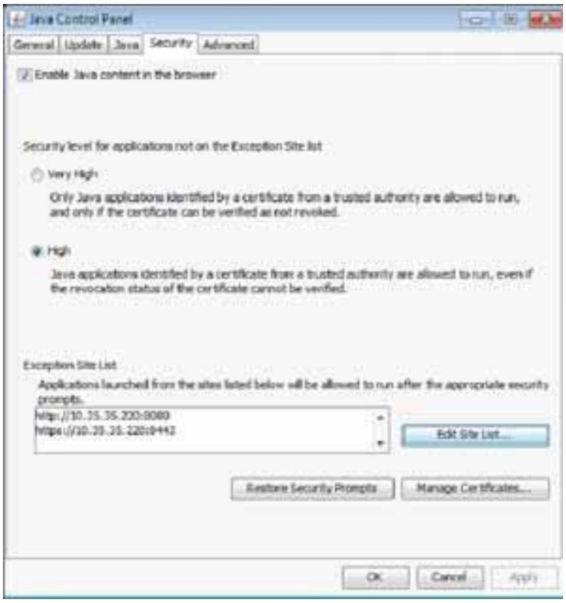

☑ Enable Java content in the browser

Security level for applications not on the Exception Site list

○ Very High

Only Java applications identified by a certificate from a trusted authority are allowed to run, and only if the certificate can be verified as not revoked.

● High

Java applications identified by a certificate from a trusted authority are allowed to run, even if the revocation status of the certificate cannot be verified.

Exception Site List

Applications launched from the sites listed below will be allowed to run after the appropriate security prompts.

```
http://10.35.35.220:8080
https://10.35.35.220:8443
```

License Viewer Window

Use the License Viewer window to examine the licensesfor your system’s capabilities, monitor license expirations, and register a license.

Access the License Viewer window by clicking the License Management button from the Settings > Application Configuration Settings portlet.

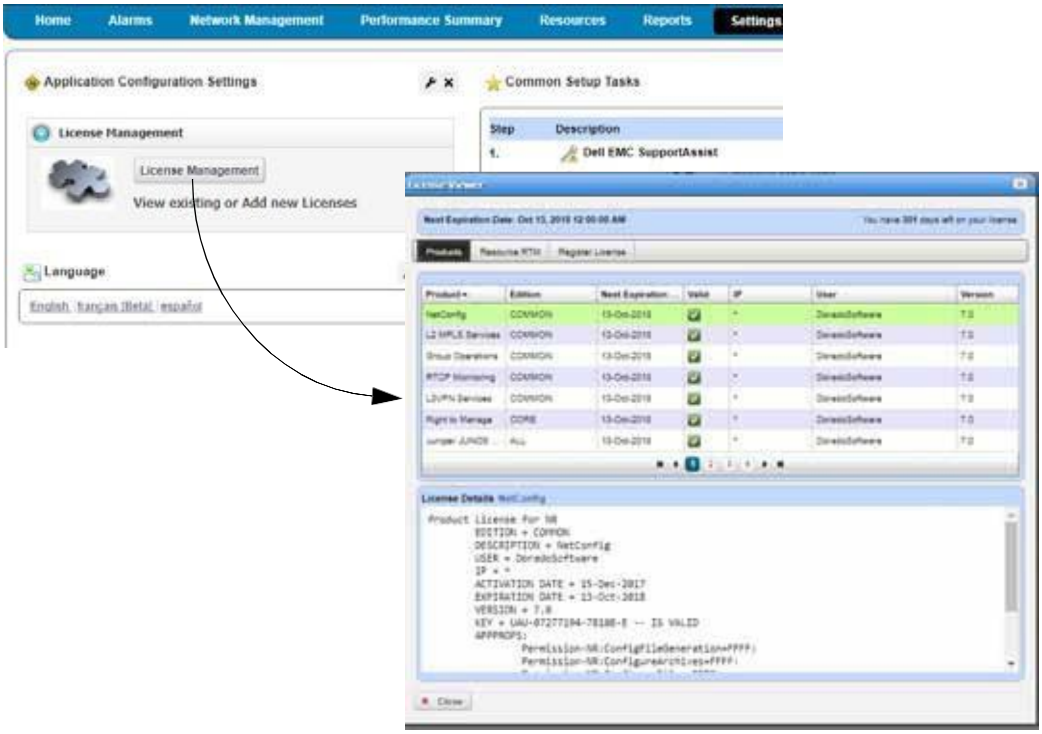

Next Expiration Date: Oct 13, 2018 12:00:00 AM

You have 301 days left on your license

Product	Edition	Next Expiration	Valid	IP	User	Version
NetConfig	COMMON	13-Oct-2018	✓	*	DoradoSoftware	7.0
L2 MPLS Services	COMMON	13-Oct-2018	✓	*	DoradoSoftware	7.0
Group Operations	COMMON	13-Oct-2018	✓	*	DoradoSoftware	7.0
RTCP Monitoring	COMMON	13-Oct-2018	✓	*	DoradoSoftware	7.0
L3VPN Services	COMMON	13-Oct-2018	✓	*	DoradoSoftware	7.0
Right to Manage	CORE	13-Oct-2018	✓	*	DoradoSoftware	7.0
Juniper JUNOS ...	ALL	13-Oct-2018	✓	*	DoradoSoftware	7.0

License Details NetConfig

```
Product License for NR
        EDITION = COMMON
        DESCRIPTION = NetConfig
        USER = DoradoSoftware
        IP = *
        ACTIVATION DATE = 15-Dec-2017
        EXPIRATION DATE = 13-Oct-2018
        VERSION = 7.0
        KEY = UAU-07277194-78188-E -- IS VALID
        APPPROPS:
                Permission-NR:ConfigFileGeneration=FFFF;
                Permission-NR:ConfigureArchives=FFFF;
```

The header displays the Expiration Date and how many days you have left on your license. For multiple licenses that are about to expire, the expiration date is for the license that expires first. Depending on the number of days left, the date and number of days change colors. See License Expiration Notification for more details.

The License Viewer provides access to the following panels by selecting the appropriate tab:

- Product
- Resource RTM
- Register License

See Registering a License on page 94 for detailed instructions.

License Expiration Notification

In addition to the expiration information provided in the License Viewer header, the System includes an alarm warning of possible license expiration (emsAppServerLicenseWillExpireSoon). If your license is about to expire, a **license expiration warning** similar to the following **example** displays in the Alarms portlet.

```
Application server license will expire in 30 days, on 2017-7-14,0:0:0.0,-
 -8:0
```

The time (0:0:0:0) indicates midnight, and --8 hours is the offset from GMT.

When this event occurs, the portal status bar color changes and displays the following message, which includes the number of days until license expiration:

```
Your Application Server license will expire in 30 days causing the system to
shutdown" in red status banner.
```

The alarm color reflects the threshold for the number of days remaining before license expiration:

- 60 days or less: yellow status bar
- 30 or less days: red status bar
- 61 or more days status bar does not change colors

A 0 days left message means that you have exceeded the license expiration. When the **Application license expires**, the following occurs:

- Any data collected up to this point remains stored and **no new data is collected.**
- Status bar does not change colors and does not display a license expired message.
- Most portlets display a message saying they cannot reach the Application server.
- License Management option is still available to register a license when the Application server is down.

  Select Settings > Application Configuration Settings from the navigation bar, click the License Management option, select the Register Product License option, and then click the Select and Register License button. A successfully registered license message is displayed.

When an **RTM license expires**, the following occurs:

- Resources under management are placed in a suspended management state until the number of resources under management meets the RTM license total.
- Suspended resources remain in inventory but cannot be accessed and no alarms are received.
- Aresourcesuspended alarm(resourceSuspended) israised foreachsuspended resource to the user visibility.

NOTE:

Extending yourlicense clears any expiration alarm and status barcolor but might leave earlier warning alarms in your event/alarm histories. Manually clear the history if you like.

Increasing your license that has 30 orless daysleft displays the following message and the status bar turns blue. The message shows and the status bar remain blue until the initial license count expires.

```
A Right to Manage (25 count) license will expire in xx days reducing your
  licensed count.
```

These behaviors applies whether the software is a trial version or yearly subscription.

Product

Use the Product Licenses list to monitor the products for which you have licenses, their expiration dates, and other details about the selected license. The most important product license listed is Oware Standard Edition, which is the Application server license. The Application server requires a valid license to run. All other components depend on a running Application server. If the Application server license expires, you cannot start the server and most applications that you start display a message that the server cannot be reached.

Access this panel by clicking the Products tab from the License Viewer window.

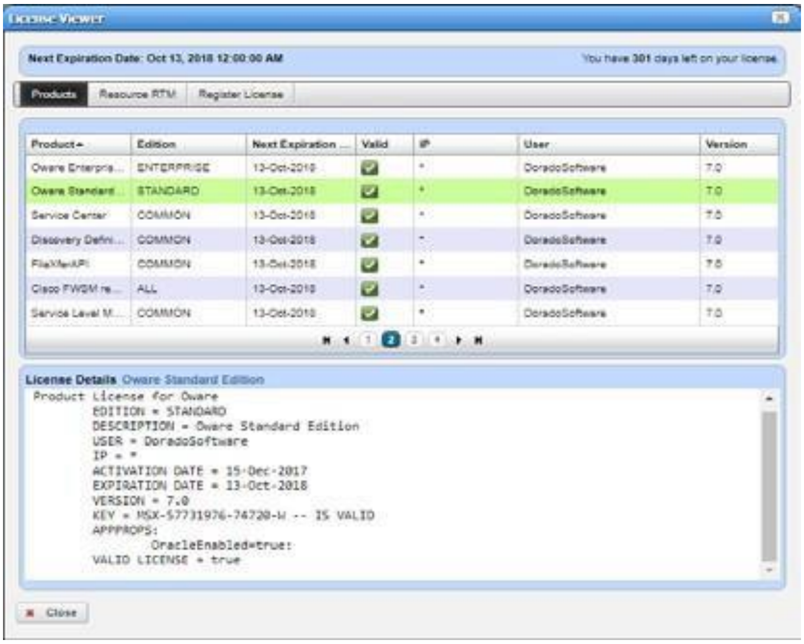

Next Expiration Date: Oct 13, 2018 12:00:00 AM

You have 301 days left on your license.

Product ▲	Edition	Next Expiration ...	Valid	IP	User	Version
Oware Enterpris...	ENTERPRISE	13-Oct-2018	✓	*	DoradoSoftware	7.0
Oware Standard...	STANDARD	13-Oct-2018	✓	*	DoradoSoftware	7.0
Service Center	COMMON	13-Oct-2018	✓	*	DoradoSoftware	7.0
Discovery Defini...	COMMON	13-Oct-2018	✓	*	DoradoSoftware	7.0
FileXferAPI	COMMON	13-Oct-2018	✓	*	DoradoSoftware	7.0
Cisco FWSM re...	ALL	13-Oct-2018	✓	*	DoradoSoftware	7.0
Service Level M...	COMMON	13-Oct-2018	✓	*	DoradoSoftware	7.0

1 2 3 4

License Details Oware Standard Edition

```
Product License for Oware
        EDITION = STANDARD
        DESCRIPTION = Oware Standard Edition
        USER = DoradoSoftware
        IP = *
        ACTIVATION DATE = 15-Dec-2017
        EXPIRATION DATE = 13-Oct-2018
        VERSION = 7.0
        KEY = MSX-57731976-74720-W -- IS VALID
        APPPROPS:
                OracleEnabled=true:
        VALID LICENSE = true
```

The Product panel provides the following license information.

Column/Field	Description
Product	The product for which you have a license.
Edition	Information useful for support, such as the distinction between core and additional.
Next Expiration Date	The date and time when your license expire. The Application server expiration date controls the Expiration Date information listed in the License Viewer header. The header displays the Expiration Date and how many days you have left on your license.
Valid	An indicator showing whether the license is valid (checkmark) or not (X).
IP	Any IP restrictions (asterisk is unrestricted).
User	The person authorized to use the application.
Version	The license version (not the software).
License Details	The selected registered license details display information, such as edition, activation date, expiration date, and so on. By default, this field displays information for the first product in the list. To see your license activation date, for example, select the Oware product license and then look in the License Details field for the ACTIVATION DATE parameter. The date is either: • The day you installed the Cruz system • The day registered a new license by increasing the license • The day the current license expires if you register a new license by extending your license • A given day if you register a new license from the command line To see your software’s Digital Service Tag parameter, select the Oware license and look in the License Details field for the APPPROPS parameter. This parameter is useful in license renewal.

Resource RTM

Use the Resource RTM panel to view Right to Manage (RTM) license information.

When the RTM license expires, any resource under management exceeding the total RTM count is put into a suspended state based upon the updated date. The most recently accessed resources are the last resources suspended due to an expired RTM license.

NOTE:

If a device's management state is Planned, Decommissioned or Suspended, the device is not included in the total RTMcount.

Also, resources classified as Other, do not count against the total RTM count.

When attempting to change a resource's management state, CruzOC verifies the current RTM licensing to ensure that the RTM license count is enough to bring the device back under management. If the RTM licensing check fails, the resource remains in a suspended state. You have the option to suspend another resource to bring a higher priority resource back under management (management state = Normal or Maintenance).

Access this panel by clicking the Resource RTM tab from the License Viewer window.

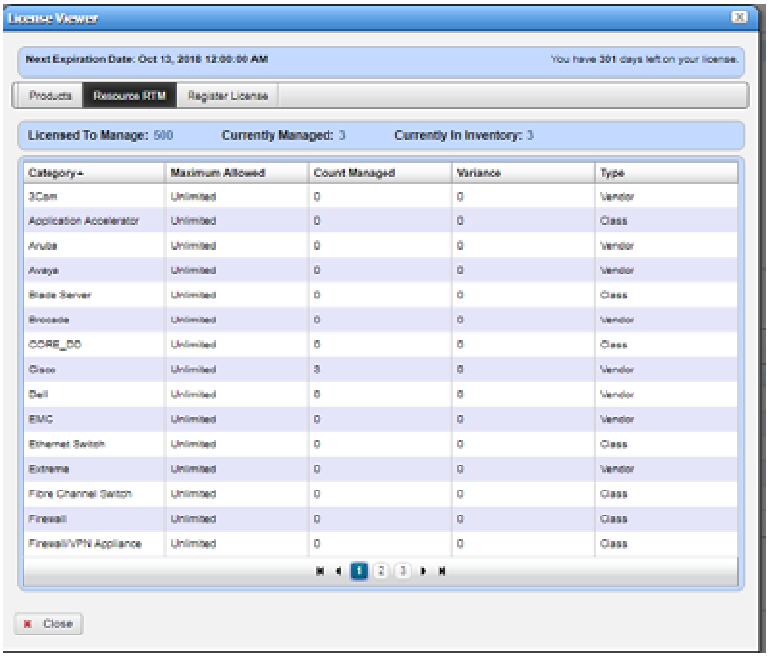

Next Expiration Date: Oct 13, 2018 12:00:00 AM

You have 301 days left on your license.

Licensed To Manage: 500 Currently Managed: 3 Currently In Inventory: 3

Category	Maximum Allowed	Count Managed	Variance	Type
3Com	Unlimited	0	0	Vendor
Application Accelerator	Unlimited	0	0	Class
Aruba	Unlimited	0	0	Vendor
Avaya	Unlimited	0	0	Vendor
Blade Server	Unlimited	0	0	Class
Brocade	Unlimited	0	0	Vendor
CORE_DD	Unlimited	0	0	Class
Cisco	Unlimited	3	0	Vendor
Dell	Unlimited	0	0	Vendor
EMC	Unlimited	0	0	Vendor
Ethernet Switch	Unlimited	0	0	Class
Extreme	Unlimited	0	0	Vendor
Fibre Channel Switch	Unlimited	0	0	Class
Firewall	Unlimited	0	0	Class
Firewall/VPN Appliance	Unlimited	0	0	Class

The Resource RTM panel provides the following license information.

Column	Description
Licensed to Manage	The number of devices you are licensed to manage.
Currently Managed	The number of devices currently under management.
Currently in Inventory	The number of devices in inventory.

Column	Description
Category	In addition to several self-explanatory categories (like Ethernet Switch, Router, and so on), these can include the following: • Other– Any type of device that is not explicitly a current licensed Category. An Other type does not impact the Licensed To Manage for Currently Managed values. **Example:** The Cruz system classifies Wireless Access Points (APs) are Other. These have an unlimited Maximum Allowed count because no license limit exists for the number of Wireless APs that can be attached to a Controller. APs attached to a controller are counted for Other. • Unknown–Any device that is successfully discovered with an unrecognized SysObjectId is licensed in the Unknown Category. • Wireless Access Point– The Cruz application licenses any Wireless access point you create by right clicking in the Managed Resources portlet and selecting the New menu item as Wireless Access Point. The Dell Instant Access Points (IAPs) are also licensed as Category Dell.
Maximum Allowed	The maximum number of managed devices. If *unlimited* appears it refers to device types, not numbers of devices licensed, and indicates no limits exist for classes of devices related to the license row.
Count Managed	The current count of managed devices.
Variance	Licensing can give you a variable number of additional Rights to Manage (RTMs) devices beyond the maximum allowed. For example, if *Maximum Allowed* is 100, a variance of 5 allows five additional RTMs. This allows some flexibility to manage additional devices if while awaiting the arrival of additional RTMs. Typically, *Variance* is zero.
Type	The Licensing scheme supports RTM by Class of device. The class of device has an RTM set. Discovery also maps Class type to a device driver. **Note:** The Cruz system does not license by vendor.

Register License

Use the Register License panel to register your license. Access this panel by clicking the Register License tab from the License Viewer window.

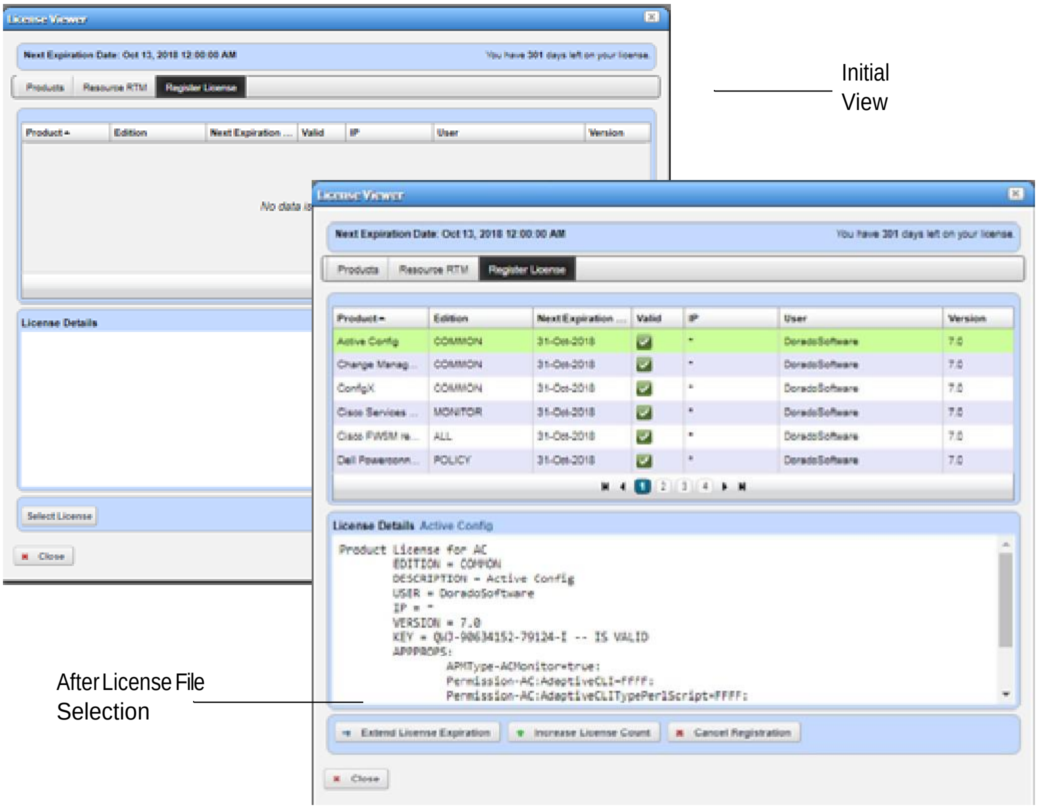

Initial View

After License File Selection

Next Expiration Date: Oct 13, 2018 12:00:00 AM You have 301 days left on your license.

Product	Edition	Next Expiration ...	Valid	IP	User	Version
Active Config	COMMON	31-Oct-2018	✓	*	DoradoSoftware	7.0
Change Manag...	COMMON	31-Oct-2018	✓	*	DoradoSoftware	7.0
ConfigX	COMMON	31-Oct-2018	✓	*	DoradoSoftware	7.0
Cisco Services ...	MONITOR	31-Oct-2018	✓	*	DoradoSoftware	7.0
Cisco FWSM re...	ALL	31-Oct-2018	✓	*	DoradoSoftware	7.0
Dell Powerconn...	POLICY	31-Oct-2018	✓	*	DoradoSoftware	7.0

License Details Active Config

```
Product License for AC
        EDITION = COMMON
        DESCRIPTION - Active Config
        USER = DoradoSoftware
        IP = *
        VERSION = 7.0
        KEY = QWJ-90634152-79124-I -- IS VALID
        APPPROPS:
                APMType-ACMonitor=true:
                Permission-AC:AdaptiveCLI-FFFF:
                Permission-AC:AdaptiveCLITypePerlScript-FFFF:
```

The Register License panel includes the following fields and options.

Fields/Options	Description
Product	The product for which you have a license.
Edition	Information useful for support, such as the distinction between core and additional.
Next Expiration Date	The date and time when your license expire. The Application server expiration date controls the Expiration Date information listed in the License Viewer header. The header displays the Expiration Date and how many days you have left on your license.
Valid	An indicator showing whether the license is valid (checkmark) or not (X).
IP	Any IP restrictions (asterisk is unrestricted).
User	The person authorized to use the application.
Version	The license version (not the software).

Fields/Options	Description
License Details	The selected registered license details display information, such as edition, activation date, expiration date, and so on. By default, this field displays information for the first product in the list. To see your license activation date, for example, select the Oware product license and then look in the License Details field for the ACTIVATION DATE parameter. The date is either: • The day you installed the Cruz system • The day registered a new license by increasing the license • The day the current license expires if you register a new license by extending your license • A given day if you register a new license from the command line To see your software's Digital Service Tag parameter, select the Oware license and look in the License Details field for the APPPROPS parameter. This parameter is useful in license renewal.
Extend License Expiration	This option activates the registered license once the current license expires, which extends the license period by the length of the new license (recommended for renewed subscriptions). **For example:** Your Cruz licensed is for 50 devices and one year starting January 1, 2018. On November 1, 2018, you apply an additional 50 devices and one-year license using the Extend option. On December 31, 2018, the original license expires, and the new license takes effect with no disruption of service. You continue to be licensed for 50 devices until December 31, 2019.
Increase License Count	This option activates the registered license immediately. Use this option for: • Upgrading a Trial version • Increasing your subscription to a higher count This option applies the new license immediately and adds the new purchased license device count to the existing device count. The new license expiration countdown starts immediately. Expiration notification continues until the existing trial/subscription device count elapses. Once expired, the trial/subscription device count is deducted from the total device count leaving only the new license device count. **For example:** Your Cruz license is for 50 devices and one year starting January 1, 2018. On July 1, 2018, you apply an additional 25 devices and one-year license using the Increase option. You can manage 75 devices from July 1, 2018 to December 31, 2018. On December 31, 2018, the original 50 devices license expires, with 25 devices remaining until June 30, 2019. See License Expiration Notification on page 56 to understand how license expiration notification works.
Cancel Registration	The option to discontinue with the license registration process.

Common Setup TasksPortlet

By default, the Common Setup Tasks portlet appears on the Home page as part of the Getting Started portlet and the Settings page. If your package does not display this portlet on these pages and you want it there, click *Add > Applications* and put it there.

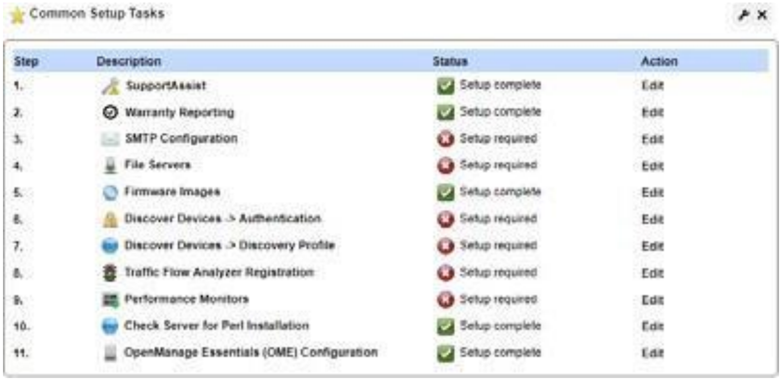

Common Setup Tasks

Step	Description	Status	Action
1.	SupportAssist	✓ Setup complete	Edit
2.	Warranty Reporting	✓ Setup complete	Edit
3.	SMTP Configuration	✗ Setup required	Edit
4.	File Servers	✗ Setup required	Edit
5.	Firmware Images	✓ Setup complete	Edit
6.	Discover Devices -> Authentication	✗ Setup required	Edit
7.	Discover Devices -> Discovery Profile	✗ Setup required	Edit
8.	Traffic Flow Analyzer Registration	✗ Setup required	Edit
9.	Performance Monitors	✗ Setup required	Edit
10.	Check Server for Perl Installation	✓ Setup complete	Edit
11.	OpenManage Essentials (OME) Configuration	✓ Setup complete	Edit

This portlet shows the setup status for the common tasks, such as:

- SMTP Configuration
- File Servers
- Firmware Images

The Status column shows whether the setup is complete (check mark) or is required (X). Click the Edit action to modify a task’s settings or complete its setup.

SMTP Configuration

Youcan use the CruzOC messaging capabilities to communicate with other users, but if you want to receive e-mails automated by actions, such as configuration file backups, the Cruz application must have a mail account. The SMTP Configuration window configures the e-mail server so the Cruz application can send such automated e-mails.

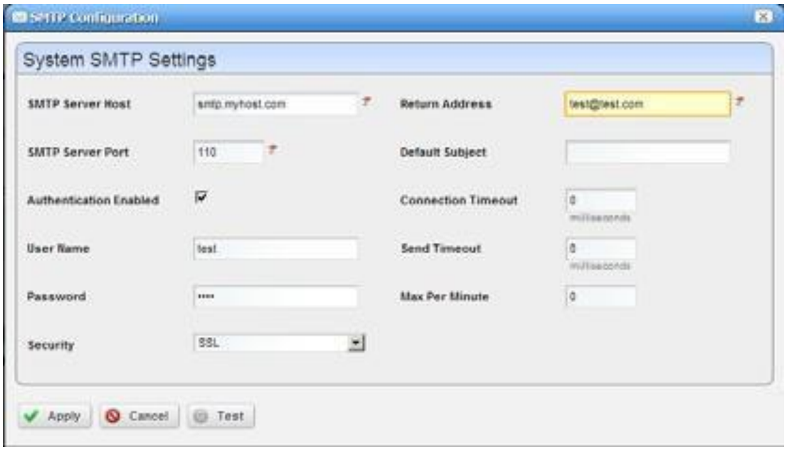

System SMTP Settings

SMTP Server Host	smtp.myhost.com *	Return Address	test@test.com *
SMTP Server Port	110 *	Default Subject	
Authentication Enabled	☑	Connection Timeout	0 milliseconds
User Name	test	Send Timeout	0 milliseconds
Password	****	Max Per Minute	0
Security	SSL		

The Apply button accepts your edits. The Test button tries your edits. The Cancel button abandons your edits and returns to the Cruz application. The SMTP Configuration window contains the following fields and options.

Field/Option	Description
SMTP Server Host	The IP address or hostname of your SMTP server.
SMTP Server Port	The port for your SMTP server. (Common ports are 25, 465, 587)
Authentication Enabled	Check this to enable authentication for this server. Checking enables the next two fields.
User Name	The login ID for the SMTP server, if authentication is enabled.
Password	The password for the SMTP server, if authentication is enabled.
Security	Enable Secure Sockets Layer (SSL) protocol to interact with your SMTP server, or Transport Layer Security (TLS).
Return Address	The return address for mail sent from Cruz.
Default Subject	Text that appears by default in the subject line of mail sent by Cruz.
Connection/Send Timeout	The time-outs for mail sent by Cruz. If your SMTP server or network is slow, increase the default timeout. **Note:** These time-outs are in milliseconds and have been critical in getting e-mail to work, so do make them long enough to handle whatever latency is normal for your network.
Max Per Minute	The maximum number of e-mails Cruz can send per minute. Two settings for e-mail servers appear in Control Panel, one in the Control Panel > Portal > Settings Mail Host Names edit screen, and another in Control Panel > Server Administration > Mail. These are for Liferay login and password reminders/resets (see Resetting a Password on page 106). The Portal-based e-mail settings help Administrators limit signups to e-mails only existing in their organization. The screen in that panel provides a list of allowed domain names, if that feature is enabled. Control Panel > Server Administration > Mail is where to configure the Main server and authentication for routing mail. **Note:** If you require a sender/reply to e-mail address on mail sent, configure that with the following property in the owareapps/installprops/lib/installed.properties file: `redcell.smtp.returnaddress.name`

File Servers

The file servers provide FTP connections for retrieving and deploying devices' configuration files, and for deploying firmware updates to devices on your network. See Configuration Management for a description of the portlet that manages file servers. If you want to configure servers from the *Common Setup Tasks* portlet, a slightly different screen appears when you click *Edit*.

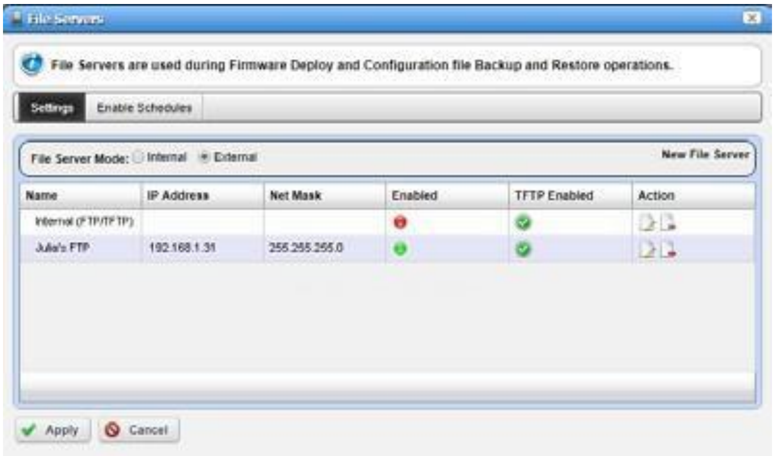

File Servers are used during Firmware Deploy and Configuration file Backup and Restore operations.

File Server Mode: ○ Internal ● External

Name	IP Address	Net Mask	Enabled	TFTP Enabled	Action
Internal (FTP/TFTP)			✗	✓	
Julia's FTP	192.168.1.31	255.255.255.0	✓	✓	

This displaysconfigured file servers. Configure new servers byclicking the *new file server* link in the upper right corner. The editing process after that is as described in External File Server Editor on page 462.

△ CAUTION:

If you select the internal file server, make sure no external file server is running on the same host. A port conflict prevents correct operation. Either turn off the external file server or use it as the FTP server. We strongly recommend using the internal file server onlyfor testing, andexternal file server(s) for production.

Cruz selects the file server protocol for backup, rest ore or deploy based on the most secure protocol the device supports.

Firmware Images

Use the Firmware Images window to copy firmware images into the Cruz database. This opens an editor like the one described in Firmware Image Editor on page 454. Refer to Configuration Management, particularly Image Repository on page 452 for more about using this capability to deploy firmware to devices you have discovered.

General Information

The following topics provide some general information you may need to know:

- Web Portal/Multitasking
- Cookies and Sessions
- Web Clients
- Network Basics

Web Portal/Multitasking

Youcan open multiple tabs to different managers in the Cruz. In most cases this does not cause any issues for read–only da ta browsing. However, **best practice** is not opening multiple tabs when creating, editing or deleting. These may report Web session information incorrectly and task completion may appear to never finish. For example, you may submit a job that appears stuck “running” when it has already finished but the status has not updated in the browser session. When this occurs, manually click the refresh button on the job status window to force an update. The recommended process is to close the job status window and move on to other tasks.
The *My Alerts* feature in the portal’s lower left corner indicates when jobs are complete, and you can view details from that status bar location.

Cookies and Sessions

The Cruz application stores cookies during the session. No personal data, such as username/ password is stored. The cookie stores the assigned session ID as well as information about the current view so that the Cruz application can do partial view updates. Cruz sessions are not persisted, they are in memoryonly.

Web Clients

With the minimum hardware for a 32-bit client, you can run **either** a Web client or Java client, **not both**. Youcanstart and stop the client without impacting the application server. Device monitoring stops when you stop the application server or turn off its host machine. The client can also be on a different machine than the application server.

NOTE:

See Starting Web Client on page 70 for more information about using web access to this software.

The Cruz application detects mobile devices and pads, such as tablets and iPads. For smaller screens, the Navigation bar collapses to the left-hand side and the page only displays a single column. Some limits apply:

- Because touch devices do not support a right-click, the first time clicking a row selects it. A repeat click launches a menu displaying the available actions. Click the menu item you want.
- All major charts are rendered as HTML 5, which are mobile-friendly. These charts are Line, Pie, Donut, Bar and Column. Some Gauges and LED charts require flash, which is not compatible with all mobile devices.

- Topology views are not available.

NOTE:

Apple products are mostly Cruz-friendly. Android is only partly supported.

CAUTION:

If your application server and client are on different machines, make sure that they have the same time settings.

Network Basics

The Cruz application communicates with devices over a network. You must be connected to a network for the Application server to start successfully. Firewalls, or programs using the same ports on the same machine where this application is installed can interfere with its ability to communicate with devices.

Your network may have barriers to communication with this software that are outside the scope of these instructions. Consult with your network administrator to ensure this application has access to the devices you want to manage with the protocols described below.

NOTE:

One simple way to check connectivity with a device is to open a command shell with Start > Run `cmd`. Then, type `ping [device IP address]` at the command line. If the device responds, it is connected to the network. If not, consult your network administrator to correct this. No useful information comes from disconnected devices.

Consider	Description

Name Resolution (equipment)	The Cruz application requires resolution of equipment names, whether by host files or domain name system (DNS).
	If your network does not have a DNS, you can also assign hostnames in the following file on a Windows system:
	`%windir%\system32\drivers\etc\hosts`
	You must assign a hostname in addition to an IP address in that file. Here are some example hosts file contents (including two commented lines where you would have to remove the # sign to make them effective):
	`# 102.54.94.97 rhino.acme.com # source server`
	`# 38.25.63.10 x.acme.com # x client host`
	`127.0.0.1 localhost`
	Caution: This software **only** supports installation to a local file system. **Do not** install to shared drives.

Protocols	The Cruz application uses the following protocols: • TCP/IP • SNMP • HTTP • UDP Multicast You can bypass multicast, if it is disabled on your network. To allow a client to connect without multicast, add the following property to the client’s owareapps/installprops/lib/installed.properties file. `oware.application.servers[Host1 IP address], [Host2 IP address]...`

Consider	Description
Fixed IP Address	The Cruz application includes a Web server and must be installed on a host with a fixed IP address. For demonstration purposes, you can rely on dynamic IP address assignment (DHCP) with a long lease. However, this is not recommended for production installations. If you need to change your host's IP address, do so by: • Changing Your Host's IP Address using the ipaddresschange script • Manually Changing Your Host's IP Address
Partition Name Limitations	First character must be a letter (a-z, A-Z). The remaining must be alphanumeric or underscore or dash characters (a-z, A-Z, 0-9 and _ or -). Maximum length is 31 characters. In a clustered installation, make `oware/synergy/data` a shared directory because local user images, documents, and uploads go there, and in a cluster environment all web servers need toaccess this directory.

Changing Your Host's IP Address

Change your host's IP address as follows.

1. Change the Virtual host IP to the new IP address in Manage > Control Panel > Portal.
2. Change the host IP address
3. Open a shell and run `oware` to set the environment
4. Run `ipaddresschange -n` in the shell followed by the new IP address
5. Restart the application server and the web server service.
6. Open a browser to see the web client at this URL: `[new IP address]:8080.`

Manually Changing Your Host's IP Address

Change your host's IP address without the `ipaddresschange` script as follows.

1. Change the Virtual host IP to the new IP address in Manage > Control Panel > Portal.
2. Change the host IP address
3. Delete the contents of `\oware\temp`.
4. Change your local IP address anywhere it appears in `\owareapps\installprops\lib\installed.properties`.
5. Change the address on your web server. Change this in `portal-ext.properties` in `\oware\synergy\tomcat-7.0.40\webapps\ROOT\WEB-INF\classes`

   Change property:

   ```
   jdbc.default.url=jdbc:mysql://[IP address]/
    lportal?useUnicode\=true&characterEncoding\=UTF-
    8&useFastDateParsing\=false
   ```

   and

   ```
   oware.appserver.ip=[IP address]
   ```
6. Restart the application server and the web server service.
7. Open a browser to see the web client at this URL: `[new IP address]:8080.`

Starting/Stopping Cruz

The Cruz Application server processes network information for W eb-based clients. The Application server monitors devices and produces the output for the Web server and then makes it available for the Web clients. When you install or upgrade the Cruz application, the installer automatically stops the servers before installation and starts them after installation. However, there are many instances where you need to stop/stop the servers. The following tasks show how to start the Application server and its subordinate process, the mediation server, which communicates directly with devices, and how to stop these processes when needed:

- Starting/Stopping Servers
- Starting Web Client

Starting/Stopping Servers

- **Starting application server**. In Windows, you can use the *Start* button (*Start* > Cruz Software > *Start application server*), or type `startappserver` in a command shell, or right-click the server manager tray icon and select *Start (*if you have installed this software as a service and that icon is red , not green*)*.

NOTE:

A message declares "Application server is now up" in *My Alerts* in the bottom left corner of the screen of the web client when application server startup is complete. You can also make server monitor appear with the `pmtray` command either in a shell or from a start menu icon.

- **Starting web server**. If this does not auto-start, you can use the *Start* button (*Start* > Cruz Software > *Synergy Manager*), or right click the web server's tray icon to start it. Youcan also double-click this icon and automate web server startup. From a command line, you can also start this manager with `[installation root]\oware\synergy\tomcat*\bin\startsynergy.`

  To start the Web server in Linux in a shell, type `/etc/init.d/synergy start.` Stop web server with `/etc/init.d/synergy stop.`

△ CAUTION:

If your Cruz environment has a firewall, ports 8080 must be open for it to function correctly. If you want to usecut-thruoutside ofyournetwork, thenports8082 –8089mustbeopen. Cruzusesthefirst one available, so typically 8082, but if another application uses 8082, Cruz uses 8083 and so on. Web Services forCruzpreviouslyusedport80,butforthisversion,theyuse8089.SeePortsUsedonpage979fora complete list of all ports impacted.

Here are the various ways to start (and stop) Cruz elements:

Windows Start Menu Program Shortcut	Windows Command Line	Linux Command Line
Server Monitor	pmtray	N/A
Start Application Server	startappserver	startappserver
Network Manager	startsynergy.com **Note:** this is in the oware\synergy\tomcat*\bin directory and is not on the path.	While no monitor display appears, you can start the Web server with these commands: startportal.sh start/ startportal.sh stop

Windows Start Menu Program Shortcut	Windows Command Line	Linux Command Line
Synergy Web Service Manager	http://[application server host IP]:8080	ht tp://[application server host IP]:8080

Starting Web Client

Open the Web client user interface from your browser. Refer to the *Cruz Installation Guide* for supported browsers and versions. The URL is:

```
http://[application server hostname or IPaddress]:8080
```

The default login user is *admin*, with a password of *admin*. The first time you log in, you can select a password reminder. If you have forgotten your password, click the Forgot Password link in the initial screen to begin a sequence that concludes by mailing your user's e-mail address a password. (See Resetting a Password on page 106.)

For this forgotten password sequence to work, you must configure users' e-mails correctly, and the portal's SMTP server in Go to > Control Panel > Server > Server Administration > Mail settings. To configure a user's e-mail settings and other things, click the usernamelink in the portal header. The same configuration settings are available in Control Panel's tabs labeled as that user's login.

The *application server hostname* is the name of the system where Cruz is installed.

△ CAUTION:

The first time you start the application after you install it, you may have to wait an additional five minutes for the Application tocompletelystart. Oneindicationyouhavestarted toosoonis that the Quick Navigation portlet does not appear properly. Also: The Webservermay indicate it has fully started before it is entirely ready. In rare instances, this may also inhibit correct communication with the client interface. If Cruzappearsstuckafter theapplication server is runningcompletely, restarttheWeb server.

Disable password reminders with `users.reminder.queries.enabled=false` to `oware\synergy\tomcat-x.x.xx\webapps\ROOT\WEB-INF\classes\portal-ext.properties`

Setting Up Secure Connections(SSL & HTTPS)

The following describes how to turn on SSL support within Cruz on single-server installations. Configure Clustered installations witha Load Balancerwith SSLOffloading. SSLOffloadingtakes advantage of hardware which has been designed to deal with quick encryption and decryption of SSL. It also lets you purchase a single SSL certificate rather than generating a certificate per server, something that can be more costly.

NOTE:

If you want a secureconnection between distributed servers (application and mediation servers, for example), the following also applies.

Enabling Secure SSL

Bestpracticeforaclustered productionenvironment is to usea Load Balancer with SSLOffloading rather than creating a private key, as described below. Refer to the *Installation Guide* for more about load balancing.

The private keyand certificate described belowprov ides identity and browser verificationagainst the CA signed root certificate. For testing and internal use, you need this step to create a Private Key and Private Signed Certificate to enabled SSL encryption.

NOTE:

Somefunctions mayfail usingthis approach sincesomethird-partylayers mayexpect avalid CAsigned root.

Creating a Private Key (Linux/Windows)

1 Open a command prompt in Wind ows or a Terminal within Linux

2 Navigate to a `<INSTALL DIR>/oware/synergy/tomcat-XX/bin/certs`

3 Enter the command: `openssl`

If this command does not find openssl, then first enter the oware environment (in Windows type `oware`, in Linux, type . ./etc/.dsienv).

4 The OpenSSL prompt appears: `OpenSSL>`

5 Enter the command:

```
genrsa -des3 -out tomcatkey.pem 2048
```

6 OpenSSL then asks for a pass phrase for the key. Enter `changeit`. See Turning on SSL Within the Web Portal on page 72 if you want to change the default password.

7 OpenSSL then creates the private key and stores it in the current directory

Creating a Certificate (Linux/Windows)

Once you have the private key created, you must create a certificate.

8 Assuming you are still running the OpenSSL program from the previous step, enter the command:

```
req -new -x509 -key tomcatkey.pem -out tomcat.pem -days 1095
```

9 OpenSSL asks for the pass phrase defined for the private key. Enter the previous pass phrase (default: `changeit`). This command creates a self-signed certificate with a lifetime of 3 years, using the private key.

This password must be identical to the one entered in the previous steps.

10 When asked the other questions such as Country Code, Organization you can enter any data you wish. When asked for the Common Name (FQN) you must enter the hostname or IP address of the server.

11 OpenSSL generates the `tomcat.pem` in the directory you were in from the previous steps.

12 Exit OpenSSL by typing `exit`

13 Two new files appear within the `//../tomcat-xx/bin/certs` directory: `tomcatkey.pem` and `tomcat.pem`

Some systems may put these files in another directory (for example `C:\users\[username]` on Windows 7). If so, copy or move them to the `oware\synergy\tomcat-7.0.40\bin\certs` directory before proceeding.

Turning on SSL Within the Web Portal

Windows: Changing the Environment:

First, update the `setenv.bat` with the SSL preferences. You must do this whether Cruz's web server starts manually or runs as a service. if Cruz runs as a service, this file automatically updates the service on the next portal service restart.

1. Stop Cruz service
2. Navigate to the `<INSTALLDIR>/oware/synergy/tomcat-xx/bin` directory.
3. Edit the `setenv.bat` file in a text editor.
4. Change the property `ENABLE_SSL=false` to `ENABLE_SSL=true.`
5. If you used a pass phrase different from `changeit` then you can set it for the `SSL_PASSWORD=changeit` value.
6. Save `setenv.bat`
7. In a command prompt navigate to `/oware/synergy/tomcat-xx/bin`, and type: `service.bat update`
8. Settings take affect after the you restart the service.

You are now ready for a secure, SSL connection to Cruz. After it has had a few minutes to start navigate to `https://[application server IP address]:8443.` (The HTTPS port is 8443, not 8080.)

Linux: Changing the Environment

1. Enter the command: " `service synergy stop`" to stop the Cruz web service.
2. Navigate to the `/oware/synergy/tomcat-xx/bin` directory
3. Edit the `setenv.sh` file.
4. Change `ENABLE_SSL` to `true`.
5. If you used a different pass phrase than the default (`changeit)` then you can set it for the `SSL_PASSWORD` property here.
6. Save the file.
7. `Enter the command: "service synergy start" to restart the Cruz web service.`

You are now ready for a secure, SSL connection to Cruz. After it has had a few minutes to start navigate to `https://[application server IP address]:8443`

Heartbleed SSL Vulnerability

Cruz is not vulnerable as shipped. If the client does not have SSL turned on with a valid certificate, then the following does not matter:

When running Linux then your system admin must keep OpenSSL up to date. This is native to Linux, not to Cruz.

For windows Cruz ships with 0.9.8d which is not affected. You can update this any time. This is *only* applicable if you are using SSL by replacing the `openssl.exe` in `oware/synergy/tomcat-xx/bin/native/windows/x64`.

Enabling Terms of Use

To Enable a "Terms of Use" statement required of each user use the following steps:

1. Login as Admin
2. Go to Control Panel
3. Click on Portal Settings and then the Users link on the right, and look in the Fields tab.
4. Check *Terms of Use Required* and save. You must then click *I Agree* to the Terms of Use document that appears.
5. Logout and attempt to login as another user to validate the Terms of Use appear.

To change the Terms of Use wording:

1. Login as Admin
2. Go to the Synergy Control Panel
3. Click on Web Content
4. Click on the TERMS-OF-USE article link which will take you to the editor where you can alter and save it.

NOTE:

Nothing prevents a user from deleting the Terms of Use article. If the Terms of Use seeded article is removed, then the static Liferay Terms of Use appears until next Cruz restart. The editable/delete-able article is a copy of the compiled static version but exposed as an article to make editing easier. The next time Cruz restarts, if the TERMS-OF-USE article does not exist, it imports a new one.

Managing Users and Permissions

Adding Users and Connecting them to Roles

When you add a new user, that user may not appear immediately. You can speed up the user's appearance by using control panel's *Server > Server Administration Resource* panel. Click *Reindex all search indexes.*

Add Users with the following steps:

1. Click *Go to > Control Panel* and navigate to Portal > Users and Organizations .
2. Click the *Add > User* menu item at the top of the *Users* screen.
3. Enter the details of the new user. If you are editing an existing user, more fields appear. *Screen Name*, and *Email Address* are required.Optionally, you can enter *Name, Job Title*, and so on.

NOTE:

Make sure you specify a *Password* when you add a user. This is not optional.

4. After you click *Save* notice that the right panel expands to include additional information.

   The first-time users log in, the application prom pts them for a security question. E-mail for password reminders/resets requires setting up the fields in Control Panel > Server Administration > Mail, not the SMTP Configuration which is for Cruz-originated e-mails. See Resetting a Password on page 106

   Also: When you make a multitenant site, Cruz automatically assigns the site prefix you select to the admin user it creates in the Site Management Editor. If you enter "Admin" as that user, and the prefix is DS-, then that user must log in as "DS-Admin." When you or the tenant site admin create tenant site users manually in control panel, you must manually add that prefix too when creating the user and when logging in as that user.

5 Notice that if you are editing an existing user, or creating a new one, you can use the links on the right to configure connections with *Roles.* Roles configure the Cruz functional permissions for that user. For example, the *Operators* role's capabilities are typically more limited than *Administrators.* See Adding and Configuring User Roles/Permissions on page 75.

6 Click *Save* again, and the user you just configured should appear listed in the *Users* screen when you select *View* > *All Users*.

7 After you have configured roles as described in Adding and Configuring User Roles/ Permissions on page 75, return to the Users and Organizations screen, edit the User, and click the *Roles* link to associate the User with the Role(s) youhave configured.

The most dramatic evidence of permission changes appears when you first remove Default User Roles — Power User from your system in Portal > Portal Settings > Users > Default User Associations (check *Apply to Existing Users* if you have already configured your user). If you impersonate your user, and Go To > Control Panel, without User and Power User roles assigned, the impersonated user can only see *My Account* and *Sites.*

NOTE:

You can *Export Users* to a comma-separated value (CSV) file.

Once you have configured a user, you can click *Action* and to do the following:

Edit—Re-configure the selected user. Select the user's Role in the editor, too. Roles configure access and action permissions.

Permissions — Manage the user's access to and control over various parts of the portal.

Impersonate User (Opens New Window or tab) — This allows you to see the effect of any configuration changes you have made on a user. The new window (typically a new tab) also lets you click the *Sign Out* link in the upper right corner where you can return to your original identity impersonation concealed.

Manage Pages— The menu described below appears when you have not installed the Multitenancy option. Configure the *Public* or *Private* pages for a user, depending on the selected tab. Possible actions here include changing the look and feel of pages (for computers and mobile browsers), adding pages and child pages, and importing or exporting page configurations. Notice that you can configure meta tags, and javascript on these pages too.

Exports are in `.lar` format, and go to the download location configured in the browser you are using. The export screen lets you select specific features, and the date range of pages to export.

NOTE:

If you want to set up several pages already configured elsewhere for another user, or even for an entire community of users, export those pages from their origin, then *Manage > Pages* menu for the user or community.

Also: On private pages, you can see the *Languages* portlet. Click a flag to translate some labels to the language represented by the flag. You can change the text in many labels (the portlet titles, for one example) by clicking and re-typing that label. Some labels do not translate, no matter what.

See also the Multitenancy chapter of the *Installation Guide* for more potential variations on page appearance.

Deactivate — Retires a user configured on your system. You can also check users and click the *Deactivate* button above the listed users. Such users are not deleted but are in a disabled state. You can do an Advanced search for inactive users and *Activate* them or permanently delete them.

Your organization has several geographic locations and you plan to manage the network infrastructure for all these locations using Cruz. Youcan define the geographic locations to which devices can be associated. This will help you manage and view your network, grouped by location or branches. See Locations Portlet on page 207 for the specifics about the portlet where you can set up locations.

NOTE:

To edit your own information as a signed-in user, simply click your login name in the upper right corner of the portal screen.

Organizations

Create Organizations just as you would create Users. You can create a *Regular* or *Location* type of organization. You can do this only if your package includes the MSP option, so this capability is not available to all users.

NOTE:

You must first create a *Regular* organization to be the parent for a *Location*. Also: These organizations are useful to organize users. They are distinct from the device-organizing Locations described in Locations Portlet on page 207, and are not available for organizing in the Hierarchical Viewportlet (see Using Hierarchies on page 249).

Adding and Configuring User Roles/Permissions

Add and configure User Roles with the following steps:

1 Click Go to > Control Panel > Portal > Roles.

2 Click the *Add* tab under the heading at the top of the page and select Regular *Roles*.

The New Role options are displayed.

You can also add site and organization roles that configure permissions.

3 Enter the details of the new role (Name, *Title, Description*), then *Save* it.

4 Click Portal > Roles > *View All* to verify that the new role was added to the list.

The actions available for a new role are Edit, Permissions, Define Permissions, Assign Members, View Users, and Delete role.

5 Click Actions > Define Permissions for the new role.

The Define Permissions editor is displayed.

Alternatively, select or delete permissions by editing the role in *Permission Manager*.

NOTE:

If you are restricting permissions for new users, you must also remove the permissions from the *User* and *Power User* roles, that Cruz assigns all new users by default. The permissions available are a combination of those configured here and the User *and* Power User roles' permissions. You can remove users from the Power User role altogether, but not from the User role. You must remove permissions from that User role if you want users not to have them.

If you have eliminated all permissions from a role by removing the Default User Roles — Power User, an intervening screen lets you copy another Role's permissions, so you do not have to enter all permissions from scratch.

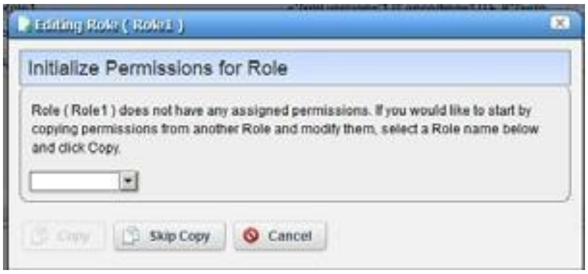

Initialize Permissions for Role

Role (Role1) does not have any assigned permissions. If you would like to start by copying permissions from another Role and modify them, select a Role name below and click Copy.

NOTE:

Defining a base role's permissions can provide the start for non-base role's permissions if you use this screen to copy them, then edit them later for the difference between the base role and non-base role. As always, planning is the key to simplifying this work.

6 Select the type of permission from the Add Permissions list.

7 Select the actions you want to enable.

8 Click the Redcell > Permission Manager to alter or enable more of Cruz functional permissions.

The Role to Permission Settings list is displayed.

9 Click the action's Edit button to see and configure available permissions.

The Editing Role window is displayed.

NOTE:

The Show Assigned/Show All options to filter what displays in the list.

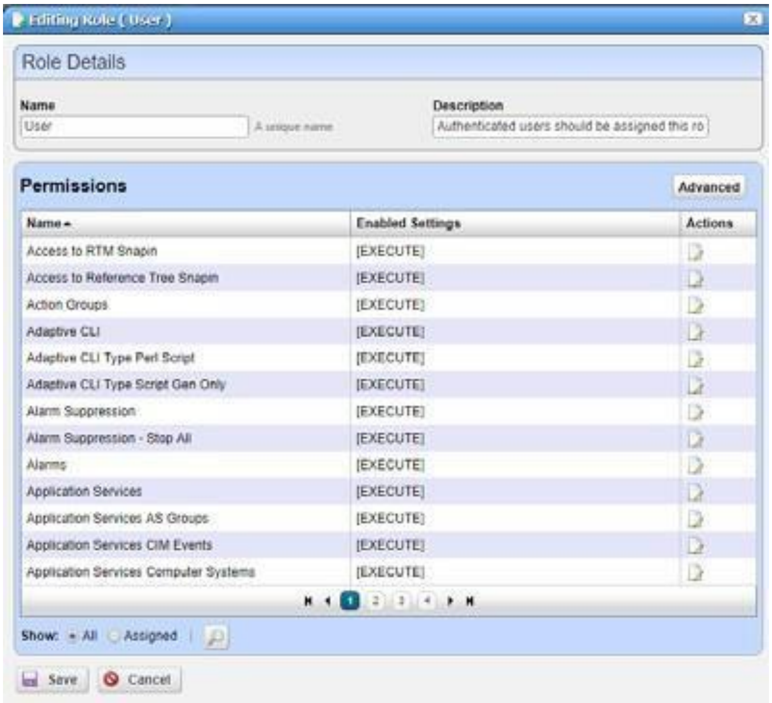

Name: User A unique name

Description: Authenticated users should be assigned this ro

Name ▴	Enabled Settings	Actions
Access to RTM Snapin	[EXECUTE]	
Access to Reference Tree Snapin	[EXECUTE]	
Action Groups	[EXECUTE]	
Adaptive CLI	[EXECUTE]	
Adaptive CLI Type Perl Script	[EXECUTE]	
Adaptive CLI Type Script Gen Only	[EXECUTE]	
Alarm Suppression	[EXECUTE]	
Alarm Suppression - Stop All	[EXECUTE]	
Alarms	[EXECUTE]	
Application Services	[EXECUTE]	
Application Services AS Groups	[EXECUTE]	
Application Services CIM Events	[EXECUTE]	
Application Services Computer Systems	[EXECUTE]	

1 2 3 4

Show: ● All ○ Assigned

10 Click *Advanced* to see available permissions organized by *Read*, *Write*, *Execute*, *Add* or *Delete* actions.

11 Select the appropriate permissions for each category.

12 Apply your changes.

13 Click the action's Edit button to modify permissions type.

The Editing Permissions window is displayed.

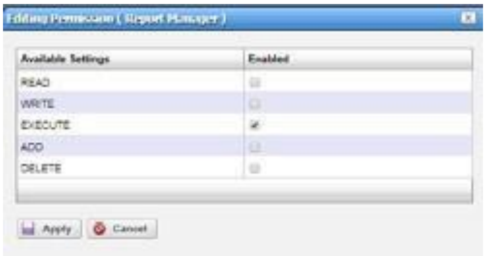

Available Settings	Enabled
READ	□
WRITE	□
EXECUTE	☑
ADD	□
DELETE	□

14 Enable and apply the appropriate permissions.

15 Save the role configuration.

Adding Individual Permissions

The Redcell > Permission Manager options are more convenient to do this in bulk, but to add individual permissions, click Portal > Roles > Actions > Define Permissions for the appropriate role, and then select the specific resource set.

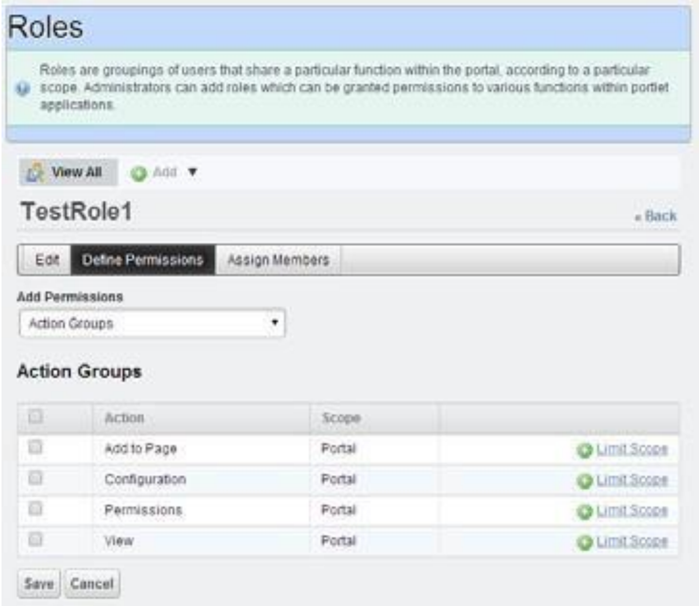

Roles

Roles are groupings of users that share a particular function within the portal, according to a particular scope. Administrators can add roles which can be granted permissions to various functions within portlet applications.

TestRole1

Add Permissions: Action Groups ▾

Action Groups

□	Action	Scope	
□	Add to Page	Portal	Limit Scope
□	Configuration	Portal	Limit Scope
□	Permissions	Portal	Limit Scope
□	View	Portal	Limit Scope

Once you select a general type of permission with the pick list below *Add Permissions*, the *Action Groups* that appear below let you check to select areas to enable, and the *Limit Scope* link to the right lets you further filter the application of this enabled permission. The *Private Pages* limitation in this final filter does not apply in Multitenant systems.

Because these screens are so granular, however, best practice is to use them to fine tune existing permission structures.

Multitenant Users

When you try to log in as a multitenant user, Cruz prepends the Screen Name Prefix if you create the user in the Site Management Editor. If you make other users for the tenant site manually in Control Panel, you must manually add the prefix to assign them to the correct site.

Adding LDAP Users

You can integrate LDAP with your Cruz installation in the Portal Settings > LDAP tabs. LDAP-added users cannot log into Cruz's Java Client and can only use the web portal.

△ CAUTION:

Before enabling an LDAP server in the Portal, you must create and assign one user from the LDAP server as the Portal administrator. You cannot access the Control Panel without a user with the administrator role. See Creating an LDAP Admin User below for details.

Make sure *Import at Startup* is turned off and in Password Policies, edit the default password policy and make sure that *Change Required* is off.

NOTE:

Noticethatseveral testbuttonsappearinthe LDAP screens, forexample, *TestLDAPConnection*. Use these to validate your entries as you make them.

Click *Add* under LDAP Servers to add the specifications of your LDAP server. After configuring your LDAP server, restart the Cruz server, and attempt to log in as an LDAP user.

LDAP Server Settings

The following settings are required (the values below are examples, only):

Connection

Base Provider URL : ldap://192.168.50.25:389

Base DN : dc =dorado-exchange,dc = oware,dc = net

Principal: dorado@dorado-exchange.oware.net [Principle user must have the necessary administrator rights in Active Directory Server or any other LDAP server]

Credentials: ********

Users

Authentication Search Filter:(sAMAccountName = @screen_name@)

Import Search Filter: (objectClass = person)

User Mapping

Screen Name: sAMAccountName

In the Portal Settings > Authentication > LDAP tab:

Authentication

Enabled

Import/Export

Import Enabled

Import on Startup Disabled

Creating an LDAP Admin User

All users imported from an LDAP server default to the *Poweruser* role. The default Cruz (login/ password: admin/admin) cannotloginto Synergyonce youenableauthentication through LDAP. Therefore, you must manually assign one of the users from the LDAP server as the Portal administrator. An example of an LDAP database user with Administrator privileges:

Screen name: ITAdmin

User password: ITPassword

First Name: Scott

Last Name: Smith

Email: scott@doradosoftware.com

NOTE:

You cannot import users without these five attributes into Cruz from an LDAP source.

Creating user ITAdmin with Administrator role:

1 As an Admin user, *Go to > Control Panel.*

2 Under the Portal category, click *Users*, then click the *Add* button.

3 Fill out the User form with name and email address and so on. Remember: screen name, first name, and email address are required. Synergy LDAP import will not overwrite existing users.

4 When you are finished, click *Save*.

5 A message appears saying that the save was successful.

6 Select the Password, enter password: ITPassword and click Save.

7 Click the *Roles* link. A screen appears showing the roles to which your ID is currently assigned. By default, all users are assigned the Power User role.

8 Remove the default PowerUser role (optional), and add the administrator role for the user, then click *Save*.

Now you can enter LDAP server info.

Stopping LDAP Authentication

1 To stop authenticating through LDAP, log in as the admin user with ITAdmin/ITPassword.

2 In control panel go to *Portal > Portal Setting > Authentication > LDAP* and uncheck the *Enabled* then *Save.*

3 When the portal re-appears, Users can login only with credentials that exist on Synergy database

Implementing DAP

Database Aging Policies (DAP) prevent the Cruz database from filling up by filling up by deleting old records. You can also save designated contents to an archive file on a specified cycle. Database Aging Policies configure which contents to archive, the archive location, and the configuration of that archive file.

To view and manage such policies, right click an item with them (for example, an alarm), or click *Manage > Control Panel,* and under *Redcell* click *Database Aging Policies.*

Policies appear in the *Aging Policies* tab of this screen, with columns that indicate whether the policy is *Enabled*, the *Policy Name, Details* (description), *Scheduled Intervals* and icons triggering three *Actions* (*Edit, Delete* and *Execute*). Notice that the bottom right corner of this page also lets you *Enable/Disable/Execute All* policies listed.

The following are steps typical for implementing DAP:

1 From the screen listing Database Aging Policies (DAP), click *Add Policy*, and select a policy from the displayed list of alternatives.

2 This opens the Aging Policies Editor.

3 In the *Aging Policies > General* tab, specify the name, schedule interval, whether this policy is *Enabled*, and so on.

4. Specify the *Archive Location.*Those listed are the *Repositories* listed on the Repositories panel. You can manage those on that tab.
5. Click the Options tab.
6. Specifyeither the archiving and retention you want, or further specify sub-policies that refine the items archived and specify archiving and retention for those sub-policy elements. Which one you can specify depends on the type of DAP you are configuring.
7. Click *Apply* until the displayed screen is the DAP manager.
8. View/Verify DAP.

   DAP archives information into the specified repo sitory under the installation root. Youcan open archived .xml data with `dapviewer`. Launch this application from a command line after setting the environment with `oware` in Windows or `. ./etc/.dsienv` in Linux.

   Archived data is deleted from Cruz's database. You can verify that by querying whether archived data still exist. Youalso can back up your database if you want to preserve records not yet archived.

Opening an Archive in dapviewer

1. First, make sure you have an archived file. One way to do this is to edit the Events DAP, make sure the archived events go to a directory you can access later, and retain them for zero days.
2. Manually run the Events DAP
3. Open a command shell. Type `oware` in Windows, or `. ./etc/.dsienv` in Linux.
4. Type `dapviewer.`
5. Select the file with the ellipsis (...).

NOTE:

dapviewer opens both compressed and uncompressed files. It does not open empty files. Also: You must have display set to the host running dapviewer if you are running it on a remote host.

6. Click the *Load* button.
7. Examine the archived data.

BackingUp the Database

To back up your database, open a command shell (*Start* > *Run* `cmd`, in Windows), and then type the following at the prompt replacing USERNAME and owbusdb. By default, the database is `owbusdb`, username is `root` and password is `dorado`.

```
mysqldump -a -u USERNAME --password=[name] owbusdb > FILENAME.mysql
```

For example:

```
mysqldump -a -u oware --password=dorado owmetadb > owmetadb.mysql
```

If you have Performance monitors or Traffic Flow Analyzer, you must also back up your stored procedures otherwise they do not get restored when you restore the database. The command line here adds `--routines`. For example:

```
mysqldump -a -u oware --password=dorado --routines owbusdb > owbusdb.mysql
```

This writes the owbusdb to a plain-text file called `FILENAME.mysql` (`owbusdb.mysql` in our examples). This file is a full backup with which you can fully restore your database in case of problems.

Defaults for the database are `oware` (login) and `dorado` (password). These are typically different from the login/password for the application.

NOTE:

To get a rough estimate of a database's size, looking at the size of the directory `\oware3rd\mysql\data.`

Here are the backup commands for all the databases:

```
mysqldump -a -u root --password=dorado owbusdb > owbusdb.mysql
mysqldump -a -u root --password=dorado owmetadb > owmetadb.mysql
mysqldump -a -u root --password=dorado lportal > lportal.mysql
mysqldump -a -u root --password=dorado synergy > synergy.mysql
```

To backup stored procedures too:

```
mysqldump -a -u oware --password=dorado --routines owbusdb >
  owbusdb.mysql
```

Restoring Databases

Restoring from `FILENAME.mysql` is a three-step process. This occurs, again, in a command shell:

1 Drop the database:

```
mysqladmin -u USERNAME -p drop owbusdb
```

or

```
mysqladmin -u USERNAME --password=[password] drop owbusdb
```

2 Recreate the database

```
mysqladmin -u USERNAME -p create owbusdb
```

or

```
mysqladmin -u USERNAME --password=[password] create owbusdb
```

3 Import the backup data

```
mysql -u USERNAME -p owbusdb <FILENAME.mysql
```

or

```
mysql -u USERNAME --password=[password] owbusdb < FILENAME.mysql
```

Here are restoration commands for all the databases:

```
mysql -u root --password=dorado owmetadb < owmetadb.mysql
mysql -u root --password=dorado owbusdb  < owbusdb.mysql
mysql -u root --password=dorado lportal  < lportal.mysql
mysql -u root --password=dorado synergy  < synergy.mysql
```

NOTE:

If you receive the error "Access Denied. Invalid Role for this device" or if the application(s) fails with error indicating that it cannot connect to the device after a network change, There may be a DNS resolution issue. You may still be able to ping or connect to a device by IP address but testing server connectivity by using hostname will confirm or rule out a DNS problem. This would only be a problem if the device(s) werediscoveredorsetto"ManagebyHostname".Ifthesystemwasmigrated toasystemthatisnot using DNS, then name resolution could fail and there would be no connectivity to devices. The solution fromthe Resource Manager is to Right Click -> edit and un-check"Manage by Hostname". This will then default to ManagebyIPaddress. Alternatively, youcan fixname resolution in yourenvironment.

NOTE:

Whenever you upgrade your system and your database is on a separate server, you must run the `dbpostinstall` script on the (primary) application server too.

Setting Up Authentication

Integrating LDAP

You can integrate LDAP with your Cruz installation in the Control Panel > Portal Settings > Authentication > LDAP

CAUTION:

Before enabling LDAP server in Portal, you must create and assign one userfrom LDAP server as Portal administrator. You will not be able to access control panel without administrator role.

Step1:Assignoneuserfrom LDAPserverasPortaladministrator

All users imported from an LDAP server default to the Poweruser role. The default Cruz (login/ password: admin/admin) cannot log into Cruz once you enable authentication through LDAP. Therefore, you must manually assign one user from the LDAP server as Portal administrator. Here is an example of an LDAP database user with Administrator privileges:

Screen name: ITAdmin

User password: ITPassword

First Name: Scott

Last Name: Smith

Email: scott@dellhardware.com

NOTE:

You cannot import users without these five attributes into Cruz from an LDAP source.

Creating user ITAdmin with Administrator role:

1 As an Admin user, Go to > Control Panel.

2 Under the Portal category, click *Users*, then click the *Add* button.

3 Fill out the User form with name and email address and so on. Remember: screen name, first name, and email address are required. Cruz LDAP import will not overwrite existing users.

4 When you are finished, click *Save*.

5 A message appears saying that the save was successful.

6 Select the *Password*, enter password: ITPassword then click *Save*.

7 Click the *Roles* link. A screen appears showing the roles to which your ID is currently assigned. By default, all users are assigned the Power Userrole.

8 Remove the default PowerUser role (optional), and add the administrator role for the user, then click *Save*.

Nowyou can enter LDAP server information. Be patient, your changes may takea moment to take effect.

Step2: Add an LDAP server

In the LDAP tab of the Authentication screen, check the *Enabled* checkbox, then click *Add* under *LDAP Servers* and fill in that screen as appropriate.

Authentication

Enabled

Import/Export

Import Enabled

Import on Startup Enabled

NOTE:

Noticethatseveral test buttonsappearin the LDAP screens, forexample, *TestLDAPConnection*. Use these to validate your entries as you make them.

LDAP Server Settings

The following settings are required (the values below are examples, only):

Connection

Base Provider URL : ldap://192.168.50.25:389

Base DN : dc= test-exchange,dc=oware,dc= net

Principal: test@test-exchange.oware.net

NOTE:

The Principal user must have the necessary administrator rights in Active Directory Server or any other LDAP server

Credentials: ********

Users

Authentication Search Filter:(sAMAccountName =@screen_name@)

Import Search Filter: (objectClass =person)

User Mapping

Screen Name: sAMAccountName

In the Portal Settings > Authentication > LDAP tab

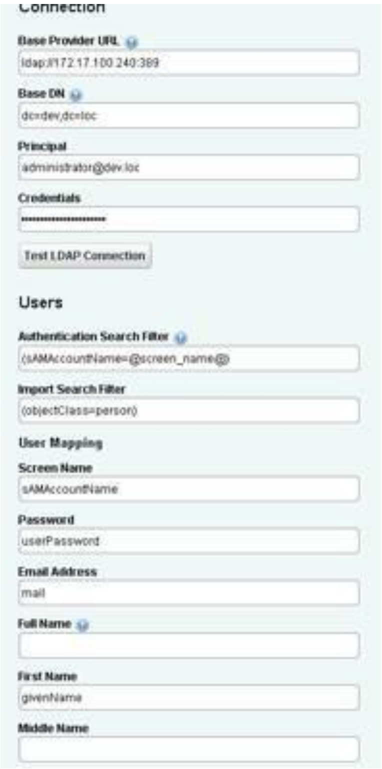

Connection

Base Provider URL: ldap://172.17.100.240:389

Base DN: dc=dev,dc=loc

Principal: administrator@dev.loc

Credentials: ••••••••••••••••••••

Users

Authentication Search Filter: (sAMAccountName=@screen_name@)

Import Search Filter: (objectClass=person)

User Mapping

Screen Name: sAMAccountName

Password: userPassword

Email Address: mail

Full Name:

First Name: givenName

Middle Name:

Step 3: Turn off default 'admin' user's local authentication. (Optional)

By default, user 'admin' able to login with local authentication even when 'LDAP' required was selected.

Toprevent user'admin' to uselocalauthentication, edit thefile.../oware/synergy/conf/server-overrides.properties and add the following line:

```
auth.pipeline.enable.liferay.check =false
```

NOTE:

user will need to rename server-overrides.properties.sample to server-overrides.properties

Step 4: Restart the webserver

Restart the Cruz server, and attempt to log in as an LDAP user.

NOTE:

IfLDAPusersarenotimportedcorrectly,youcancheckthelogunder.../oware/synergy/tomcat-7.0.40/logs/

LDAPand Multitenancy FAQs

The followingare answers to some of the frequently asked questions about LDAP, particularly related to multitenancy (see the User Guide for more about Multitenancy).

Disabling logins after a preset number of failed attempts —Cruz supports this for both local and LDAP users.

Reporting login attempts — Supported from report: User Login Report/Last 30 Days.

All users log in with LDAP —This is supported. Tenant site users must prepend the site prefix. For example: The full screen name for user *admin* in customer site with the prefix *BP*, logs in as *BP-admin.*

Are Passwordsstoredasplaintext?—Passwordsarestoredinencryptedforminthedatabase, even for imported users. LDAP users can authenticate through Active Directory (AD) or OpenLDAP. If you do not want locallystored password, manuallycreate users. Alternatively, import users, then disable import, and change the local passwords so they are different than the one from AD.

Roles and Users — You must locally configure different roles for users within Cruz.

Authentications — By default Redcell authenticates from the local server(s). If you add `auth.pipeline.enable.liferay.check=false (in [installation root]\oware\synergy\tomcat-7.0.40\webapps\social-networking-portlet\WEB-INF\classes\portal.properties)` and enable LDAP required, it uses LDAP to authenticate. Liferay does have multiple entries for AD and OpenLDAP.

Configuring a CAS Server with RADIUS

Cruz does not support RADIUS for authentication directly, however it does support LDAP (see Integrating LDAP on page 83), CAS, NTLM SSO, OpenID, Open SSO and Siteminder. If you are notdoing NTLM/LDAP/Active Directory, Central Authentication Service(CAS) isawidelyused, open source central authentication solution.

NOTE:

For more information on NTLM SSO (Microsoft Single Sign-On Authentication), please reference https://dev.liferay.com/discover/deployment/-/knowledge_base/7-0/ntlm-single-sign-on-authentication

Thisfeatureimportsuserswiththedefault levelofpermissions. You mustmanually alterpermissions and create groups if you want to differentiate between user permissions.

CAS can also use various authentication schemes like LDAP, or RADIUS, so Cruz supports those indirectly. Web applications like Cruz only need to know about the CAS server, not the various authentication protocols CAS uses to provide the final authentication mappings.

One popular CAS Server is available at: http://www.jasig.org/cas

Configure access to CAS in the Portal > Portal Settings > Authentication > CAS tab, which includes a *Test CAS Configuration* button. Other tabs are available here for authentication too, for example LDAP and Active Directory (see Integrating LDAP on page 83 for instructions about how to enable LDAP).

Liferayprovides foundation classes for Cruz's web client. Liferay Wikiinstructions about setting up CAS appear here: http://www.liferay.com/community/wiki/-/wiki/Main/CAS +Liferay +6 +Integration.

The example we tested uses two devices running Tomcat 7.x and Java 6 on one device (DeviceA) and Cruz on the second device (DeviceB). You must access DeviceA using its fully qualified hostname (example: QA002.test.loc, not QA002). Youmustcreate `cas-web.war`for Cruz's CAS server to support this. Instructions about how to do this are on the CAS open source site at wiki.jasig.org/display/CASUM/Best +Practice + -+Setting+ Up+CAS +Locally+using+the + Maven +WAR+ Overlay+ Method. Your preferred search engine may find other instructions for compiling or downloading `cas-web.war` file.

Configuring DeviceA

Follow these steps:

1 Install tomcat 7.x (example: `apache-tomcat-7.0.37-windows-x64.zip`)

2 Insert `cas-web.war` into the `..\tomcat\apache-tomcat-7.0.37\webapps` directory. Start Tomcat (run `startup.bat`in `tomcat\bin` directory). Thisextracts `cas-web.war`, creating the `cas-web` folder with subcomponents.

3 Shut down Tomcat (`shutdown.bat)`

Creating RADIUS configuration setup:

Follow these steps (inserting the correct path when `[path]` appears):

1 Edit the `deployerConfigContext` file located in the `..tomcat\apache-tomcat-7.0.37\webapps\cas-web\WEB-INF` directory.

2 Search for the `RadiusAuthenticationHandler` section of that file.

3 Replace index ="0" with the IP address of the RADIUS server.

4 Replace index = '1' with the global RADIUS server password.

5 We tested a RADIUS server using mschapv2 protocol. If your radius server uses a different protocol replace index = '2' value with the correct RADIUS protocol value.

6 *Save* this file.

Create, Export, Import Certificates using Java

Follow these steps:

1 Run the following from the Java location on your computer (typically under `c:\Program Files` in Windows):

```
..Java\jdk1.6.0_26\bin>keytool -genkey -alias cascommon -keyalg RSA
```

```
..Java\jdk1.6.0_26\bin>keytool -export -alias cascommon -file
 casserver.crt
```

```
..Java\jdk1.6.0_26\bin>keytool -import -trustcacerts -alias cascommon -
 file casserver.crt -keystore "C:\Program
 Files\Java\jdk1.6.0_26\jre\lib\security\cacerts"
```

2 Uncomment `connector port="8443"` section in the `..\tomcat\apache-tomcat-7.0.37\conf\server.xml` file

3 And add keystorefile, keystorepass, truststorefile properties

```
<Connector port="8443" protocol="HTTP/1.1" SSLEnabled="true"
               maxThreads="150" scheme="https" secure="true"
               clientAuth="false" sslProtocol="TLS"
          keystoreFile="C:\users\TestUser\.keystore"
          keystorePass="changeit"
          truststoreFile="C:\Program
 Files\Java\jdk1.6.0_26\jre\lib\security\cacerts"
          />
```

4. In same file comment out `<Listener className="org.apache.catalina.core.AprLifecycleListener" SSLEngine="on" />`
5. Restart tomcat
6. Copy the `casserver.crt`file that wascreated during `keytool -export`fromdeviceA to deviceB

Do the following configuration on deviceB

7. Import the certificate:

```
/cygdrive/c/[path]/oware3rd/jdk1.6.0_45nt64:keytool -import -trustcacerts
 -alias cascommon -file "c:\ssl\casserver.crt" -keystore
 "[path]\oware3rd\jdk1.6.0_45nt64\jre\lib\security\cacerts"
```

8. Add near the end of `set "JAVA_OPTS" -Djavax.net.ssl.trustStore="[path]\oware3rd\jdk1.6.0_45nt64\jre\lib\security\cacerts"` to the `setenv.bat` file located in `..\oware\synergy\tomcat_xxx\bin`

   When finish the line should look like this:

```
set "JAVA_OPTS=%JAVA_OPTS% -Dfile.encoding=%PORTAL_ENDCODING% -
 Djava.net.preferIPv4Stack=%PORTAL_IP_STACK% -
 Dsynergy.https=%ENABLE_SSL% -Dssl.certfile=%SSL_CERTFILE% -
 Dssl.certkeyfile=%SSL_CERTKEYFILE% -Dsynergy.http.port=%PORTAL_PORT% -
 Djavax.net.ssl.trustStore="[path]\oware3rd\jdk1.6.0_45nt64\jre\lib\secu
 rity\cacerts" -Xms%PORTAL_MAX_MEM% -Xmx%PORTAL_MAX_MEM% -
 XX:MaxPermSize=%PORTAL_PERMGEN%"
```

9. On deviceB edit the `portal-ext.properties`file located in `\oware\synergy\tomcatxxx\webapps\root\web-inf\classes`

   Set property settings as follow:

```
live.users.enabled=false
com.liferay.portal.servlet.filters.sso.cas.CASFilter=true
default.landing.page.path=/group/root
company.default.home.url=/group/root
```

10. Restart web service using a command window to run this command:

```
oware/synergy/tomcat-7.0.40/bin/startup.bat
```

11. `Go to > Control Panel` in Cruz
12. Select *Portal* > *Portal Settings*
13. Click on the *Authentication* link on the right
14. Click on CAS tab.
15. Check the *Enabled* check box.
16. Change the login URL to `https://deviceA_hostname:8443/cas-web/login`
17. Change the logout URLto `https://deviceA_hostname:8443/cas_web/logout`
18. Change the server URL to `https://deviceB_IPAddress:8443/cas_web`
19. Click the *Test CAS Configuration* button.
20. If the test passes, click *Save*

21 Touse RADIUSwith Cruz, createusers(no password) that alreadyexiston the RADIUS server in the *Portal > Users and Organizations* portion of the Control Panel.

22 Logout from Cruz.

23 in a web browser go to the URL of your Cruz:8080.

24 A CAS authentication page appears where you enter credentials of users created on radius server.

Setting Up Radius Authentication

The Cruz application supports radius authentication using an external radius authentication server. To set up direct radius support, you need a Cruz user with the same name as the radius username.

Set up direct radius support as follows.

1 Configure the Cruz system to recognize your radius server.

a. Navigate to the installDir/owareapps/installprops/lib/installed.properties file.

b. Open the installed.properties file with a text editor.

c. Add the following properties:

```
com.dorado.server.radius.server=<serverIP>  required
com.dorado.server.radius.port=<radius port#> optional, default =1812
com.dorado.server.radius.secret=<radius secret value> required
com.dorado.server.radius.timeout=<timeout in ms> optional, default =
    1000 ms
```

For example:

```
com.dorado.server.radius.server=192.168.54.137
com.dorado.server.radius.port=1812
com.dorado.server.radius.secret=testing123
com.dorado.server.radius.timeout=1000
```

2 Enable radius authentication.

a. Navigate to the *installDir*/oware/synergy/tomcat-7.0.40/webapps/netview/WEB-INF/classes/portal.properties file.

b. Open the portal.properties file with a text editor.

c. Add thefollowingline:

```
auth.pipeline.pre=com.dorado.nva.auth.CustomRadiusAuthenticator
```

This property causes the Cruz application to authenticate against the radius server prior to logging in through Cruz authentication. To skip Cruz authentication, continue with step 3. Otherwise, continue with step 4.

3 Edit the server-overrides.properties file if you only want to atheniticate against the radius server.

a. Navigate to the *installDir*/oware/synergy/conf/server-overrides.properties.sample file.

b. Rename the server-overrides.properties.sample file to server-overrides.properties.

c. Open the server-overrides.properties file with a text editor.

d. Add the following line:

```
auth.pipeline.enable.liferay.check=false
```

If this line is set to false, the Cruz system authenticates against the radius server only. If set to true (or not set), authentication must pass at both the radius server and local

Cruz authentication.

4 Verify that authentication work as defined.

Configuring Pages and User Access

This section describes adding pages to your installation and configuringrole-based user views. This is a way to manage user access to the Cruz features in a more complex environment. Configuring pages and user access involves:

- Creating Users
- Creating and Rearranging Pages
- Setting Page-Level Permissions
- Setting Portlet-Level Permissions
- Setting Resource-Level Permissions

Pagesdisplayportletsinthefollowingways:

Display Mode	Description
Summary/ Minimized Mode	Any portlet's that have the *Settings* toolbar option (Filters and Max Results) can save/toggle the Current Filter, Max Results, Max Items Per Page, and column choices. See Portlet Tools and Options on page 128. **Note:** The Max Results settings for summary portlets differ from those for maximized/expanded portlets. If you are an Admin and are on the Main portal site, Cruz saves these permanently. If you are a REGULAR user, they are only saved temporarily unless the portlet is on your personal Public/Private pages. See Public/Private Page Behavior on page 34 for details.
Maximized/ Expanded Mode	The *Settings* button in expanded portlets lets you configure displayed columns and their order, and the number of items to display. If the number of items in a list exceeds the maximum specified, a *[limit reached]* message appears next to the number of items listed in the bottom right corner of the page. **Note:** For large list, filters are a more efficient use of computing resources than large maximum settings. See Defining Advanced Filters on page 145 for more about configuring filters.

Creating Users

Create users as follows.

1 Go to the Control Panel.

2 Under the Portal category, click *Users and Organizations*,then click the *Add > User* menu item.

3 Fill out the User form with name and email address and so on.

4 Click Save when finished.

The save was successful message appears.

NOTE:

The expanded form lets you fill out more information about the user.

5 Select the Password, enter a password for the user, and then click *Save*.

6. Click the *Roles* link.

   Ascreenappearsshowing the roles to which your ID is currentlyassigned. Bydefault, all users are assigned the Power User role, and to the User role.

7. Optionally remove the default PowerUser role, add the appropriate new role for the user with the + *Select* link, and then click *Save.*

   You can optionally fill out other details later.

   (You may want to do this step after configuring roles. See Adding and Configuring User Roles/Permissions on page 75.)

8. Select Redcell > Permission Manager.
9. Remove any permissions from the User role you do not want the user to have.

Creating and Rearranging Pages

Create and rearrange pages as follows.

1. Sign in as an administration user.
2. Click *Add* > *Page* from theportal.

   A field appears.

3. Enter the page name.

   That creates a page with a blank title in the doc.

4. Click on that tab page to see it.
5. Click *Manage* > *Page.*

   An editor appears that lets you configure the page, add child pages, and so on.

6. Re-arrange the portal pages from the tree on the left.
7. Click *Save* when configured as needed and then click the X in the upper right corner of this editor.

   Your page should appear in the portal after you refresh it.

8. Click the page label to open any new page and click *Add* > *Applications* to add portlets to that page.

   You can also drag and drop the portlets within the page to rearrange them. The applications under the *Portal* node are open source, and not documented here. The rest are Cruz-connected, and are documented in this guide.

NOTE:

Use the *Search Applications* field at the top of the *Add > Applications* menu to find portlets nested within that menu's categories. The *Portal Applications* and *Global* categories includes generic portlets; the remaining categories are for Cruz portlets.

Setting Page-Level Permissions

Settingpage-levelpermissionsprovideauser/group/role/organizationaccesstoadefined Cruzpage. Set page-level permissions as follows.

1. Sign in as an administration user.

2 Click *Manage > Page.*

3 Expand the Page Layout tree.

This represents the page layout as seen in the portal.

4 Select a page where you want to restrict access.

5 Click the *Permissions* button.

6 Deselect the *View* permission for Guest and Community members.

7 Make sure Owner and PowerUser can still view the page.

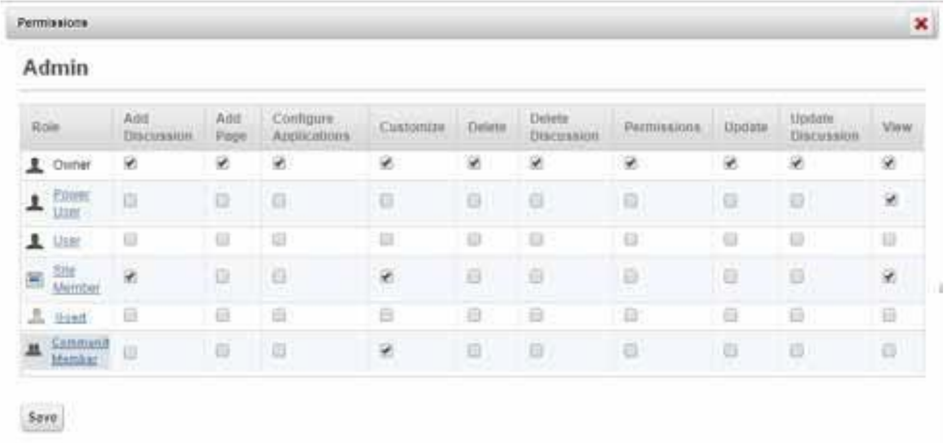

Permissions

Admin

Role	Add Discussion	Add Page	Configure Applications	Customize	Delete	Delete Discussion	Permissions	Update	Update Discussion	View
Owner	☑	☑	☑	☑	☑	☑	☑	☑	☑	☑
Power User	□	□	□	□	□	□	□	□	□	☑
User	□	□	□	□	□	□	□	□	□	□
Site Member	☑	□	□	☑	□	□	□	□	□	☑
Guest	□	□	□	□	□	□	□	□	□	□
Community Member	□	□	□	☑	□	□	□	□	□	□

8 Now select *View* for any other roles you want to give access.

9 Click *Save*.

10 Verify that the user cannot see restricted pages by logging out and then logging in as the new user.

Setting Portlet-Level Permissions

The portlet-level permissions provide a user/group/r ole/organization access to a defined portlet. Configure portlet-level permissions as follows.

1 Sign in as an administration user.

2 Click the Configuration tool (wrench) > Configuration for the appropriate portlet.

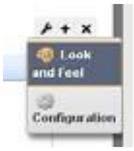

The Managed Resources - Configuration displays a list of permissions.

3 Deselect View permission for Guest and Community members.

4 Makesure Ownerand PowerUserstillhave Viewpermissions.

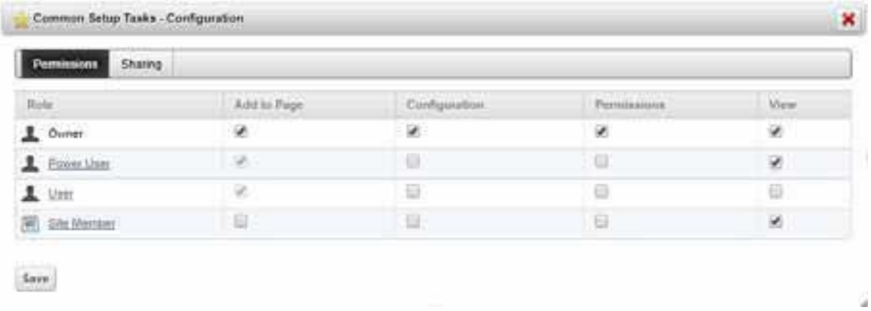

Role	Add to Page	Configuration	Permissions	View
Owner	☑	☑	☑	☑
Power User	☑	□	□	☑
User	☑	□	□	□
Site Member	□	□	□	☑

5 Select View for the relevant roles (for example, Silver Group).

6 Click Save.

7 Verifyyour configuration bylogging out and then log in as Guest or another community member.

Setting Resource-Level Permissions

Theresource-levelpermissions provideauser/group /role/organization access to adefinedresource. Configure resource-level permissions as follows.

1 Create a hierarchy for each Customer from the Hierarchical View Manager portlet.

a. Right-click to select New.

The Creating New Root Hierarchical View window is displayed.

b. Create a hierarchy for the desired Customer, naming and describing it.

c. Click the Authorizations tab.

d. Add authorization for admin and select limited authorization for the desired customer (goldcustomercp, for example).

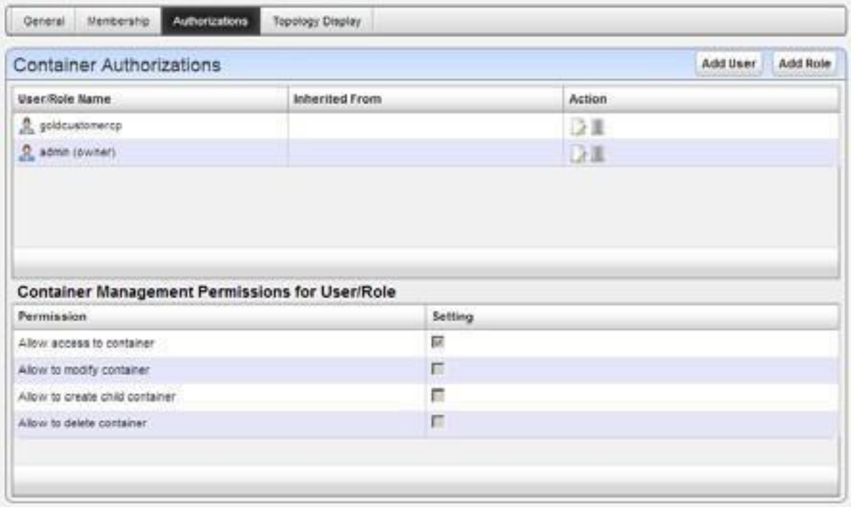

User/Role Name	Inherited From	Action
goldcustomercp		
admin (owner)		

Container Management Permissions for User/Role

Permission	Setting
Allow access to container	☑
Allow to modify container	□
Allow to create child container	□
Allow to delete container	□

2 Configure the hierarchy’s membership by selecting and adding a group ofdevices.

3 Set up a page for Device Level View.

a. Add a Hierarchical Viewportlet to the page ofinterest with portlets for which you want to restrict access.

Currently the Hierarchical View is enabled for the Managed Resources, Alarms, Ports, and Audit Trails portlets.

b. Log out as admin and then log back in as a user with Gold Customer permissions.

c. Confirm your permission configuration is operating on this page.

NOTE:

If you add sub-hierarchies, the admin and goldcustomercp permissions "trickle down," so that users see the lower hierarchies' contents. If you added, for example, silvercustomercp customer, that customer does not see the parent Gold hierarchy.

Registering a License

When it is time to register a license, you have the option to register the license using the Cruz application or using a command line interface.

Using the CruzApplication

Register a license from the Cruz application as follows.

1 Select Settings > Application Configuration Settings > License Management option.

The License Viewer is displayed.

2 Click the Register License tab.

Ablank license list is displayed.

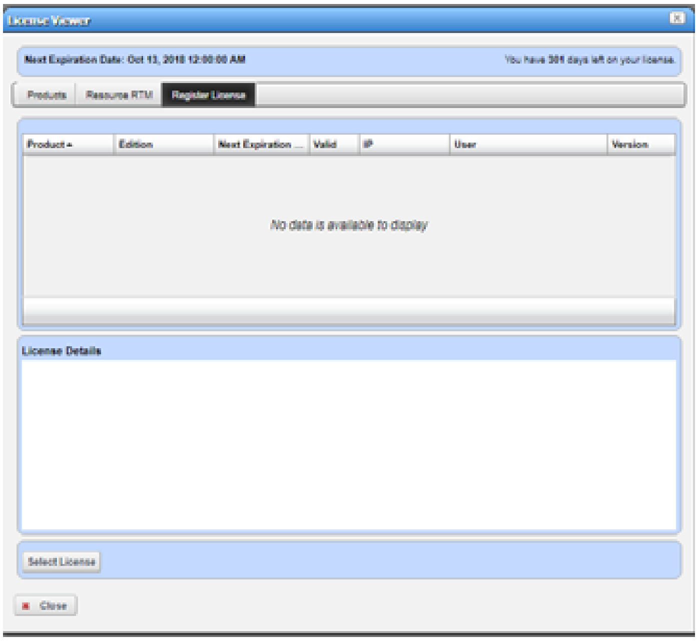

Next Expiration Date: Oct 13, 2018 12:00:00 AM

You have 301 days left on your license.

Product ▲	Edition	Next Expiration ...	Valid	IP	User	Version

No data is available to display

License Details

3 Click Select License.

4 Navigate to the license file, for example:

`license6.5_2000dd_180D_200tfa_ServiceTag_byGen.xml`

5 Select the file and then click Open.

The license information is displayed in the register license list.

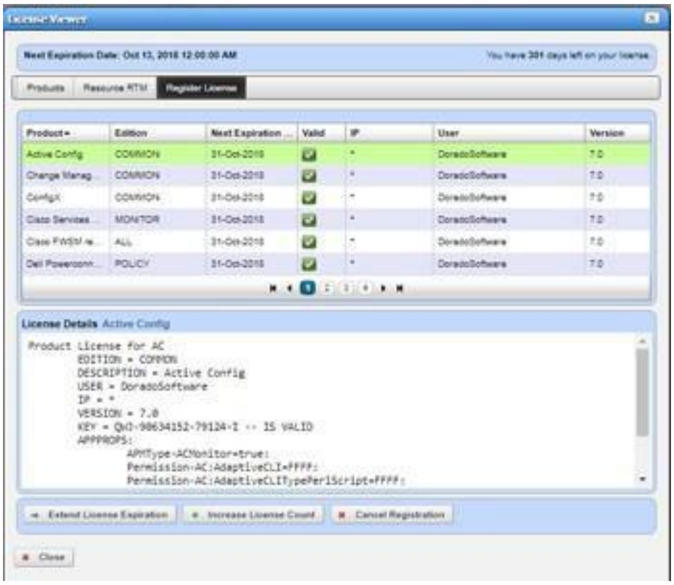

Next Expiration Date: Oct 13, 2018 12:00:00 AM

You have 301 days left on your license

Product ▴	Edition	Next Expiration ...	Valid	IP	User	Version
Active Config	COMMON	31-Oct-2018	☑	*	DoradoSoftware	7.0
Change Manag...	COMMON	31-Oct-2018	☑	*	DoradoSoftware	7.0
ConfigX	COMMON	31-Oct-2018	☑	*	DoradoSoftware	7.0
Cisco Services ...	MONITOR	31-Oct-2018	☑	*	DoradoSoftware	7.0
Cisco FWSM re...	ALL	31-Oct-2018	☑	*	DoradoSoftware	7.0
Dell Powerconn...	POLICY	31-Oct-2018	☑	*	DoradoSoftware	7.0

License Details Active Config

```
Product License for AC
        EDITION = COMMON
        DESCRIPTION = Active Config
        USER = DoradoSoftware
        IP = *
        VERSION = 7.0
        KEY = QHJ-98634152-79124-I -- IS VALID
        APPPROPS:
                APMType-ACMonitor=true:
                Permission-AC:AdaptiveCLI=FFFF:
                Permission-AC:AdaptiveCLITypePerlScript=FFFF:
```

6 Click one of the following registration types:

- Extend License Expiration
- Increase License Count

  See Register License on page 61 if you need more details about these options.

The license is registered.

7 Verify that the license was activated.

a. Click the Products tab.

b. Select the Oware product.

c. Look at the ACTIVATION DATE.

If you selected Increase License Count, the activation date is today’s date. If you selected Extend License Expiration is the day the current license expires.

Using a Command Line Interface

Register a license from a command line interface (CLI) following the instructions for the appropriate operating system:

- Registering a License from Windows CLI
- Registering a License from Linux CLI

Registering a License from Windows CLI

1 Start the command line tool.

2 Enter `oware`.

3 Navigate to the *licenseFilename*.xml file.

4 Enter the following command:

```
licenseimporter licenseFilename.xml
```

If you want to specify an activation date, enter the following command using the activation option (-a*YYYY-MM-DD*):

```
licenseimporter -aYYYY-MM-DD licenseFilename.xml
```

Registering a License from Linux CLI

1 Start the command line tool.

2 Enter `. /etc/.dsienv`

3 Navigate to the *licenseFilename*.xml file.

4 Enter the following command:

```
licenseimporter licenseFilename.xml
```

If you want to specify an activation date, enter the following command using the activation option (-a*YYYY-MM-DD*):

```
licenseimporter -aYYYY-MM-DD licenseFilename.xml
```

Creating an IPv6 Discovery Profile

Cruz system functions independently of the underlying IP protocol communication, so you can expect the Cruz application behavior to be the same regardless of whether resource management is through I Pv4 or IPv6.

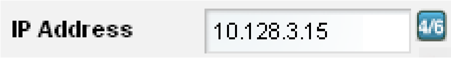

IP Address: 10.128.3.15

The following conditions indicate how to use such IP addresses in Cruz:

- When the blue *4/6* icon appears to the right of the IP address field means both IPv4 and IPv6 addresses are acceptable. When no icon appears, then only IPv4 addresses work.
- When a blue *6* icon appears, then only IPv6 addresses work. Notice that IPv6 does not support capital letters; lower case only.
- If a field supports IPV4/V6 and you enter something like `::ffff:192.2.2.2` then Cruz converts it to a standard address format after you tab off the field, perhapsadding some zeroes you did not enter.

IP v6 support exists for the following:

- Discovery—IPv6 Discovery profiles support single IPv6 entries, not IPv6 ranges and subnets. Cruz discovers devices configured to have an IPv6 management interface. If you discover a device with an IPv6 management interface, the discovery data collected on its ports and interfaces defaults to IPV4 if both IPv6 and IPv4 are configured on the same port or interface.

NOTE:

The ability to exclude a specific IP address on discovery profiles is limited to IPv4 only. IPv6 discovery at this time does not support ranges or subnets so exclusion is not necessary.

- Filters/Filtering/Target selection by IP address
- Editing the device management interface for IPv4 or IPv6
- Network tool portlet includes support IPv4 and IPv6
- Alarms, inventory reports, and other screens support IPv6 too.

CAUTION:

Ifyouhaveadistributedinstallation, inter-servercommunicationmustbeIPv4. Also: IPv6isenabledby default. Addthefollowingpropertytothe owareapps/installprops/install.properties filetodisableIPv6:

```
discovery.supports.ipv6=false
```

The following features are not supported:

- You cannot install to an IPv6 server.
- Distributed installations must communicate with IPv4

Create an IPv6 discovery profile asfollows.

1 Create and name a Discovery Profile (see Discovering Your Network on page 100 for a more complete description).

2 Enter the IPv6 address in the second screen of a discovery profile in the editor.

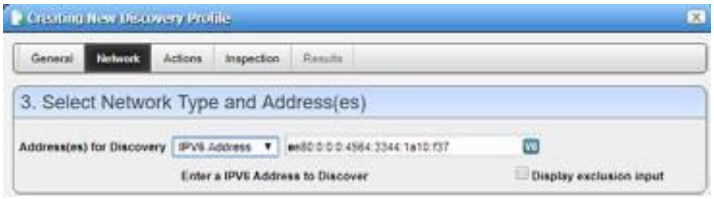
Creating New Discovery Profile

3. Select Network Type and Address(es)

Address(es) for Discovery: IPV6 Address ▼ ee80:0:0:0:4564:3344:1a10:f37

Enter a IPV6 Address to Discover

□ Display exclusion input

Notice that an address like `de80::4564:3344:1a10:f37` becomes `de80:0:0:04564:3344:1a10:f37`, inserting zeros, when you tab off the address field.

3 Configure the discovery as you would like and inspect to validate the device is ready to discover.

Noticethatifyournetworkhasswitcheswithout IPv6 enabledbetween your Cruzinstallation and the device you are discovering, inspection fails.

4 Execute discovery.

The device appears in Managed Resources with the IPv6 interface in its IP Address column.

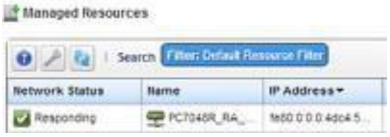
Managed Resources

Network Status	Name	IP Address ▾
✓ Responding	PC7048R_RA_...	fe80:0:0:0:4dc4:5...

It also appears in filters (like selecting device targets)

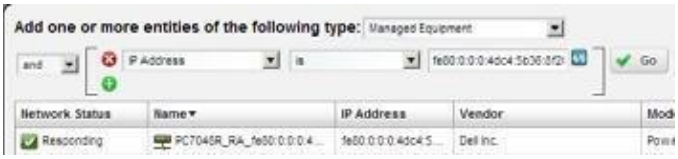
Add one or more entities of the following type: Managed Equipment

and | IP Address | is | fe80:0:0:0:4dc4:5b36:8f2

Network Status	Name ▾	IP Address	Vendor	Mod
✓ Responding	PC7048R_RA_fe80:0:0:0:4...	fe80:0:0:0:4dc4:5...	Dell Inc.	Pow

...and in various Alarm screens

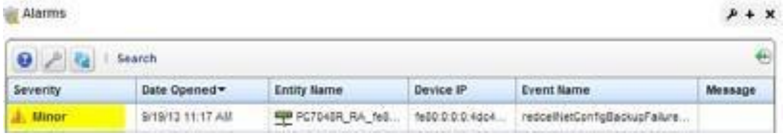
Alarms

Severity	Date Opened ▾	Entity Name	Device IP	Event Name	Message
Minor	9/19/13 11:17 AM	PC7048R_RA_fe8...	fe80:0:0:0:4dc4...	redcellNetConfigBackupFailure...	

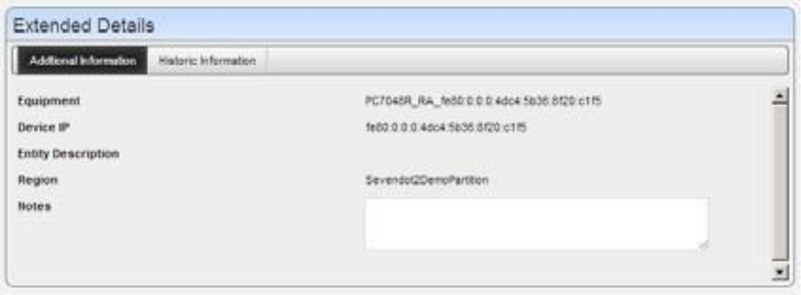
Extended Details

Equipment	PC7048R_RA_fe80:0:0:0:4dc4:5b36:8f20:c1f5
Device IP	fe80:0:0:0:4dc4:5b36:8f20:c1f5
Entity Description	
Region	SevendotZDemoPartition
Notes	

...and in Details panels

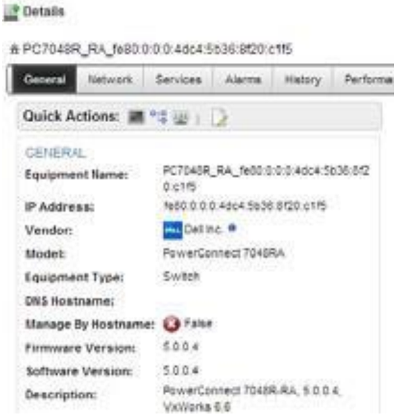

PC7048R_RA_fe80:0:0:0:4dc4:5b36:8f20:c1f5

GENERAL

Equipment Name: PC7048R_RA_fe80:0:0:0:4dc4:5b36:8f20:c1f5

IP Address: fe80:0:0:0:4dc4:5b36:8f20:c1f5

Vendor: Dell Inc.

Model: PowerConnect 7048RA

Equipment Type: Switch

DNS Hostname:

Manage By Hostname: False

Firmware Version: 5.0.0.4

Software Version: 5.0.0.4

Description: PowerConnect 7048R-RA, 5.0.0.4, VxWorks 6.6

...and in Reports...

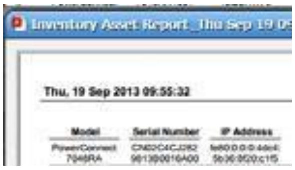
Inventory Asset Report_Thu Sep 19 0

Thu, 19 Sep 2013 09:55:32

Model	Serial Number	IP Address
PowerConnect 7048RA	CN02C4CJ282 98130001GA00	fe80:0:0:0:4dc4:5b36:8f20:c1f5

...and view topologies

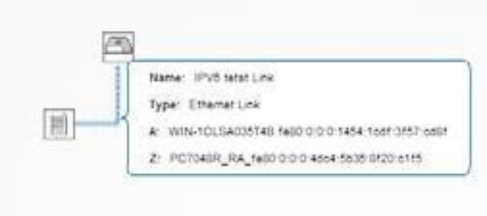
Name: IPV6 tetst Link

Type: Ethernet Link

A: WIN-1OLSA035T4B fe80:0:0:0:1454:1cdf:3f57:cd8f

Z: PC7048R_RA_fe80:0:0:0:4dc4:5b36:8f20:c1f5

...and Performance Dashboards

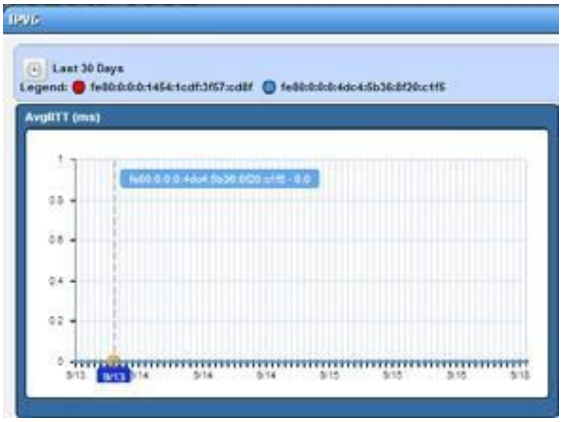

Discovering Resources

Discovery profiles configure equipment discovery for Cruz.

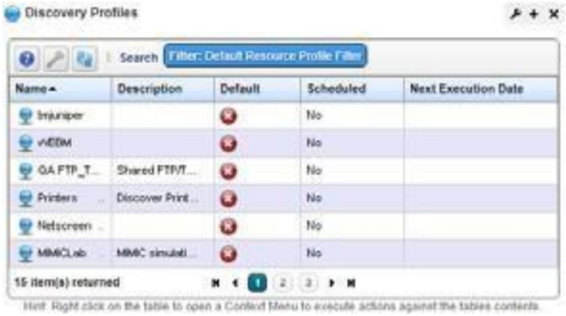

Discovery Profiles

Search Filter: Default Resource Profile Filter

Name ▲	Description	Default	Scheduled	Next Execution Date
lmjuniper		✗	No	
vVEEM		✗	No	
QA FTP_T...	Shared FTP/T...	✗	No	
Printers	Discover Print...	✗	No	
Netscreen		✗	No	
MIMICLab	MIMIC simulati...	✗	No	

15 item(s) returned

Hint: Right click on the table to open a Context Menu to execute actions against the tables contents.

The summary view displays the *Name*, *Description*, *Default* (the green check indicates the default profile), whether the profile is *Scheduled* and *Next Execution Date* for scheduled discovery.

The Expanded portlet adds a Reference Treesnap panel that displays a tree of associations between selected profiles and authentication and tasks that they execute. See Discovery Profiles on page 163 for more about this portlet.

Youcan import discovery profiles to target mu ltitenant domains with a command line importer. Thecommand is `importprofiles`and isin the `owareapps/redcell/bin directory.` This command takes the import file name an argument. The required domains should be available in the Cruz system before import occurs. Before importing discovery profiles to domains, any referenced authentications should be available in the domains or should be imported first by using the`importauths`command (thesamewayyouwouldimportdiscoveryfiles).Inotherwords, you should either manually create authentications for domains or import authentication files using importauths command before importing discovery files to those domains. Example XML files (with the `<customer>` tag for domains) are in `owareapps\redcell\db.`

NOTE:

The date format follows the operating system's conventions for the location and language selected. Restarting the system changes system menus to the new language. If you want to revert back to the original language in Linux, you may also need to update the cache file under `/var/cache/gdm.`

Tuning Discovery Ping

During discovery, Cruz pings the devices specified in profiles. Youcan alter the defaults with the following properties:

```
redcell.discovery.timeout
redcell.discovery.retries
redcell.discovery.defaultport
```

Add these to the `owareapps/installprops/lib/installed.properties` file with the values you want preceded by an equal sign (=).

Discovering Your Network

Thefollowingsteps describe howto discover deviceson your network. Youcanalso edit any seeded authentications and discovery profiles to see what they look like.

1 Right click the Discovery Profiles list and select *New*.

In a multitenant environment, youcancreate pro filesexclusivelyforspecific tenant sites (and the master site). A selector appears with a list of available sites if you have other sites configured.

2 The Discovery Profile Editor appears, with a step-by-step set ofscreens to configure resource discovery. You can navigate through it by clicking the screen tab names at the top, or by clicking the *Next* button at the bottom of the page.

This editor appears with the following panels:

General

3 **General Parameters** —Configure the *Name, Description* and whether this profile is the baseline default. Baseline discovery finds the baseline configuration to compare to later discoveries.

4 **Profile Options** — Select the *Device Naming Format* (how the device appears in resource lists, once discovered), whether to *Manage by* IP address or hostname, and check whether to *Resolve Hostname(s), ICMP Ping Device(s), Manage ICMP-only Device(s),* or *Manage Unclassified Device(s).* This last checkbox determines whether Cruz attempts to manage devices that have no device driver installed. Management may be possible, but more limited than for devices with drivers installed, provided this capability is one you have licensed.

NOTE:

Some packages disable ICMP ping by default.

The Filters (by *Location, Vendor,* or *Device Type*) let you narrow the list of devices discovered by the new profile. As the screen says, this filtering will not have any impact on the processing that occurs during the Inspection step.

Make sure you *Save* profiles you alter, or these selections have no impact when you execute discovery.

Network

5 After you click *Next,* the *Network* screen appears.

Network Type and Addresses —Select the type of entry in the pick list (*IP Address(es), IP v6 Address, CIDR Address, Hostname, Subnet*).

NOTE:

You can specify an IPv4 Address range byseparating the beginning and end with adash. For example: 192.168.1.1-192.168.1.240.

Hover your cursor over the data entry field and the tooltips describe what valid entries look like.

You can exclude IP addresses, or ranges of IP addresses if you check the *Display exclusion input* checkbox and input the addresses you want excluded as you did for those you entered in the *Address(es) for Discovery* field. Such exclusions only apply to the profile where you enter them. To exclude an address or range, use the `com.dorado.redcell.discovery.exclude` property. Examples of how to enter such exclusions appear in the `redcell.properties` file under `owareapps\redcell\lib`. As always, best practice, if you want such properties to persist is to put the property in `owareapps\installprops\lib\installed.properties.`

6 **Authentication** — You can *Create new* or *Choose existing* authentications. (See Authentication on page 114 for more about creating authentications.) Notice that authentications appear with *Edit/Delete* icons and *Up/Down* arrows on their right. The *Up/ Down* arrows order authentications, so Cruz tries the top authentication first, then the next, and so on.

If you have an authentication like `admin/abc123` and one that is identical with an enable-level login/password (`admin/abc123/enable/enable123`), make sure the authentication with `enable` appears first in the list, otherwise, discovery finds the device, but does not access its enable functionality.

△ CAUTION:

If you do not get to the correct level of authentications—for example the “enable” user—then Cruz’s full functionality is not available. The functionality will not be available for backup, restore, deploy, seeded or created actions that require a device enable login, Proscan, some performance monitoring that requires for Cruz to log into the device to retrieve an information using command line interface (CLI). Also, some device information will be missing like firmware versioning, serial number, or port attributes when Cruz cannot retrieve that information using SNMP and needs to use CLI to retrieve that kind of information

NOTE:

Best practice is to avoid special characters, particularly # and > (command line prompts) in device banners so terminal access is unambiguous.

The *Edit* icon opens the authentication editor. Click the arrows to arrange the order in which the software tries credentials (top first). Ordering onlyapplies when two credentials areof the same type.

Actions

7 Youcan configure Actions to run as part ofdiscovery. Bydefault, the actions screen includes the *Resync* action. For more about that, see Configuring Resync on page 104.

Use *Add Action* to select others to enter here. You can also edit parameters (if available), delete and re-order the actions listed here by clicking the icons to the right of them. Cruz executes them in top-to-bottom order.

By default, discovery now automatically updates monitor targets with discoveredequipment. For example, if you have a monitor targeting the dynamic All Dell Devices group, and discover a Dell device, discovery automatically adds the discovered device to the monitor’s target list.

Device discovery initiated by web services does not require an existing discovery profile, however, if adefault discoveryprofileexists, then discoveryinitiated byweb services uses it. If you have updated your system, you must add the Refresh Monitor Targets action to any existing discovery profiles you have created before this default behavior occurs in upgraded discovery profiles.

Youcan change this default bychanging the settings in the `/owareapps/redcell/lib/redcell.properties` file’s `redcell.discovery.taskactivity.order` property. See also *Refresh Monitor Targets for Newly Discovered Devices* on page 412.

Inspection

8 **Inspect Network using your current settings**—This screen lets you preview the discovery profile’s actions and access to devices. If you clicked *Next* rather than *Inspect* at the bottom of the previous screen, click *Start Inspection* to begin the inspection process for selected authentications that validates the device’s credentials.

Notice that the *Inspection Status* fields below listed authentications indicates the success or failure of ping (if not disabled), Hostname resolution, and the listed Authentications.

If the device does not match all required authentications, you can click the *Fix it* icon (a wrench with a red or yellow dot) to edit them for the selected device. You can also click *Test Device*, *Create New*, or *Choose Existing* authentications while in the editor clicking the *Fix it* icon displays the authentication selection panel. The yellow dot on the *Fix it* icon means an optional authentication is missing. A red dot means a required one is missing.

When authentications are unsuccessful, you can remove or edit them in this editor too. Click the icons to the right of listed authentications to do this.

When they test successfully, the authentications appear in a nested tree under the *Discover* checkbox (checked when they test successfully).

9 **Save—** Click *Save* to preserve the profile. You can then right-click it to select *Execute* and begin discovery. If you select *Execute* from the profile editor, Cruz does not save the profile to execute later.

Results

10 **Execute**—Clicking *Execute* begins discovery, confirm you do not mind waiting, and the message traffic between Cruz and the device appears on the *Results* screen.

This is a standard *Audit* screen. See Audit Trail Portlet on page 150 for more about it. The Audit Trail portlet saves its contents if you want to see the message traffic between Cruz and the device(s) later.

11 A message (*Discovery Profile Execute is complete*) appears in the *Messages* at the bottom left of the status bar.

NOTE:

You can also schedule discovery profiles to runperiodically, updating your database with any network changes. For more, see the options description in Schedules on page 153.

12 The devices in your network nowappear in the Managed Resources portlet, and elsewhere (in the Network View topology, for example).

See Discovery Profiles on page 163 for more about these capabilities.

NOTE:

Cruz automatically adds discovered devices to the default ICMP monitor.

LLDP Warnings

Warning messages sometimes appear in the Discovery Audit Trailabout LLDP. This occurs when the device's LLDP information indicates that a component exists with LLDP enabled at a certain IfIndex, but that IfIndex doesn't actually exist in the IfTable.

In other words, the information in the two SNMP tables does not match perfectly, Cruz warns about the IfIndexes that are missing. This incongruity in the tables is rather common and is not normally a problem, but Cruz still displays the warning.

The only time it would be of concern is if the warning claimed that it could not match IfIndex numbers that do exist in the IfTable, otherwise it's just a harmless warning that the LLDP table has some bad data in it.

Incomplete Discovery

If the device is detected and responds to ping but does not respond to actions (for example: Resync or Adaptive CLI), you may have only partially discovered it. Right-click the device in the Managed Resources portlet and select *Direct Access > Telnet*. If that menu option does not exist, the device is typicallypartially discovered with SNMP only. Right-click to edit the device and add a both Telnet Management Interface and Authentication in those two tabs of theeditor.

Configuring Resync

Resync nowcan retrieve, and if it is obsolete, upda te, the information gathered on initial discovery. Fields retrieved/updated include SysDescription, Contact, and Location.

Cruz can also retrieve the device’s SysName to update the device's *Name* field on resync, depending on a system property. By default, this is false, so no name retrieval occurs on resync.

The property determining this behavior is in `owareapps\ddbase\lib\ddbase.properties`. Here is how that property looks in `ddbase.properties`:

```
##Update Device Name on resync - 'false' turns this behavior off.
##Other options: sysname_ip , hostname_ip, sysname, hostname, and ip
##For example, to set this to use sysname + IP naming format, use
##com.dorado.devicedriver.base.updateName=sysname_ip
com.dorado.devicedriver.base.updateName=false
```

Best practice is to override the default in `owareappse\installprops\lib\installed.properties`.

△ CAUTION:

This property overrides the discovery profile’s *Device Naming Format* convention selected. This can be useful if you want to force discovery to use a particular naming convention regardless of how others may have configured the Discovery Profile that initially retrieves information about the device discovered. It may be less-than-useful if a Resync action that follows or is part of the discovery process provides naming you do not want.

If you elect the default, resync still updates SysDescription, Contact and Location, but does *not* update SysName (*Name*).

If you have selected *Manage by Hostname* in your discovery profile, then resync will also retrieve any IP address changes, for example in a network with DHCP. Any override to resync’s retrieved *Device Naming Format* may also change the device’s *Name* when resync occurs if the IP address has changed, but it does not override the *Manage by Hostname* selection, and Cruz still keeps the originally retrieved hostname to refer to the device, even though it may not appear in the Managed Resources portlet and elsewhere.

Finally, you can manually override all retrieved *Name* information. Right-click a device and selecting *Edit*, then make the alterations. See Resource Editor on page 184. Such edits only alter Cruz’s database, not data on the device. To alter data on the device, you can create an Adaptive CLI. See .

Resync overwrites any manual alterations if you make the resync property anything but the default.

Zero-Touch Provisioning and Auto-Discovery

Zero-touch provisioning is a process through which devicescan beautomatically configured and provisioned. Auto-discovery is a related process through which Cruz automatically discovers unmanaged devices that have been configured to send traps.

To enable zero-touch provisioning within your network, you will need an external DHCP server. This server needs to be configured to automatically provision new devices with a basic configuration file. You can include any basic con figuration settings you want within this file.

If you want to enable auto-discovery, then you need to configure the devices in your network to send traps to Cruz and you will also need to log into the Cruz web portal and follow a few configuration steps in order to activate this feature. Note if you want to use zero-touch provisioning in conjunction with auto-discovery, then your basic configuration file that is provisioned to the device will need to include a setting that tells the device to send traps to the Cruz server. However, auto-discovery can be used independently of zero-touch provisioning and vice-versa.

To activate auto-discovery of network devices within Cruz, follow these steps:

1. Log into the Cruz web portal.
2. Click on the Discover tab.
3. Within the Discovery Profiles portlet, find the entry named “Device Auto-Discovery”. Thisis the discovery profile that will be executed when a trap is received from an unmanaged device, but it needs to be configured properly first.
4. Right-click and Edit this entry.
5. You might want to change some of the settings on the General tab. Please see the Discovery Profiles section for more information about this edit screen and the available options.
6. Click on the Network tab.
7. Note that unlike the other discovery profiles, wi thin this record, the IP Address(es) field is intentionally left blank because the IP address of the target device will come from the source IP address of the trap that is received from the unmanaged device.
8. Youwill need to configure the credentials within the device(s) that you wish to auto-discover. If youare usingzero-touch provisioning, then youcould include authentication credentials in the basic configuration file that is provisioned to the device.
9. Thisdiscoveryprofilewascreatedwithplaceholdersfortheauthentication credentials, asseen in the “Select Authentication” list. You will need to remove these entries from the list and create new entries that match the authentication credentials that the device is configured with.
10. Save the discovery profile.
11. Youwill now need to enable the event processing rule that will execute this discovery profile so that it will be triggered when a trap is received from an unknown device.
12. Expand the Alarms menu item and click on Definitions and Rules.
13. Within the Event Processing Rulesportlet, find theentry named “Execute Auto-Discovery when a trap is received from an unmanaged device”.
14. Right-click and Edit this entry.
15. Everything within this event processingruleshould alreadybeconfigured properly so that the only thing you need to do would be to check the Enabled box and the click Save.

16 Once the event processing rule is enabled it will be triggered any time a trap (or syslog) is receivedfroman unmanaged device. Ifauto-discovery is usedinconjunction with zero-touch provisioning, and if all the pieces are configured properly, then Cruz will automatically discover a device shortly after it has been provisioned by the DHCP server.

The Managed Resources portlet displays all discovered devices.

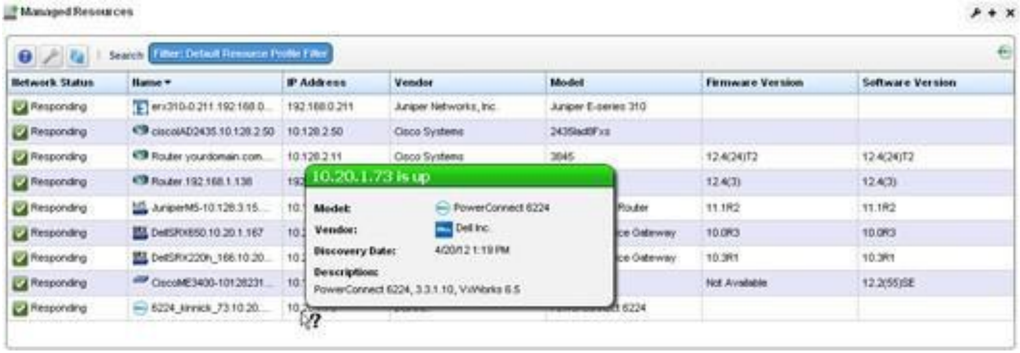

Network Status	Name ▾	IP Address	Vendor	Model	Firmware Version	Software Version
✓ Responding	erx310-0.211.192.168.0...	192.168.0.211	Juniper Networks, Inc	Juniper E-series 310		
✓ Responding	ciscoIAD2435.10.128.2.50	10.128.2.50	Cisco Systems	2435iadBFxs		
✓ Responding	Router.yourdomain.com...	10.128.2.11	Cisco Systems	3845	12.4(24)T2	12.4(24)T2
✓ Responding	Router.192.168.1.138	192			12.4(3)	12.4(3)
✓ Responding	JuniperM5-10.128.3.15...	10.1		Router	11.1R2	11.1R2
✓ Responding	DellSRX650.10.20.1.167	10.2		ce Gateway	10.0R3	10.0R3
✓ Responding	DellSRX220h_166.10.20...	10.2		ce Gateway	10.3R1	10.3R1
✓ Responding	CiscoME3400-10128231...	10.1			Not Available	12.2(55)SE
✓ Responding	6224_kinnick_73.10.20...	10.		t 6224		

10.20.1.73 is up

Model: PowerConnect 6224

Vendor: Dell Inc

Discovery Date: 4/20/12 1:19 PM

Description:
PowerConnect 6224, 3.3.1.10, VxWorks 6.5

Resetting a Password

You can reset a user's password two ways. One is to login as admin and change the user's password in Portal Settings > Users and Organizations. For additional information please refer to Portal > Users and Organizations on page 33.

For the second method, users themselves can request an email be sent to them with instructions to set a new password. Follow the steps below.

1 Login fails. At the bottom of the login screen is the Forgot Password link.

2 A prompt appears for user to enter a Screen Name.

3 A prompt appears to enter the answer to the reminder question (their Father's middle name) that they set when logging in the first time.

4 After entering the correct answer for their account, Cruz sends an email to the user’s email address. E-mail for password reminders/resets requiressetting up the fields in Control Panel > Server Administration > Mail, not the SMTP Configuration which is for Cruz-originated e-mails.

After entering an incorrect answer, a request failed screen appears, with another chance for entering a correct answer.

5 The e-mail provides a link where the user can enter a new password and confirm it.

Deploying Add-On Capabilities

Cruz add-on capabilities come in the following forms:

- Deploy Updates
- Extensions
- .ocp and .ddp files

These add-on capabilities do not require a complete re-installation of the application. If you are upgrading to an entirely new Cruz package, refer to the *Cruz Installation Guide* for the upgrading from a previous version instruct ions. If you are updating your operating system, refer to the *Cruz Installation Guide* for instructions.

The following sections describe how to update your initial system with them.

Add-On Capability	Description
Deploy Updates	Updates to Cruz can come in .war files. For example, a new helpset (nvhelp.war), that updates the information about the program. To deploy such files, copy them to the *installDir*\oware\ synergy\deploy directory. A few minutes later, the Cruz application deploys them.
Extensions	Extended capabilities for the Cruz application may appear in .jar files. For example, the synergy-msp.jar file. To deploy these, copy the file into the *installDir*\oware\synergy\extensions directory.

.ocp and .ddp files	Device drivers and additional application capabilities come in files with the .ddp and .ocp extensions, respectively. These install automatically during the full Cruz installation when they are in the owareapps directory. To install them after your system is already up and running, use the following command line programs:
	`ocpinstall -x [filename.ddp or filename.ocp]`
	`ocpinstall -l [filename.ddp or filename.ocp]`
	`ocpinstall -s [filename.ddp or filename.ocp]`
	Note: You must install these to all application servers in a distributed environment.

Using Device Drivers

For complete communication with devices, the application requires a device driver. For example, to communicate with Dell EMC devices, you must have a Dell EMC driver installed. That does not mean you cannot discover and communicate with other vendors' devices without a driver installed. See. ocp and .ddp files on page 107 for driver installation instructions. The following sections include discussions of some of these drivers:

- Base Driver
- Supported PowerConnect Models
- Windows Management Instrumentation (WMI) Driver
- Web-Based Enterprise Management (WBEM) Driver

Base Driver

Ifyou have no driver installed, the application provides the following functionality, depending on the devices' supporting and providing data from the SNMP system group (sysDescr, sysObjectID, sysUpTime, sysContact, sysName, sysLocation) and the ifT able, which provides list of device interface entries from the RFC1213-MIB. The Cruz application also depends on the entPhysicalTable in the ENTITY-MIB, which provides a list of physical entities contained on a device. If the device does not support the ENTITY -MIB, then the Cruz application bases sub-component creation entirely on ifT able contents.

Confirm that a device is not part of those supported by installed drivers when part of its OID is 3477.

Functionality	Description
Top Level Resource	The application creates top level resource for discovered devices with the following attributes: Equipment Name, Description, IP Address, Location, Contact, Vendor, Model, System Object Id, Date created, Creator, Discovery date, Last Modified.
Subcomponents	The application creates subcomponents (modules, ports, interfaces, power supplies, fans, and so on) for discovered device based on contents of ent Physical Table.
Port/Interface Attributes	The application sets Port/Interface Attributes depending on port/interface type: Name, Port Description, MAC Address, Administrative State, Operational State, Port Type, Speed, Encapsulation, Operation Type , Switch Mode, CLI Name , If Index, Port Number, and Slot Number.
Direct Access	SNMP and Ping (ICMP) are enabled.
Monitors	The application automatically adds discovered device instances to the Default ICMP Monitor to indicate their Network Status. Support for SNMP based performance monitors using discovered port s and interfaces as targets is al so possible. For example, Bandwidth Utilization.
Reports	You can execute reports like the Port Inventory Report or Device Inventory and results should include discovered device and device port entities.
Network View	Discovered devices and their sub-components appear, regardless of whether a device driver exists for them.
Events	The application supports standard MIB-II traps for discovered device and or sub-components. For example, linkUp, linkDown, coldStart, warmStart, and so on.
MIBs	The application can import MIBs for use within MIB Browser and performance monitoring so you can query dev ice specific OID values on discovered device.
Hierarchies	Depending on the licensing, device and or contained sub-components are selectable and manageable in filters and portlets like Hierarchical View.
Links	You can manually create Links using discovered device or device subcomponents as end points which are then visible in Network View.
Attributes	You can manually populate or modify device/port attributes. For example, Serial Number, Firmware Version, Port Type, Notes etc. Attribute values should then be included in reports based on a given report template.

Supported PowerConnect Models

Refer to release notes for a list of supported devices. From the application portal, select Manage > Show Versions to view information about supported devices and operating systems.

Windows Management Instrumentation (WMI) Driver

The Windows Management Instrumentation (WMI) driver currently supports any Windows-based operating system that supports the WMI. This driver must install supported versions of Windows. Refer to the *Cruz Installation Guide* for supported operating systems.

This driver supports global group operations. Discovery may display benign retry warning messages in the application server shell or log. You can safely ignore these.

Before installing the WMI driver, make sure that completed all prerequisites. Then, prevent WMI requests from being blocked by configuring the firewall.

Prerequisites

Before installing the application to manage other computers with a WMI drive r, you must download and install the Microsoft .Net framework version 3.0 or later on the application server if it is not already installed there. For complete functionality, the WMI login for this software must be a login for a domain user who also belongs to the administrator group on the WMI device. Both are requirements for any installation managing WMI devices.

NOTE:

Ifyouhavecomplied with the prerequisites for installation and do notneed the basicinstallation instructionsthatappearinthenextsection,refertothemore detailedinstallationinstructionsinthe other manuals for information about how to install Cruz in more complex environments.

The following are common Windows prerequisites:

Credentials — Use administrative credentials to manage the computer system with WMI.

Firewall—Allow those WMI requests that you want to manage as some firewalls installed may block WMI requests. (See Configuring the Firewall.)

License—Make sure that you have the proper licenses installed to discover the devices you want. If you have a Dell-only license and are discovering a non-Dell device, discovery does not work. Or if you have a Dell license for desktop discovery only, you cannot discover a server.

See License Viewer Window on page 56 for more about licenses.

Configuring the Firewall

Configure the firewall between your server and the Internet as follows.

1. Deny all incoming traffic from the Internet to yourserver.
2. Permit incoming traffic from all clients to TC P port 135 (and UDP port 135, if necessary) on your server.
3. Open Port 445 (WMI).
4. Permit incoming traffic from all clients to the TCP ports (and UDP ports, if necessary) on your server in the Ports ranges specified.
5. Permit incoming traffic on all ports where the TCP connection was initiated by your server if you are using callbacks.

   WMI queries succeed only if you add the user account to a local admin group. Refer to the Microsoft knowledgebase articlesfor the wayto do this. Forexample, Leverage Group Policies with WMI Filters .

   For user rights for WMI access, see: www.mcse.ms/archive68-2005541196.html

   See also: *Service overview and network port requirements for the Windows Server system* (support.microsoft.com/kb/832017/)

Web-Based Enterprise Management (WBEM) Driver

The Web-Based Enterprise Management driver currently supports operating systems supporting the Web-Based Enterprise Management interface (WBEM).

WBEM is always installed on the following operating systems versions, and later:

- Red Hat Linux and/or CentOS 6.2, 6.4
- VM Ware (ESX) with WBEM installed.

Youcan install Web-Based Enterprise Management on some other systems if they do not already use it but monitored devices must have this installed.

NOTE:

ToverifyWBEMisrunningonyoursystem, runthefollowingcommand:`ps -e | grep cim`. You should see a process labelled `cimserver`.

Installing WBEM on Red Hat

You can download and install WBEM support for Red Hat Linux. For example, for Red Hat 6.2, a release for WBEM is `tog-pegasus-2.12.0-3.el6_4.x86_64.rpm`. This is what you need to download once you have logged into the Red Hat network.

Install this as follows:

Install: `rpm -ih tog-pegasus-2.12.0-3.el6_4.x86_64.rpm`

Upgrade: `rpm -Uh tog-pegasus-2.12.0-3.el6_4.x86_64.rpm`

To determine if wbem is running, run `ps -ef | grep cimserver` in a shell.

To start | stop | get status of the WBEM service:

```
tog-pegasus start | stop | status"
```

WBEM Prerequisites

The following are common prerequisites:

Credentials—WBEM credentials have a role in discovering the device. Your system must have access to the computer using Administrative only credentials. These are the same credentials as the user installing WBEM on the device.

Telnet/SSH credentials are necessary for other supported applications.

Forfullfunctionality, this WBEM devicedriverrequiresadministrative (root) access. Many devices may only allow root logins on a localconsole.

In such cases, configure the Telnet/SSH authentication for these devices to login as a non-root user—and, in Authentication Manager, enter `su` in the *Enable User ID* field and enter the root user's password in *Enable User Password* in that same authentication. This enables full device management functionality with root access.

Credentials for Telnet/SSH should have a privilege level enough to stop services and to restart the computer system.

Firewall—Some firewallsinstalled on thecomputer mayblock WBEM requests. Permitaccessfor those you want to manage.

License—Make sure you have the correct WBEM driver license installed. Licenses come in the following types:

- Major Vendor by Name - Such as Dell, Compaq, HP, Gateway.
- Server/Desktop individual license support.
- Generic computers - non-major vendors.

- ALL - this gives the driver all capabilities for any computer system.

NOTE:

If you discover an Amigopod host that does not have its SNMP agent turned on, Cruz labels it a WMI or WBEM host rather than an Amigopod host.

Secure WBEM Access

Some monitoring capabilities require root access, even if you securely log into the Linux host. In this case, when configuring a secure (SSH) lo gin, configure a telnet authentication with `su` as an *Enable User ID*, and the root user’s password as the *Enable Password.* For other WBEM access, configure authentication as an HTTP/HTTPS login/password, and select WBEM as the protocol after you have selected the WBEM authentication.

VMWare ESX and KVM Controller Support

Basicsupport for management of VMWare ESXand KVM Controller devices has been added. The controllers are discovered/managed via WBEM protocol and require WBEM and SSH authentication protocol at discovery time. VMs will appear in the controller's reference tree but can also be discovered standalone.

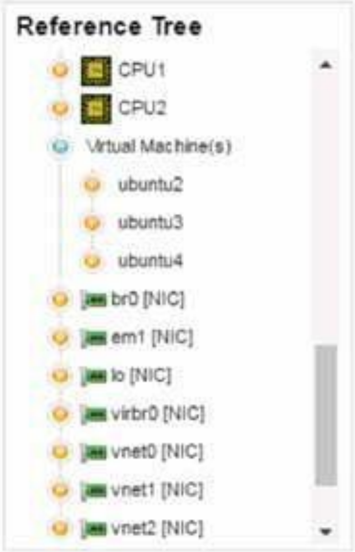

Example Reference tree for a KVM Controller showing three hosted VMs

Supported Functionality for VMWare ESX and KVM Controllers

VMWare ESX and KVM Controller devices support the following functionality:

- Discovery/Resync
- Limited KPI Support
- Direct Access/Terminal
- ACLI Support
- VM Management Actions

VM Management Actions

VMWare ESXand KVM Controller deviceshavespecialized VM Management Actions that are available to them. These actions are accessible via the "Actions" menu that is available on a Managed Resource.

The following actions are supported for managing hosted VMS:

- Start
- Stop
- Suspend
- Resume
- Reboot

NOTE:

In order for some actions to function, the target hosted VM must be configured correctly. For VMWare ESX devices, please see: https://www.vmware.com/support/developer/vcli/ and ensure that the vSphere CLI tools are installed on the target guest VM. For KVM devices, please see: http://virt-tools.org/learning/start-stop-vm-with-command-line/, and ensure that the target guest VM is configured to respond to ACPI requests.

Optimizing Your System

Overriding Properties

Youcan fine-tune various features of the application. Rather than lose those changes when you upgrade or patch, best practice is to override changes.

Todo this for the Web portal, first rename the provided file `\oware\synergy\conf\server-overrides.properties.sample` to `server-overrides.properties`, and enable the properties within it by uncommenting them, and altering them to fit your needs. The comments in this file provides more information.

You can also override application server-related properties in

`\owareapps\installprops\lib\installed.properties.`

Both of these properties files remain as you previously configured them if you install an upgrade, but upgrades overwrite the `server-overrides.properties.sample`, so keep a copy if it has anything you want to preserve.

Screen names — One possible configuration property (`com.synergy.validation.screenmame.min.length`) specifies a minimum length for user screen names. For the existing user base then any screen names that are shorter than the value must change to the required length on the next edit/save for that user.

Tuning Memory (Heap & Portal)

Youcan adjust the memoryfootprint ofany installed server's virtual machine (VM) byconfiguring it in the Heap configuration installation screen that appears during most package installations. Within limits, using more memory, if it is available, generally means better performance.

Launching a server without sufficient memory produces the following error: `Error occurred during initialization of VM Could not reserve enough space for object heap.`

You can re-set these after installation too, with the following properties in `\owareapps\installprops\lib\installed.properties`:

```
oware.server.min.heap.size=3072m
oware.server.max.heap.size=3072m
```

For Windows and Linux valid settings range from 512m to the limit of available RAM minus operating system needs.

While you can enter any number within these constraints, the following are values that are supported during upgrade. Other values are ignored during upgrade and you must choose again from supported value list during installation/upgrade.

Portal Memory Settings

To manually change Cruz web portal heap settings, change the `setenv.sh` (Linux) or `setenv.bat`(Windows) file:

```
set "PORTAL_PERMGEN=512m"
set "PORTAL_MAX_MEM=3072m"
set "PORTAL_INIT_MEM=768m"
set "PORTAL_32BIT_MAX_MEM=768m"
```

These filesare in the `Tomcat***/bin`directory. After you change their settings, for Linux, restart the portal service to apply new memory settings. In Windows, besides updating `setenv.bat` you must run `service.bat update`in that same directory.

You can increase these to even higher figures if your system has the memory available.

NOTE:

Make sure only one Tomcat process is running, otherwise your web server may exhibit poor performance.

Memory Limits Advice

RAM size on hardware can increase virtually without limit. However, if you configure application server so it does not use half of the configured Heap, then having such a large Heap degrades performancesince Java(thissoftware'sprogramminglanguage) scansand sizesgarbagecollection with the pre-allocated large Heap in mind.

That is not to say servers cannot have large amounts of memory. As the applications goes into productionand usagegrow, larger RAMlets you adjust to meet demand as it grows. Having asmall amount of RAM does not allow for growth when needed.

Another thing to remember: Suppose a host has 32GB for an Application Server. Say the Application Server Heap is 28GB. This limits the amount of Threads you can run simultaneously in Thread Pools as well as I/O forking. Every time Java executes a thread, it allocates memory outside of its VM for native calls Since the server only had 32GB and the operating system must use some, very little remains for these processes.

Best practice: Lower the Heap Memory in favor or leaving some more available to the operating system so you can take advantage of more threads if you have the CPU cycles.

Too much Heap RAM impacts only excessive garb age collection which can equate to application pauses as garbage collection moves memory around. Application pauses degrade performance.

Tuning Application Features’ Performance Impact

Resync, Performance, and so on use different pools than backups, so limiting the size of the backups pool would have little impact on resync since these applications does not compete for resources.

If you need to configure these, you can configure pool sizes using properties in `installed.properties`. On startup the application server creates a file called `mbean_attr_overrides.template` in `owareapps/installprops/lib`. This includes text descriptions of Mbeans and their settable properties. It allows you to copyproperties you want to override and add them to the existing `install.properties`, so the settings persist even if you upgrade the software. This replaces a previous tuning method that required editing `mbean-settings.xml` and did not persist pastupgrade.

For more about performance settings for monitors, see Understanding Performance Monitoring on page 351.

Authentication

For successful discovery of the resources on your network, Cruz requires authenticated management access to devices. Toget such access, you must provide the correct SNMP community strings, WMI login credentials, and any other command-line (Telnet/SSH) or browser (HTTP/HTTPS) related authentication, and SNMP must active on devices, if that is not their default. Some devices require pre-configuration to recognize this management software. Consult your network administrator or device manuals for instructions about how to enable them and authorize Cruz as the management console. See Authentication on page 114 for more about authentication.

△ CAUTION:

If you do not get access to the deepest level of authentications—for example the “enable” user’s—you cannot access all of Cruz’s functionality.

MySQL Resizing, Starting and Stopping

MySQL Database Sizing— Installation includes the chance to select a size for your embedded database. This should reflect expected use and should be small enough that you leave enough RAM for the application and operating system (at least 4G, typically).

NOTE:

The default MySQL login command line is: >mysql -u root --password=dorado

By default, installation optimizes the embedded database for the minimum hardware requirement. This may not be sufficient for some environments when your database size grows. You can set the database size during installation, and further tune performance parameters in `..oware3rd\mysql\[version number]\my.cnf`. Have your MySQL operational expert review the links cited below to determine the best values for your environment.

1. `innodb_buffer_pool_size=512m` to `16382m`

   Best practice is to make your buffer pool roughly 10% larger than the total size of Innodb TableSpaces. You can determine total tablespace size with the (free download) MySQL Workbench application.

   If your database size is 30G, ideally have a buffer size of 33G or more. You can also investigate limiting database size or consider adding extra RAM. For dedicated database server, we recommend 70%-80% of system server's RAM, for example use 16G of RAM for a server with 24G RAM total.

To avoid operating system caching what is already cached by this buffer you may have to make additional adjustments. This is not necessary on Windows, but for Linux you need to set `innodb_flush_method=O_DIRECT`.

You may want to make MySQL to use Large Pages for allocating Innodb Buffer Pool and few other buffers. Tuning your VM to be less eager to swap things with `echo 0 > /proc/sys/vm/swappiness` is another helpful change though it does not always save you from swapping.

The optimal setting for Inno DB buffer is to have buffer pool hit rate of 1000/1000)

```
mysql> SHOW ENGINE INNODB STATUS\G
------------------------------
BUFFER POOL AND MEMORY
--------------------------------
Buffer pool hit rate 1000/1000
```

You may need to modify system settings, increase or decrease application server heap, web server heap, and innodb buffer to fit your needs. This depends on whether you use the webserver heavily.

2. `innodb_log_file_size = 256 M`to `1024m`

   A larger file improves performance, but setting it too large will increase recovery time in case of a crash or power failure. Best practice is to experiment with various settings to determine what size is best for performance.

   To change the log file size, you must move existing the log files named ib_logfile0, ib_logfile1, and so on. See Changing InnoDB Log Files in MySQL on page 116 for step-by-step instructions. The database may not start if you configure a log file size mismatch.

3. `max_connections=100` to `1000`

   Best practice is to configure 200 or more connections per server (application server + web server), especially if you are adding more servers.

   The number of connections permitted in this version of MySQL defaults to 100. If you need to support more connections, set a larger value for this variable. Windows is limited to (open tables × 2 + open connections < 2048) because of the Posix compatibility layer used on that platform.

   Log in to mySQL to check current settings:

   ```
   mysql -u root --password=dorado
   mysql> show variables like 'max_connections';
   ```

   To check open connections:

   ```
   mysql> SHOW STATUS WHERE `variable_name` = 'Threads_connected';
   ```

NOTE:

You may need to reduce `table_cache` if you increase `max_connection`.

`4. table_cache = 1024` (increase the default as appropriate).

`table_cache` is related to max connections. For example, for 200 concurrent running connections, you should have a table cache size of at least 200 * N, where N is the maximum number of tables per join in any of the queries which you execute. You must also reserve some extra file descriptors for temporary tables and files.

If the value is very large or increases rapidly, even when you have not issued many FLUSH TABLES statements, you should increase the table cache size.

5. **Monitors:** If you enabled and configured the default SNMP interface monitor, it would typically consume most of the space in owbusdb (thedatabase).

Unless you have reason to do otherwise, best practice is to disable *Retain polled data* on the default SNMP interface monitors. The graphs do not need these data for display. Cruz only uses retained data to derive the calculated metrics attributes. In most cases, saving only calculated data for the default SNMP interface monitor suffices.

For example, if you have 16 polled data attributes and 27 calculated attributes, not saving polled data can reduce the table size about 35%.

You can further reduce the table size if you only poll/save the relevant calculated attributes in the default SNMP interface monitor. To accomplish this, you must remove calculated/polled attributes that you do not want to retain from the monitor configuration, Cruz does not support selectively choosing which attributes to keep. Retained attributes can be all calculated or no calculated.

NOTE:

Best practice is to archive the modified database sizing file somewhere safe. Upgrading or patching your installation may overwrite your settings, and you can simply copy the archived file to the correct location to recover any configuration you have made if that occurs.

If you want to change the size of your database after you have installed it, edit the `my.cnf` file in `/oware3rd/mysql/[version number]/`. Alter the last number on the following line:

```
[path]/oware3rd/mysql/ibdata/ibdata1:1024M:autoextend:max:102400M
```

To start MySQL, run the following in a shell:

```
[path]oware3rd/MySQL/[version number]/bin/mysqld" --console
```

When it starts successfully, the consoleincludes a `ready for connections` message. Without the `--console` parameter, MySQL writes diagnostic output to the error log in its `data` directory.

To stop MySQL, run the following in a shell on the path with MySQL:

```
mysqladmin shutdown
```

Other operatingsystem-specificshutdowninitiation methodsarepossibleaswell: Theservershuts down on Linux when it receives a SIGTERM signal. A server running as a service on Windows shuts down when you shut it down in Windows’ services manager.

NOTE:

The default MySQL login command line is: >mysql -u root --password=dorado

Changing InnoDB Log Files in MySQL

To change the Number or Size of InnoDB redo log Files, follow these steps:

1 If `innodb_fast_shutdown` is 2, set it to 1:

```
mysql> SET GLOBAL innodb_fast_shutdown =1;
```

2 After ensuring that `innodb_fast_shutdown` is not set to 2, stop the MySQL server and make sure that it shuts down without errors (to ensure that there is no information for outstanding transactions in the log).

3 Copy the old log files into a safe place in case something went wrong during the shutdown and you need them to recover the tablespace.

4 Delete the old log files from the log file directory.

5 Edit `my.cnf` to change the log file configuration.

6 Start the MySQL server again. mysqld sees that no InnoDB log files exist at startup and creates new ones.

SNMP in Multi-Homed Environment

Trap listener, Inform listener and all outbound SNMP requests must bind to a specific interface in a multi-homed environment. This interface is considered appropriate to use for all network-facing SNMP activity. By default, this is `localhost`, interpreted as the application's local IP value (the NIC selected at installation time). The following text in `installed.properties` provides a specific IP address to control outbound SNMP interface binding on the local machine:

```
# specific interface used for all NMS initated
# communications to the network
com.dorado.mediation.outbound.address=localhost
```

Include the following and provide a specific IP a ddress to control inboun d (listener) interface binding on the local machine:

```
#
# specific interface used for binding mediation
# listeners such as SNMP trap listener
com.dorado.mediation.listener.address=localhost
```

Events with no corresponding definition appear as alarmsofindeterminate severity. The only way to change behavior of an unknown event in this version would be to locate the missing MIB and load it into the system. This creates the missing event definition(s) needed to specify explicit behaviors.

Maintaining and Repairing Your System

The following describes ongoing tasks that keep your Cruz functioning without interruptions. Some of these take advantage of pre-seeded features, and others take manual intervention.

Mini Troubleshooting

Suggested mini-troubleshooting steps for a balky application that is already installed and running:

1 Refresh the browser. If that does not work...

2 Clear the browser’s cache (Firefox in particular loves persistent old pages), then refresh. If that does not work...

3 Stop and start the browser. If that does not work...

4 Stop and start the Web server

ForWindows, to start the webserver manager: `oware\synergy\tomcat-X.X.X\bin\startsynergy`. For Linux.

`/etc/init.d/synergy start` or `/etc/init.d/synergy stop`

Worth noting: The tray icon for the web server (is "optimistic" about both when the web server has completely started and completely stopped. Youcannot re-start web server when its Tomcat process still lingers. If you lack patience, kill the (large) Tomcat process then re-start web server. The smaller one is that tray icon.

If that does not work...

5 Stop and start application server. Command lines for this:

`stopappserver` and `startappserver`

If that does not work...

6 Delete the contents of the `oware/temp` directory and restart application server. If that does not work...

7 Reboot the host andre-start the application server, web server and browser.

When troubleshooting (or contacting technical support), you may find pertinent information in logs located in the following directories:

```
..\oware\jboss-3.0.8\server\oware\log
..\oware\temp\soniqmq.log
..\app_setup.log
..\db_setup.log
```

Best practice is to run the getlogs script from a command line. It packages relevant logs in a logs.jar file in the root installation directory, and moves any existing copyof the logs.jar file to oware\temp. The logs.jar file compresses all logs necessary for troubleshooting. Read the jar yourself or forward this jar to technical support for help in troubleshooting.

Searching\oware\jboss-5.1\server\oware\log\server.log for "error" is one way to lookfor the rootof application server problems.

NOTE:

If you see errors that say your Linux system has too few threads, make sure that you have set the file handles correctly.

Database Aging Policies (DAP)

DAP policies automatically purge or archive stale data so the database can maintain its capacity. Several pre-defined and pre-seeded DAPs come with Cruz. Youmay need to revise these to fit your system. These start at specific times—see the Schedules portlet for specifics aboutwhen.

DAPsamount to preventative maintenance since they help to maintain the database's capacity. Best practice is to do the following regularly:

1 In the Audit Trail Manager, create a Filter for Creation Date = prior Month and Action = DAP Executed.

2 Review the records for Status – Failed. These indicate that a DAP job failed. If the following DAP jobs execute, no immediate action is required. If any DAPs are repeatedly failing, then consult the troubleshooting document or Cruz support.

3 Review the DAP jobs entries and compare to the scheduled DAP start times. Confirm that audit records are displaying a corresponding audit record for each scheduled execution.

Scheduled Items

Reviewing schedules to ensure that scheduled task is executing as expected.

1 In the Audit Trail portlet, create a Filter for each scheduled action and confirm that the schedule action is successful. The schedule “Type” is the Audit Filter Action.

2 Investigate any failure of a scheduled action.

Database Backup and Administration

Backup your database regularly. Best practice is for the Oracle database administrator to perform a monthly Full backup with Daily incremental backups. Best practice is for the Database Administrator to check monthly to ensure that the application has enough resources. See Backing Up the Database on page 81 and Restoring Databases on page82.

Log Cleanup

The `server.log` files may accumulate over time. Best practice is to purge these once a month. The server.log files are in the directory `<installation root>\oware\jboss-5.1\server\oware\log`. Archivedthese if youneed historicallogdata. Best practice is to store only a maximum of 30 days’ worth of log files.

Installation Logs may also accumulate in `<installation root>\logs`. Best practice is to review this directory monthly and purge it as needed. Best practice is to retain at least 6 months to a year of log data in this directory.

File Cleanup

When you turn on CLI trace, this software generates a log file for every CLI transaction and stores them in the `oware\temp` directory. This directory may accumulate many files if you leave CLI trace on for an extended period. Youcan delete or archive these files. Best practice is to inspect the directory and take the appropriate action at least once permonth.

When backing up/ restoring configuration files or deploying OS images using the internal file server, a copy of the file may remain in `oware/temp`. Youcan delete the contents of this directory since it is auto created.

External file servers used in the production environment may also accumulate copies of transferred files. Best practice is to review space on these file servers monthly to ensure the file server space is adequate.

2

Portal Configuration

The following explains how to customize the Cruz portal. Because this portal can be so flexible, and comes from open source features, this is not a comprehensive catalog of all its features. The following discussion covers only those features significant for using the application. If you already have a good understanding of the portlets and editors, go directly to the tasks you want to perform.

The Audit Trail and Schedules portlets are not part of the portal configuration. However, they are fundamental to the functionality and accessible from many different portlets within the application. Therefore, they are described in this section.

Portal Overview

This section describes the portal window , general portlet information, and features that are common for multiple portlets.

The **portal header** has the comp any logo, product name, and its menu options . The *username* is an active link to the Manage My Account options, where you can configure your name, job title, image, email and so on. You can access these same opt ions from the Go to > Control Panel menu option. The header also includes:

- Sign Out, which logs you out of the Cruz portal. When you sign in, Password Reminder portlet is displayed if your administrator has set thisoption.
- Full Page Toggle, which expands the window full page or returns to a manually adjustable window.
- The page navigation bar , which provides access to the pages, subpages, and portlets.

Thestatusbar opens My Alert/Action Historywindow, displays subscriptionexpirationalert, access to settings, conferencing, anda list of colleagues on the system.

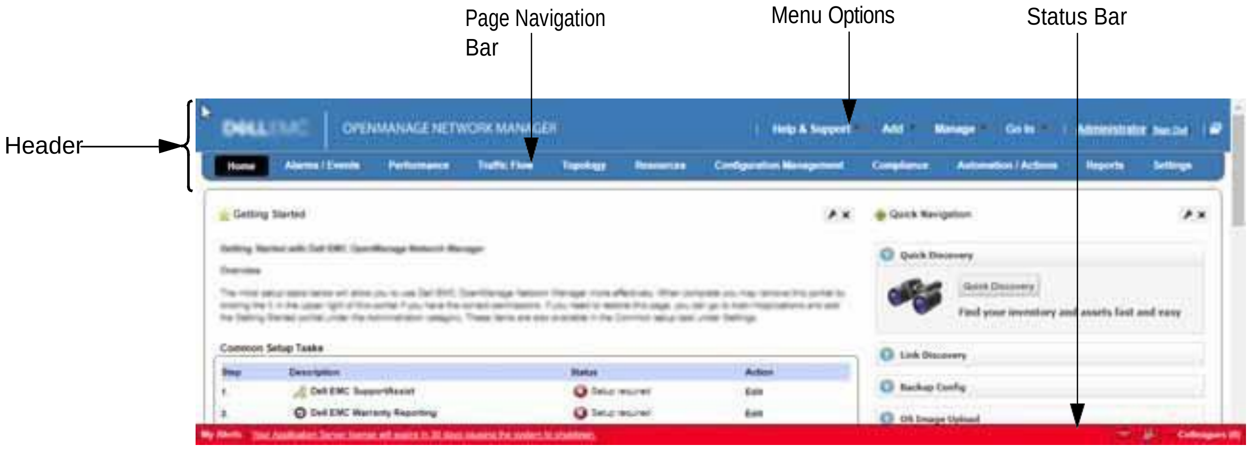

Getting Started with Dell EMC OpenManage Network Manager

Overview

Find your inventory and assets fast and easy

Common Setup Tasks

Step	Description	Status	Action
1	Dell EMC SupportAssist	Setup required	Edit
2	Dell EMC Warranty Reporting	Setup required	Edit

Your Application Server license will expire in 30 days causing the system to shutdown.

Menu Options

The following menu options are available from the Cruz portal header.

Menu	Description
Help	Opens online help and provides access to any other available sources of information.
Add	Let’s you create Pages or add Applications (portlets) to an existing page. Select Add > Page and a field opens on the navigation bar when you enter the page name and then click the checkmark. Select Add > Applications and a list of available applications (portlets) is displayed. Navigate to the portlet, click Add, and the portlet is added to the currently selected page. See General Portlet Information on page 127 for more information about pages, child pages, portlets, editors, and the applications list .

Menu	Description
Manage	Opens the Manage Page window where you alter the page organization or its layout when you select the Page or Page Layout options. From the Manage Page window, you can drag and drop page locations in the tree and add child pages. See Creating and Rearranging Pages on page 91 for instructions. The Show Versions option displays information about your current Cruz installation, such as product details, installed extensions, and driver information. This option is not available from the message board portlet. You can create an HTML version of this information by running the drvrpt (Linux) or drvrpt.cmd (Windows) command from the /owareapps/ddbase/bin directory.
Go to	Provides access to the Control Panel on page 32 and any other pages (public/private) defined by you or your system administrator. When you add a new community, its configured pages appear in this menu too. Creating public pages requires administrative rights. However, users that do not have administrative rights can only configure their private pages. Changes to a page persist after you make them, provided you have the rights to make changes to that page. See Public/Private Page Behavior on page 34 for the details. Best practice is to use multiple pages within the Cruz system rather than multiple tabs. If you have the Multitenancy (MSP) option installed, these options are not available, and MSP offers a different security model.

Page Navigation Bar

The page navigation bar provides access to the pages and child pages installed by default or the pages and child pages you configured from the Manage > Page menu option. The page navigation bar appears horizontally below the portal header by default, or vertically on the left of the screen. The pages that appear on this bar vary depending on which Cruz package you installed.

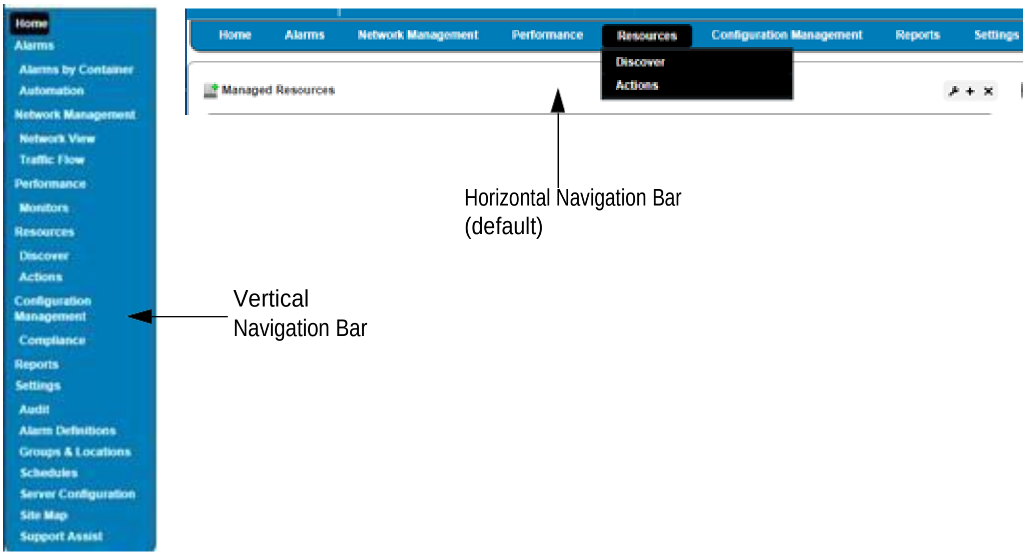

If you narrow your browser window, the navigation bar collapses to the left. If you want the vertical navigation bar as the default, change the regular browsers menu position settings to vertical from the Manage Page window.

Status Bar

The status bar is at the bottom of the portal window. It contains the following elements:

- My Alerts
- Settings
- Settings
- Conferencing
- Colleagues

My Alerts

My Alertsopens the MyAlert/Action History wind ow when you clickon it. The MyAlert/Action Historywindowdisplayscataloged messagesand notifications youreceived, including generated reports. From this window, you can:

- Delete selected messages or notifications by selecting those you want to delete and then clicking Selection.
- Delete all messages and notifications by clicking All.
- Viewadditional information about the message or notification by clicking the magnifying glass.

You can see the portal when the Web server is up, but the application server is not. When the applicationserver starts after the Webserver, an alert appearsin the MyAlert/Action History list announcing that application server is up.

△ CAUTION:
If My Alerts is not receiving messages or notifications, make sure that your firewall is not blocking ports the application uses. See Ports Used on page 979 for a list of ports and more about configuring Linux firewalls.

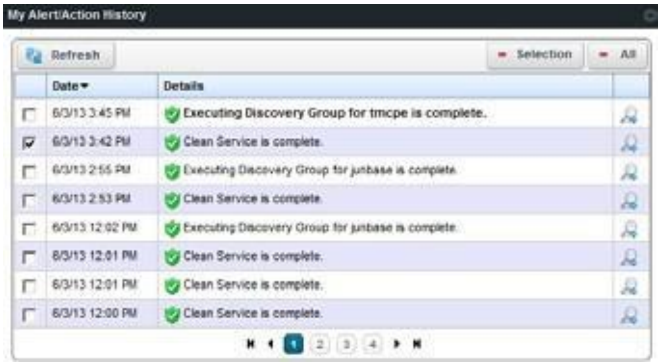

	Date	Details
□	6/3/13 3:45 PM	Executing Discovery Group for tmcpe is complete.
☑	6/3/13 3:42 PM	Clean Service is complete.
□	6/3/13 2:55 PM	Executing Discovery Group for junbase is complete.
□	6/3/13 2:53 PM	Clean Service is complete.
□	6/3/13 12:02 PM	Executing Discovery Group for junbase is complete.
□	6/3/13 12:01 PM	Clean Service is complete.
□	6/3/13 12:01 PM	Clean Service is complete.
□	6/3/13 12:00 PM	Clean Service is complete.

Settings

The Settings tool in the status bar lets you configure your user settings for any online chat with your colleagues, including the saying, whether your online presence appears, and whether to play a sound when messages arrive.

When you have a message from another user, that user’s name appears on the status bar to the left of the settings tool.

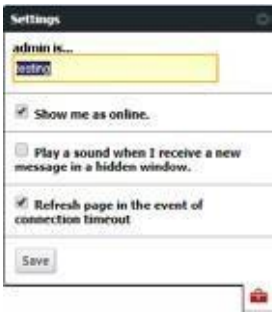

admin is...
testing

☑ Show me as online.

□ Play a sound when I receive a new message in a hidden window.

☑ Refresh page in the event of connection timeout

Conferencing

The Conferencing tool in the status bar lets you share whatever is of concern with other Cruz users and collaborate with more than one person. F or example, if a discovered device has problems, you can create a link to the device’s from the Connected Devices on page 189 and share it with other users with the Cruz internal instant messaging/chat system (see Sharing a Resource on page 146).

Use the Conferencing tool to configure your user settings for any online chat with multiple colleagues.

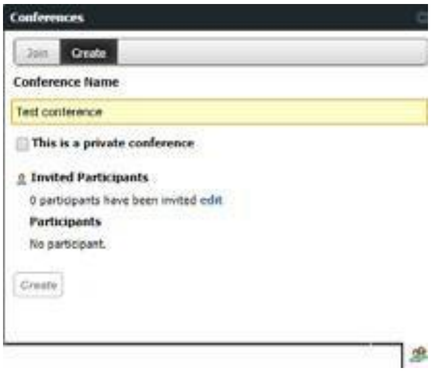

Conference Name
Test conference

□ This is a private conference

Invited Participants
0 participants have been invited edit
Participants
No participant.

The Join options become active when you are invited to a conference. An online chat window appears after you join.

The Create options let you name the conference, specify whether the conference is private (for invites only), and then invite colleagues and con figure an invitation message by clicking edit.

Conferencing also opens a screen that both records text and provides a virtual white board where participants can draw.

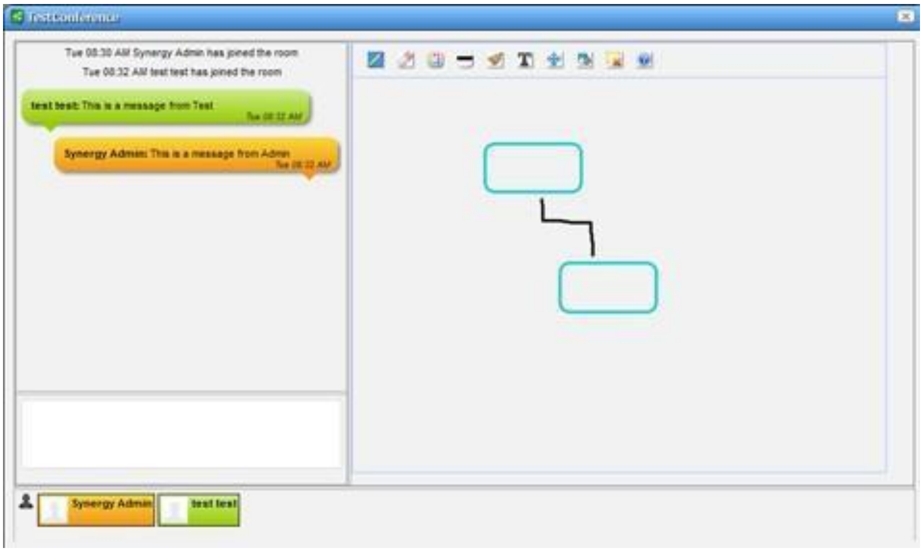

Tue 08:30 AM Synergy Admin has joined the room

Tue 08:32 AM test test has joined the room

test test: This is a message from Test
Tue 08:32 AM

Synergy Admin: This is a message from Admin
Tue 08:32 AM

Hover the cursor over the white board tools at the top to see what they do. Enter text in the lower left corner, and it appears on the left after you click Enter. Conference participants appear with icons and colors keyed to their text in the lowest portion of thescreen.

NOTE:

Ifappearanceorperformanceconcernsimpedeyourconferencing,clearyourbrowser’scache,thentry conferencing again.

Colleagues

The Colleagues tool has a red dot when you are online alone and a green dot when other users are online. It also shows the number of colleagues online. Click the Colleagues tool and a list of Cruz users is displayed. The title bar shows the number of users online. A green dot indicates the users that are online, and a red dot indicates the users that are not online. Click on a colleague and enter text in the popup that appears to send messages. Previouschat historyalso shows above any current text in the chat popup.

Click the upper right corner (-) to close the window.

NOTE:

You can turn off chat for the application with special branding available through your sales representative, but not for a single user. Chats are stored in the Cruz database, but as blobs, so reading chat history, except the date of chats, is problematic.

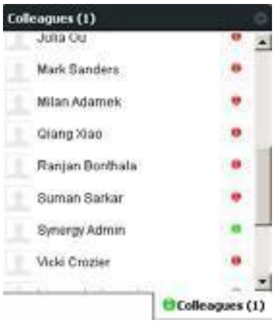

General Portlet Information

Portlets are the elements of any page within the portal. Whether you have access to a portlet and from which page a portlet is accessed is defined by your system administrator. In your environment, a portlet could be accessed from multiple pages. See the description of each portlet for more details.

Here are some common portal features that are related to portlets:

- **Modifying** pages by dragging portlets to a different location, adding portlets, or deleting portlets. The applications list shows all the available portlets when you select the Add > Applications menu option.
- Viewing data from both a summary view and an expanded portlet view for most portlets.
- **Sorting** on a column by clicking on that column’s heading. Reverse the sort order by clicking it again. These only sorts what appears in the port let, whether expanded or not. The application remembers each user’s choice saving the last Sort Column and Order on any page. The arrow to the right of that heading’s text displays the sort direction (ascending or descending). When the arrow appears in a heading, the selected column is the basis for sorting.
- **Resizing** columns by dragging the header border.
- **Print** a portlet’s content by exporting by exporting its expanded view content to a PDF, Excel, or CSV file, open the exported file in the appropriate application, and then printing it.
- All **time stamps** are based on the users timezone. It is based upon Greenwich Mean Time (GMT) and changes depending upon the timezone that you are in when viewing a record.
- Some portlets include editors that appear after you select an item, right-click, and then select either New or Edit.

  Mandatory fields in these editors appear with a red flag icon to their right. That flag may disappear once you fill in the field. Mandatory fields in an Action appear with a red flag icon to their right. That flag disappears once you add the action to an Action Group.

- Password fields do not support copy/paste operations. When you type into a password field, the text shows briefly and then disappears.

- Standard **tooltips** show for most portlet field's content, which is useful when the field cannot show all its content. Windows and editors have help tooltips when you click the question mark next to a field or what goes into a field in an editor when you hover over the cell/field. Forcells where a question mark appears when you hover over a listed item, a more detailed tooltip is displayed or a graph (in the Top N portlets) is displayed.

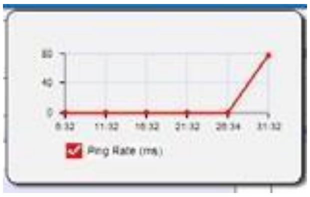

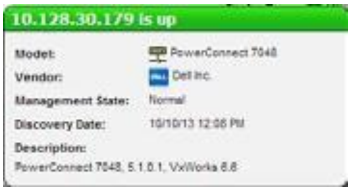

10.128.30.179 is up

Model: PowerConnect 7048

Vendor: Dell Inc.

Management State: Normal

Discovery Date: 10/10/13 12:06 PM

Description:

PowerConnect 7048, 5.1.0.1, VxWorks 6.6

If a graph is displayed, it can line, bar or pie graph, depending on the portlet, device and activity monitored. The more detailed tooltips require the latest Adobe Flash for full functionality.

This section describes common portlet features, such as the portlet tools and options , summary, footer, and hint.

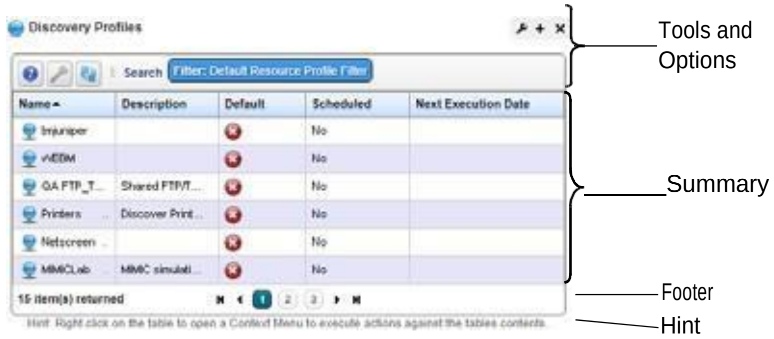

Name	Description	Default	Scheduled	Next Execution Date
Injuniper		☒	No	
vWEBM		☒	No	
QA FTP_T...	Shared FTP/T...	☒	No	
Printers	Discover Print...	☒	No	
Netscreen		☒	No	
MIMICLab	MIMIC simulati...	☒	No	

15 item(s) returned

Hint: Right click on the table to open a Context Menu to execute actions against the tables contents.

Portlet Tools and Options

The following portlet tools and options apply to all Cruz portlets, where available. Specific portlet features and functions are described in their respective section.

Tool	Description
🔧	Provides access to the followingoptions: **Look and Feel**, where you configure the portlet text styles, background, border, margin and padding, advanced, and WAP styles. **Configuration**, where you specify whether users have view, add to page, configuration, and permissions based on their assigned role. You can also specify whether to share the portlet with a website, Facebook, OpenSocial gadget, Netvibes, or Friends.
+	Displays a expanded portlet and provides additional information and options. Click Return to previous to return to the initial portlet view. This tool shows only for those portlets that have the Maximize option.
X	Removes the portlet from the existing page.
Return to previous	Takes you back to the previous page (initial portlet). Some examples where this option shows are the expanded portlet and details portlets. Although browsers have a back button, the Return to previous option is the most dependable way to return to a previous page, window, portlet, and so on within the Cruz portal.

Tool	Description
?	Displays online help for the selected portlet. Once you access the online help, you also have the option to search the help for other topics.
🔧	Opens a Settings window, where you filter the results and specify which columns to include and their settings. In some portlets, such as Alarms, this option can configure whether charts or graphs appear. See Modifying Column Settings on page 144 for instructions. Column selections from the portlet are not propagated to the maximized view. If you want the same columns for both views, you must select them for each view.
🔄	Isolates the browser's page refresh to the selected portlet so that you do not have to load the entire window again. After you modify a portlet's settings, click the refresh tool to see your changes.
Search	Filters the list to show more specific items. Select from the list of user-defined filters. This searches all available items in the database, whether they appear listed or not.
⏮	Displays the first page in the list.
◀	Displays the previous page in the list.
▶	Displays the next page in the list.
⏭	Displays the last page in the list.

Summary

The portlet summary supports displaying up to 500 rows, the expanded portlet supports up to 1000 rows. Using the portlet filtering capability makes more sense than trying to see more rows. (See Defining Advanced Filters on page 145 forinstructions.)

To act on listed items, right-click the item and then select from the popup menu. The menu options available depend on the portlet and the row selected.

To see information about listed items in a portlet, hover over the row until a large tooltip displays more details.

NOTE:

Portlet for Reports, Report Templates, Action and Compliance have a pre-seeded drop-down list filters to allow quick filtering of portlet rows.

Footer

The footer lists the number of items returned, the limit reached in red if you specified a maximum for the list, and the navigation buttons. The navigation buttons are the standard go to first/last page, move back/forward one page at a time, and select the page to go to.

Hint

The hint at the bottom of the portlet lets the user know that right clicking the table opens a popup menu (context menu) listing the available actions. By default, this hint is there for every portlet. If you **do not** want the hint showing, set the show.portal.hints attribute to false in the server-overrides.properties file. See Hiding Portlet Hint Text on page 143 for detailed instructions.

Expanded Portlet

The expanded portlet lets you display more information, do quick searches, do more advanced filters, or export the list to a PDF document or Excel or CSV formats. Youcan also see details about a selected row in the Widgets panel. Access the expanded portlet from the summary portlet's title bar by clicking the expand (+) tool. Return to the summary portlet by clicking Return to previous.

You can perform all the same actions from the expanded portlet, and you can from the summary portlet and navigating the list is the same as the summary portlet.

User permissions may limit access to the expanded portlets. Forexample, the Cruz can have many communities and limit users' memberships. Such users can lightly browse other Communities' screens without full privileges.

See Control Panel on page 32 for more about setting up user privileges for portlets.

NOTE:

Note that Screen size limitations may require youto expand the browser tosee expanded screens correctly. You must have at least 1250 pixels in width.

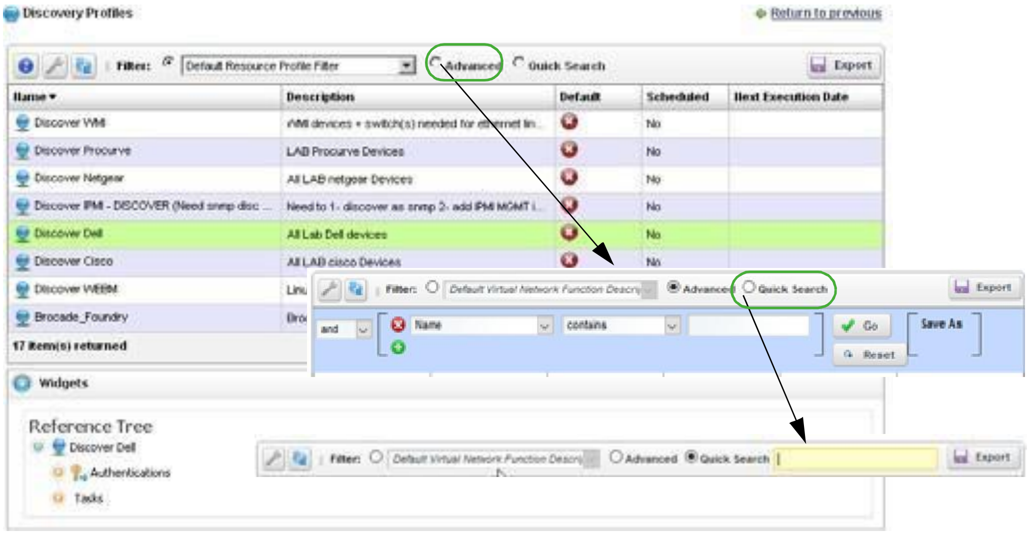

Name ▾	Description	Default	Scheduled	Next Execution Date
Discover VMI	VMI devices + switch(s) needed for ethernet lin...	✖	No	
Discover Procurve	LAB Procurve Devices	✖	No	
Discover Netgear	All LAB netgear Devices	✖	No	
Discover IPMI - DISCOVER (Need snmp disc ...	Need to 1- discover as snmp 2- add IPMI MGMT i...	✖	No	
Discover Dell	All Lab Dell devices	✖	No	
Discover Cisco	All LAB cisco Devices	✖	No	
Discover WEBM	Linu			
Brocade_Foundry	Broc			

17 item(s) returned

In addition to the options the summary viewprovides, the expanded portlet provides the following options.

Option	Description
Return to previous	Takes you back to the previous page (initial portlet).
Filter	Displays the default filter and any user-defined filters available. Select a filter from the list. Otherwise, select the Advanced or Quick Search option.
Advanced	Provides options to create conditional (AND, OR) filter statements. Once you create a conditional statement, click GO to view the results. Click Save As to save the filter. Click Reset to create another conditional filter. Otherwise, use the Filter or Quick Search option.
Quick Search	Activates a field, where you type a phrase and then press ENTER. Otherwise, use the Advanced or Filter option.

Option	Description
Export	Saves the current table to PDF, Excel, or CSV format. Click Export, select the format type, and then click Generate Export.
Widgets	Displays one or more detail widgets, such as the Reference Tree. You have the option to show or hide the widget details from the title bar. The Reference Tree also expands/collapses tree elements.

Settings Window

Use the Settings window to set the return results parameters, define a filter, and select the context mode. You can also specify which columns the portlet shows and the settings for each column.

The Filter/Results panel is displayed by default. Click the Columns tab to display the Columns panel.

NOTE:

Column selections from the portlet are not propagated to the maximized view. If you want the same columns for both views, you must select them for each view.

Click the portlets Settings tool to access this window.

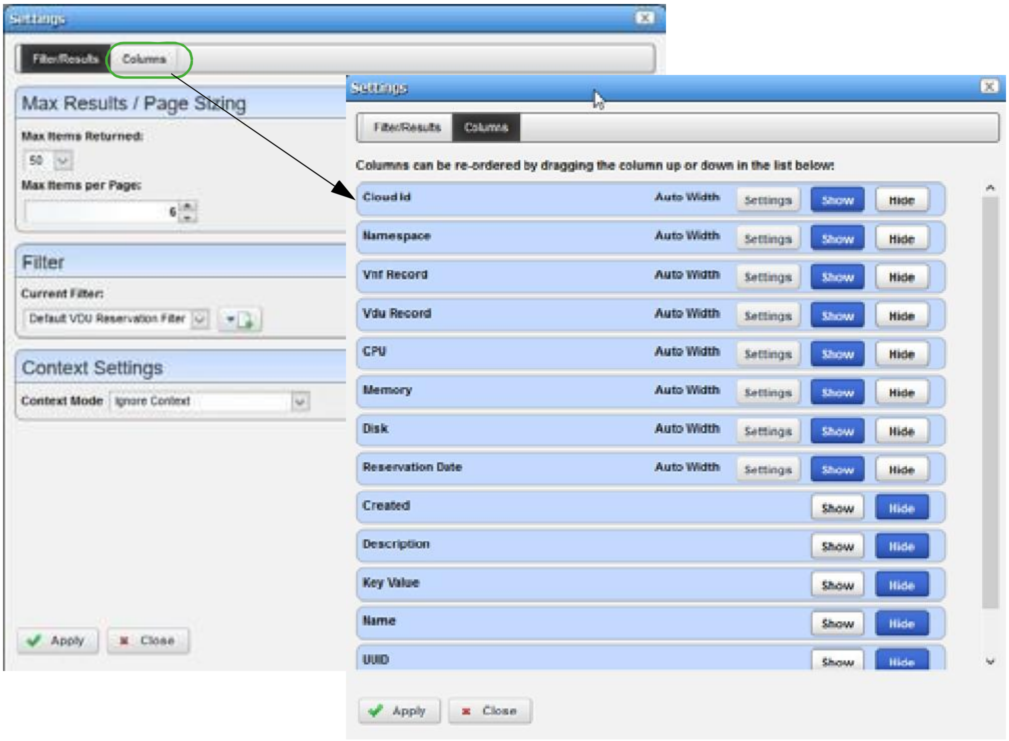

Max Items Returned: 50

Max Items per Page: 6

Current Filter: Default VDU Reservation Filter

Context Mode: Ignore Context

Columns can be re-ordered by dragging the column up or down in the list below:

The Settings window provides the following fields and options.

Field/Option	Description
Max Items Returned	Sets the maximum number of items the filter returns. Valid values are 2, 25, 50, 75, 100, 200, 300, or 500. For performance reasons, this default value is often relatively low.

Field/Option	Description
Max Items per Page	Sets the maximum number of it ems displayed per page. Valid values are 4 to 100. For performance reasons, this default value is often relatively low.
Current Filter	Provides a list of user-defined filters from which to choose. The product ships with a default filter. Depending on your permissions, you have the option to create a filter or copy the currently selected filter from which you create a filter. *The* expanded portlet does not include this option. See Defining Advanced Filters on page 145 for alternatives. **Note:** An administrator configures a portlet's default display filter and then clicks the portlet name and renames it. For example, make the default filter in the Managed Resources portlet to display only Dell Routers, then click the Managed Resources portlet name to rename it to *Dell Routers.* If you are not an administrator, you must make a personal page for such portlets if you want the filter settings to persist (not applicable in multitenant environments).
Context Mode	Specifies which of the following context modes the selected portlet uses: • Listen for Context • Ignore Context (default) • Broadcast Context • Broadcast and Listen for Context • Not all portlets have all option.
Columns	Displays a list of columns, their settings, and whether a column is hidden. Select the Column tab to display the columns. The default columns displayed, and the available columns vary with each portlet. If you have many instances of a portlet in your environment, the changes made to the portlet or maximized portlet are not global, they apply only to the current portlet or its maximized portlet. Click the Settings button to specify whether the width is automatically set (Auto Width) or specify a specific pixel width. Auto Width is the default. The grayed-out Show/Hide button indicates the selection for the column.

Applications List

The applications list is displayed when you select the Add > Applications menu option. The application (portlet) list shows some portlets with a purple icon and others with green icons.

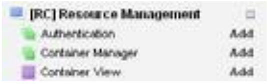

The **purple icon** indicates that you can add **only** one instance to a community, and it displays the same data, even if it appears on more than that one page, such as the Hierarchical View portlet. These are referred to as non-instanceable portlets. Once you add a non-instanceable portlet to a page, its entry in the applications list is disabled (grayed out).

The **green icon** indicates that you can add the portlet to many pages with each instance displaying different information, such as the Authentication or Hierarchical View Manager portlets. These are referred to as instanceable portlets.

Category	Application	Default Location/Notes
Actions	Action Group	Automation/Actions page **Note:** Can create only one instance to a community.
	Actions	Automation/Actions page
	Tasks	Available for use
Administration and Settings	Application Configuration Settings	Settings page **Note:** Can create only one instance to a community.
	Audit Trail	Settings > Audit page
	Common Setup Tasks	Settings page and Home page as part of Getting Started **Note:** Can create only one instance to a community.
	Getting Started	Home page **Note:** Can create only one instance to a community.
	Quick Navigation	Home page **Note:** Can create only one instance to a community.
	Schedules	Settings > Schedules page
Alarms, Events, and Automation	Alarms	Home, Alarms/Events, Alarms/Events > Hierarchical View, Topology > Hierarchical View pages
	Automation and Event Processing Rules	Automation/Actions page
	Event Definitions	Settings > Alarm Definitions page
	Event History	Alarms/Events, Automation/Actions pages
	Variable Binding Definitions	Settings > Alarm Definitions page
Compliance	Compliance Policies/Proscan	Compliance page
Configuration Management	Configuration Alarms	Configuration Management page
	Configuration Files	Configuration Management page
	Configuration Management Schedule	Configuration Management page
	File Servers	Configuration Management page
	Image Repository	Configuration Management page

Category	Application	Default Location/Notes
Hierarchical Views	Hierarchical View	Alarms/Events > Hierarchical View, Topology > Hierarchical View pages **Note:** Can create only one instance to a community.
	Hierarchical View Manager	Alarms/Events > Hierarchical View page
	Map Context	Available foruse **Note:** Can create only one instance to a community.
Multitenancy	Access Profile Templates	Available for use
	Site Management	Available for use
	User Site Access	Available for use
Network	Connected Devices	Resources page
	Links	
	Network Tools	Resources page
	Search by IP or Mac Address	Home, Resources pages **Note:** Can create only one instance to a community.
	VLAN Domain Assignment	Available for use
	VLAN Domains	Available for use
	VLAN Membership	Available for use
	VLANs	Available for use
Network Virtualization	Network Service Descriptors	Available for use
	Network Service Records	Available for use
	Physical Network Function Descriptors	Available for use
	Physical Network Function Records	Available for use
	Software Images	Available for use
	VIM Images	Available for use
	Virtual Network Function Descriptors	Available for use
	Virtual Network Function Records	Available for use
	Virtual Requirements	Available foruse **Note:** Can create only one instance to a community.
	Virtual Reservations	Available for use
	Virtualized Infrastructure Managers	Available for use
Performance Monitoring Management	Application Server Statistics	Settings > Server Configuration page **Note:** Can create only one instance to a community.
	Dashboard Views	Performance > Dashboard Management page
	Performance Dashboard	Available for use
	Resource Monitors	Settings > Monitor Management page
	System Dashboards	

Category	Application	Default Location/Notes
Performance Top N	Top Bandwidth Received	Available for use
	Top Bandwidth Received (bps)	Performance page
	Top Bandwidth Transmitted	Available for use
	Top Bandwidth Transmitted (bps)	Performance page
	Top CPU Utilization	Performance page
	Top Configuration Backups	Available for use
	Top Disk Utilization	Available for use
	Top Egress Packet Loss	Performance page
	Top Ingress Packet Loss	Performance page
	Top Input Discards	Available for use
	Top Input Errors	Performance page
	Top Interface Bandwidth	Performance page
	Top Interface Errors	Performance page
	Top Jitter	Available for use
	Top MOS	Available for use
	Top Memory Utilization	Performance page
	Top Output Discards	Available for use
	Top Output Errors	Performance page
	Top Packet Loss	Performance page
	Top Ping Response (Slowest)	Performance page
	Top Problem Nodes	Available for use
	Top RT Delay	Available for use
Portal Applications > News	Alerts	Available foruse **Note:** Can create only one instance to a community.
	Announcements	Availablefor use **Note:** Can create only one instance to a community.
	RSS	Available for use
	Weather	Available foruse **Note:** Can create only one instance to a community.
Portal Applications > Sample	IFrame	Available for use
	Web Proxy	Available for use

Category	Application	Default Location/Notes
Portal Applications > Tools	Dictionary	Available foruse **Note:** Can create only one instance to a community.
	Language	Available foruse **Note:** Can create only one instance to a community.
	Network Utilities	Available for use **Note:** Can create only one instance to a community.
	Password Generator	Available foruse **Note:** Can create only one instance to a community.
	Quick Note	Available for use
	Search	Available for use **Note:** Can create only one instance to a community.
	Sign In	Available foruse **Note:** Can create only one instance to a community.
	Unit Converter	Available foruse **Note:** Can create only one instance to a community.
Reports	Report Templates	Reports page
	Reports	Reports page
Resource Management	Authentications	Resources > Discover, Settings pages
	Cards	Available for use
	Contacts	Settings > Groups & Locations page
	Discovery Profiles	Resources > Discover page
	Interfaces	Resources page
	Locations	Settings > Groups & Locations page
	Managed Resource Groups	Settings > Groups & Locations page
	Managed Resources	Home, Resources, Configuration Management pages
	Ports	Resources page
	Vendors	Settings > Groups & Locations page
Services	Customers	Available for use
Storage	System Arrays	Available for use
Topology	System Topology	Topology, Topology > Hierarchical View pages
	Topology Views	Available for use

Category	Application	Default Location/Notes
Traffic Flow Analysis	Traffic Flow Applications	Traffic Flow page **Note:** Can create only one instance to a community.
	Traffic Flow Autonomous	Traffic Flow page **Note:** Can create only one instance to a community.
	Traffic Flow Conversations	Traffic Flow page **Note:** Can create only one instance to a community.
	Traffic Flow Endpoints	Traffic Flow page **Note:** Can create only one instance to a community.
	Traffic Flow Exporters by Managed Equipment	Traffic Flow page **Note:** Can create only one instance to a community.
	Traffic Flow Exporters by Subcomponent	Traffic Flow page **Note:** Can create only one instance to a community.
	Traffic Flow Protocols	Traffic Flow page **Note:** Can create only one instance to a community.
	Traffic Flow Receivers	Traffic Flow page **Note:** Can create only one instance to a community.
	Traffic Flow Senders	Traffic Flow page **Note:** Can create only one instance to a community.
	Traffic Flow Snapshots	Available for use

Manage Page

Use the Manage Page window to configure, add, or delete pages and to manage their appearance and permissions. You must refresh any altered page before edits take effect. Use the Copy Portlets from Page *option* to duplicate another page’s portlets on the selected page.

Alter the following portal components:

- Page permissions, appearance, order (using drag-and-drop), and so on
- Child page creation
- Page configuration import/export
- Page layout column configuration

  This option is not available if you have an expanded portlet open because the focus is not in the context of a page.

Access this window by selecting the Manage > Page or Manage > Page Layout menu option.

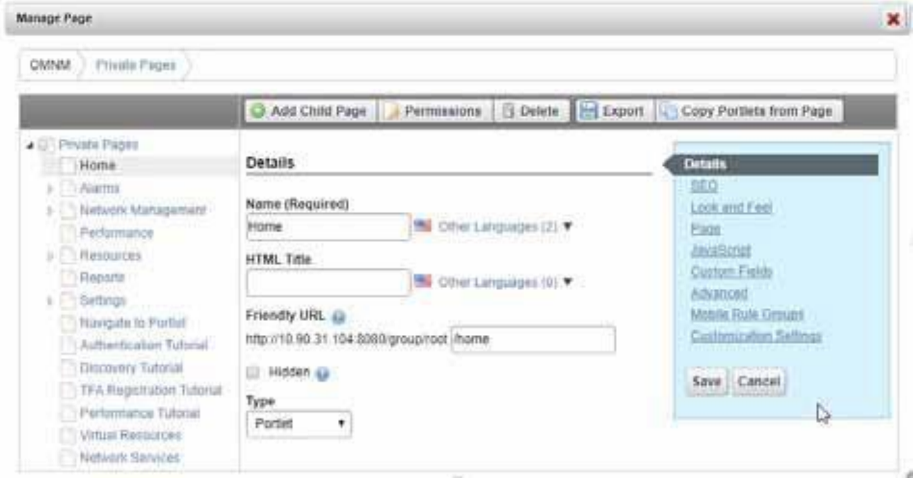

Name (Required): Home

HTML Title:

Friendly URL: http://10.90.31.104:8080/group/root /home

☐ Hidden

Type: Portlet

Show Versions

Use the Show Versions portlet to see which products and versions are installed. This portlet has the following panels:

- Product Details displays the installed package and modules, as well as their version numbers.
- Installed Extensions displays any installed presentation layer enhancements.
- Driver Information displays individual drivers (see Base Driver on page 107).

  The Profile Details outlines the supported device models, identifiers (OIDs), types and interfaces, and the OSVersionssupported bythedriver(although not device-by-device). This information is important when you need technical support.

NOTE:

You can also produce an HTML version of this information from a command line. Run `drvrpt` (drvrpt.cmd in Windows) from the \owareapps\ddbase\bin directory. The drvrpt command saves the HTML version in the *installRoot*\reports\drivers directory.

Access this portlet by selecting the *Manage > Show Versions* menu option.

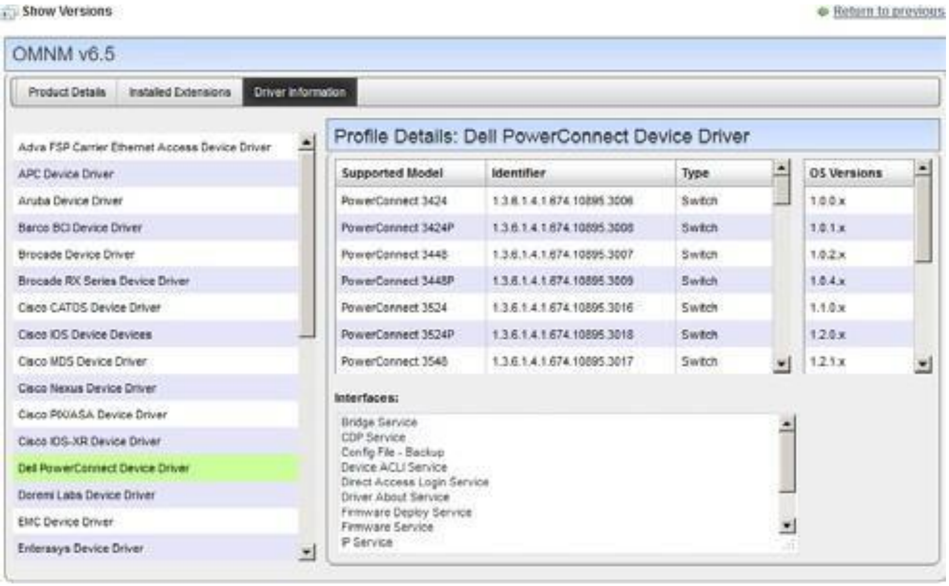

Adva FSP Carrier Ethernet Access Device Driver
APC Device Driver
Aruba Device Driver
Barco BCI Device Driver
Brocade Device Driver
Brocade RX Series Device Driver
Cisco CATOS Device Driver
Cisco IOS Device Devices
Cisco MDS Device Driver
Cisco Nexus Device Driver
Cisco PIX/ASA Device Driver
Cisco IOS-XR Device Driver
Dell PowerConnect Device Driver
Doremi Labs Device Driver
EMC Device Driver
Enterasys Device Driver

Profile Details: Dell PowerConnect Device Driver

Supported Model	Identifier	Type
PowerConnect 3424	1.3.6.1.4.1.674.10895.3006	Switch
PowerConnect 3424P	1.3.6.1.4.1.674.10895.3008	Switch
PowerConnect 3448	1.3.6.1.4.1.674.10895.3007	Switch
PowerConnect 3448P	1.3.6.1.4.1.674.10895.3009	Switch
PowerConnect 3524	1.3.6.1.4.1.674.10895.3016	Switch
PowerConnect 3524P	1.3.6.1.4.1.674.10895.3018	Switch
PowerConnect 3548	1.3.6.1.4.1.674.10895.3017	Switch

OS Versions
1.0.0.x
1.0.1.x
1.0.2.x
1.0.4.x
1.1.0.x
1.2.0.x
1.2.1.x

Interfaces:

Bridge Service
CDP Service
Config File - Backup
Device ACLI Service
Direct Access Login Service
Driver About Service
Firmware Deploy Service
Firmware Service
IP Service

Password Reminder

Use the Password Reminder portlet to add another level of security. If set by your administrator, this portlet displays when you sign into the Cruz portal.

You have the option to pick a question from the list or create your own question.

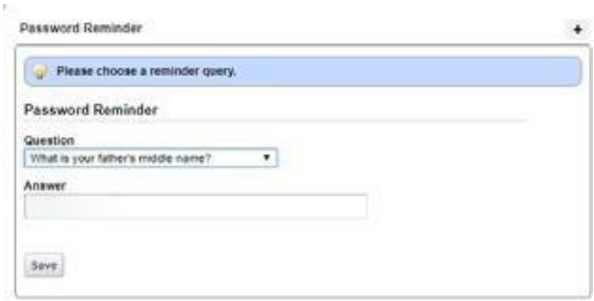

Please choose a reminder query.

Password Reminder

Question: What is your father's middle name? ▾

Answer:

Setting Time Formats

To set the time display in various locations (alarms, schedules, and so on), set the operating system's timeformatasyouwouldlike. Theseexamplestepsshowhow to set the Australiandefault day, month, year for the Windows 10 operating system.

Set Windows time formats as follows.

1. Navigate to the Control Panel.
2. Select Clock, Language and Region.
3. Click Change Date, Time, or number formats.

   The Region window displays the format settings.
4. Verify that the format is: English (Australia).
5. Click the Administrative tab.
6. Click on Copy settings.
7. Select both these options:
   - Welcome screen and system accounts
   - New user accounts
8. Click OK.
9. Restart the application server.
10. Verify that the Day/Month form at appears in your portlets.

Defining a DebugFile

For more advanced users, any component under owareapps can define a log4j.xml debug file for each component matching the following pattern:

```
owareapps\<component-dir>\server\conf\*log4j.xml
```

Consult these files for categories you want to change and copy those altered properties to the file you created in the owareapps\installprops directory. The categories altered in this file override any others. Changing such properties can produce enhanced error output in server logs. See also Application Server Statistics on page 359.

ActivatingLog4J Email Feature

The activation steps vary depending on whether you are activating the log4j email feature on the application/mediation server or on the Web server.

Defining Log4J on Application or Mediation Servers

Theapplication and mediation servers use JBoss, which defines Log4J settings through XML. Define the log4j email feature on Application or Mediation servers as follows.

1 Go to the .../oware/conf/ directory

2 Find the server-log4j.xml file.

3 Add the following alongside the other < appender > tags:

```
<!-- If this is present the processing of SMTP will be asynchronous.  It
 is not required -->
  <appender name="ASYNC" class="org.apache.log4j.AsyncAppender">
    <errorHandler class="org.jboss.logging.util.OnlyOnceErrorHandler"/>
    <appender-ref ref="SMTP"/>
  </appender>


  <!-- These are the main settings. Note that "SMTP" here is just a name.
 You choose any name you want
     and in fact you can have more than email appender -->
   <appender name="SMTP" class="org.apache.log4j.net.SMTPAppender">
<errorHandler class="org.jboss.logging.util.OnlyOnceErrorHandler"/>
    <param name="Threshold" value="ERROR"/>
    <param name="To" value="destination@email"/>
    <param name="From" value="sender@email"/>
    <param name="Subject" value="Testing Log4J Email feature"/>
    <param name="SMTPHost" value="email.com.au"/>
 <!-- you might need this <param name="TLS" value="true"/> -->
    <param name="SMTPUsername" value="myusername"/>
    <param name="SMTPPassword" value="mypassword"/> -->
    <param name="BufferSize" value="10"/> <-- find an appropriate value
 for this -->
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%m"/>  <!-- read more at
 https://logging.apache.org/log4j/1.2/apidocs/org/apache/log4j/
 PatternLayout.html -->
    </layout>
  </appender>
```

4 Find the server-log4j-tail.xml file.

5 Change the file to something like this:

```
<root>
    <appender-ref ref="CONSOLE"/>
    <appender-ref ref="FILE"/>
    <appender-ref ref="SMTP"/>
</root>
```

6 Restart the server.

Defining Log4J on Web Servers

Web servers use Tomcat, which defines Log4J settings through a properties file. Define the log4j email feature on Web servers as follows.

1. Go to the following directory:

   ```
   .../oware/synergy/tomcat-7.0.40/webapps/netview/WEB-INF/classes
   ```

2. Find the log4j.properties file.

3. Change the file to something like the following example, asappropriate.

   You may want to use some of the same values from the XML file used for the application server.

   ```
   log4j.logger.com.dorado=INFO, CONSOLE
   log4j.logger.com.dorado.netview.social=DEBUG, CONSOLE
   log4j.logger.com.icesoft=WARN, CONSOLE
   log4j.logger.com.icesoft.faces.async.render=TRACE, CONSOLE

   log4j.appender.CONSOLE=org.apache.log4j.ConsoleAppender
   log4j.appender.CONSOLE.layout=org.apache.log4j.PatternLayout
   log4j.appender.CONSOLE.layout.ConversionPattern=%d{ABSOLUTE} %-5p
     [%c{1}:%L] %m%n
   log4j.rootLogger=ERROR, EmailAlertsAppender
   log4j.appender.EmailAlertsAppender=org.apache.log4j.net.SMTPAppender
   log4j.appender.EmailAlertsAppender.From=sender@email
   log4j.appender.EmailAlertsAppender.To=destination@email
   log4j.appender.EmailAlertsAppender.Threshold=ERROR
   log4j.appender.EmailAlertsAppender.SMTPUsername=myusername
   log4j.appender.EmailAlertsAppender.SMTPPassword=mypassword
   log4j.appender.EmailAlertsAppender.SMTPHost=email.com.au
   log4j.appender.EmailAlertsAppender.BufferSize=10
   log4j.appender.EmailAlertsAppender.Subject=Testing Log4J Email feature
   log4j.appender.EmailAlertsAppender.layout=org.apache.log4j.PatternLayout
   log4j.appender.EmailAlertsAppender.layout.ConversionPattern=%m
   log4j.appender.EmailAlertsAppender.EvaluatorClass=TriggerLogEvent
   log4j.appender.EmailAlertsAppender.TLS=true
   ```

4. Restart the Web server

NOTE:

Your settings are overwritten on upgrade. Make sure that you back up the appropriate files before upgrading and then restore the files after upgrading.

Hiding Portlet Hint Text

By default, all portlets display hint text at the bottom. If you do not want the hint text to show, it can be hidden.

Hide the portlet hint text as follows.

1 Stop the Webserver.

For example, enter the following from the command line:

```
sudo service synergy stop
```

2 Navigate to the following directory:

```
installDir/oware/synergy/conf
```

3 Copy the server-overrides.properties.sample file to server-overrides.properties.

4 Open the server-overides.properties file with a text editor.

5 Uncomment the show.portal.hints attribute.

6 Set the show.portal.hints attribute to false.

7 Restart the Webserver.

For example, enter the following from the command line:

```
sudo service synergy start
```

8 Login and verify that the hint no longer shows.

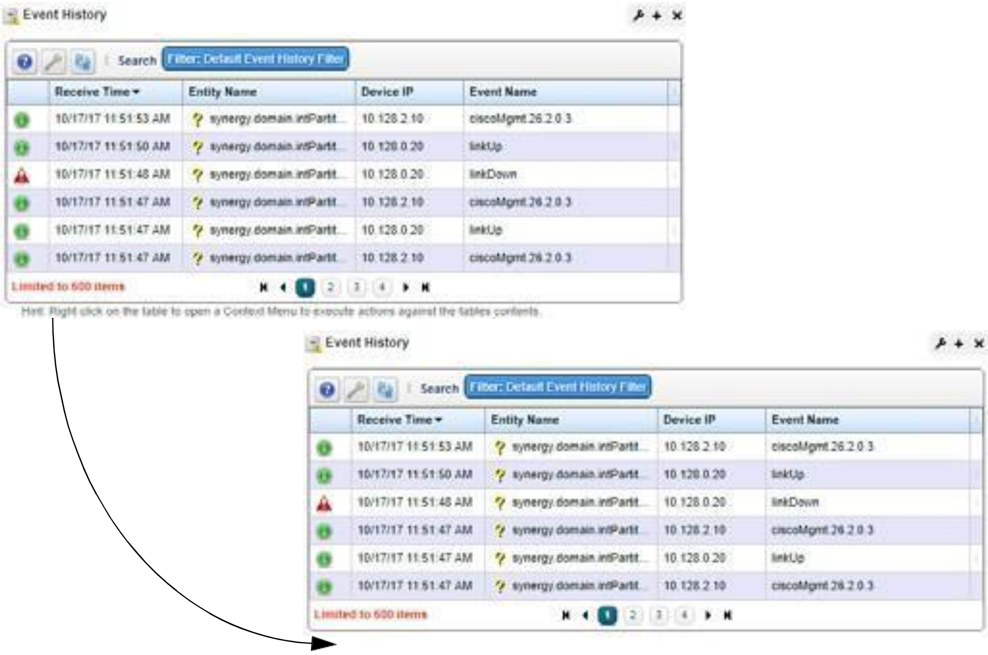

	Receive Time ▾	Entity Name	Device IP	Event Name
●	10/17/17 11:51:53 AM	? synergy.domain.intPartit...	10.128.2.10	ciscoMgmt.26.2.0.3
●	10/17/17 11:51:50 AM	? synergy.domain.intPartit...	10.128.0.20	linkUp
▲	10/17/17 11:51:48 AM	? synergy.domain.intPartit...	10.128.0.20	linkDown
●	10/17/17 11:51:47 AM	? synergy.domain.intPartit...	10.128.2.10	ciscoMgmt.26.2.0.3
●	10/17/17 11:51:47 AM	? synergy.domain.intPartit...	10.128.0.20	linkUp
●	10/17/17 11:51:47 AM	? synergy.domain.intPartit...	10.128.2.10	ciscoMgmt.26.2.0.3

Hint: Right click on the table to open a Context Menu to execute actions against the tables contents

Event History
Search
Filter: Default Event History Filter

	Receive Time ▾	Entity Name	Device IP	Event Name
●	10/17/17 11:51:53 AM	? synergy.domain.intPartit...	10.128.2.10	ciscoMgmt.26.2.0.3
●	10/17/17 11:51:50 AM	? synergy.domain.intPartit...	10.128.0.20	linkUp
▲	10/17/17 11:51:48 AM	? synergy.domain.intPartit...	10.128.0.20	linkDown
●	10/17/17 11:51:47 AM	? synergy.domain.intPartit...	10.128.2.10	ciscoMgmt.26.2.0.3
●	10/17/17 11:51:47 AM	? synergy.domain.intPartit...	10.128.0.20	linkUp
●	10/17/17 11:51:47 AM	? synergy.domain.intPartit...	10.128.2.10	ciscoMgmt.26.2.0.3

Limited to 600 items
1 2 3 4

Modifying Column Settings

You can make changes to the columns that are displayed, including to show additional columns, hide certain columns, change the column left-to- right orientation shown, or change the column width.

Modify column settings from an expanded or summary portlet as follows.

1 Click the Settings tool.

The Settings window is displayed.

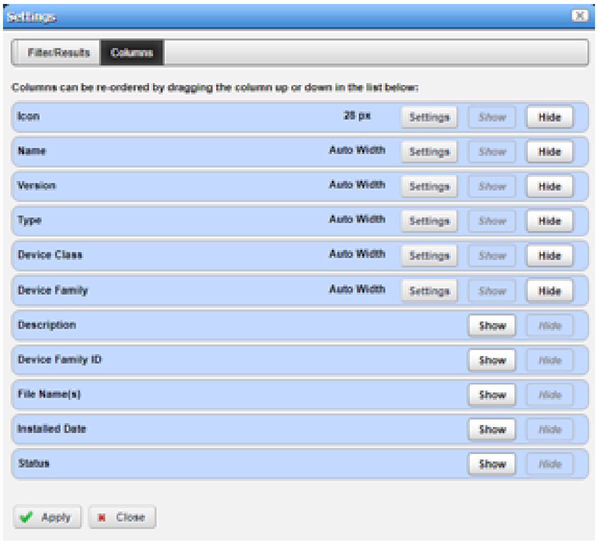

Settings

Filter/Results | Columns

Columns can be re-ordered by dragging the column up or down in the list below:

Icon	28 px	Settings	Show	Hide
Name	Auto Width	Settings	Show	Hide
Version	Auto Width	Settings	Show	Hide
Type	Auto Width	Settings	Show	Hide
Device Class	Auto Width	Settings	Show	Hide
Device Family	Auto Width	Settings	Show	Hide
Description			Show	Hide
Device Family ID			Show	Hide
File Name(s)			Show	Hide
Installed Date			Show	Hide
Status			Show	Hide

2 Click the Columns tab.

All data attributes that are available for the portlet are displayed.

3 Click the appropriate button to show/hide a column.

If you select to showa column, the settings option is activated, where you specify column width.

4 Click Settings to change the column width settings asneeded.

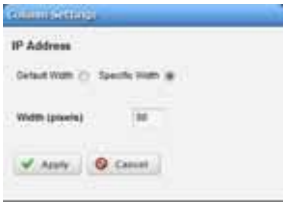

Column Settings

IP Address

Default Width ○ Specific Width ●

Width (pixels) 80

5 Change the order in which columns appear usingdrag-and-drop.

The top-to-bottom attribute orientation corresponds to the left-to-right column orientation within the portlet.

6 Click Apply and exit the window.

The changes appear instantaneously when you return to the portlet.

Defining Advanced Filters

Among other places, filters appear at the top of expanded portlets. Many pre-installed filters come from driver packages you installed. Filters match vendors and/or entity types but may not necessarily make sense in the context of a portlet. You can pick filters from already- configured filters list or click Advanced and create your own filter.

NOTE:

> You can also filter what appears on a page with the Hierarchical View. Select a hierarchy, and the rest of the portlets on that page confine displayed data to reflect the selected hierarchy's contents.

Define advanced filters as follows.

1. Select Advanced.

   The advanced filter fields and options are displayed.

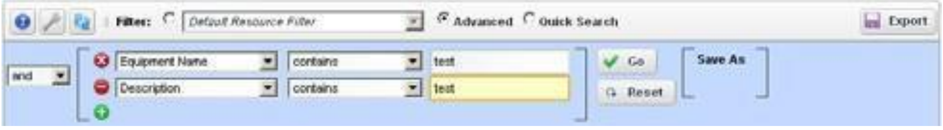

Filter: ○ Default Resource Filter ● Advanced ○ Quick Search

and

Equipment Name contains test

Description contains test

2. Select an operator (and/or) if combining more than one filter.
3. Provide the field, condition, and text to filter.
4. Click the add button (+) to define another filter.
5. Click Go to see the list after the filter acts on it.
6. Click Reset if you want to return the list to itsoriginal state.
7. Click Save As to preserve a filter you have configured for future use.
8. Enter a name and description.
9. Click Save

   The new filter is now available from the Filter list.

   See Redcell > Filter Management on page 43 for directions to the screen that catalogs all such filters.

Exporting/Importing PageConfigurations

Export/Import also appears as a tab in screens that manage pages (Manage > Page and Manage > Control Panel screens display these tabs). For example, click Manage > Settings in the Dock. Use the options on the Export/Import page to select exactly what elements to export. The automated file name includes your login identity, the date, and the lar extension. The file itself is a compressed collection of XML file configuration settings for the Pages/Portlets you have elected export. Its destination is the browser's default download location. Use the *More Options* link at the bottom of the Export screen to expose more export options. Use this same page to import such exported files, if it is enabled for your user.

Sharing a Resource

You can share elements within the Cruz system with colleagues when more than one user exists on your Cruz system, and consult with them using the conferencing described in Status Bar on page 124.

Share a resource with colleagues as follows.

1 Select a resource listed from the appropriate portlet.

2 Right-click and the select Share with User.

The Share with User window is displayed.

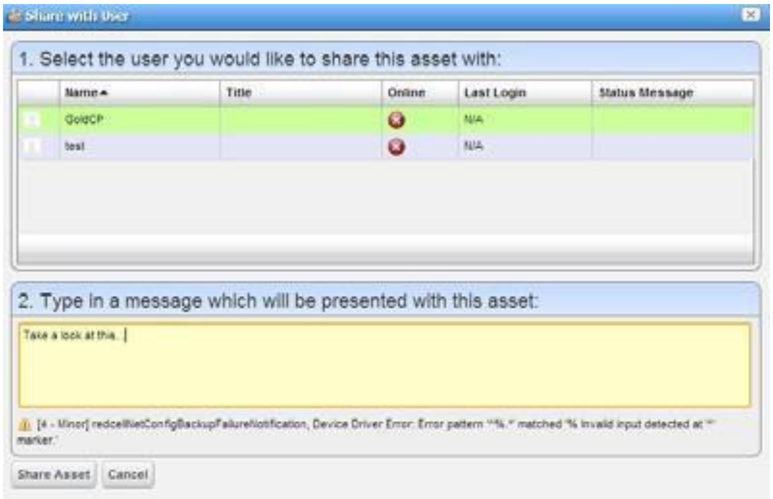

1. Select the user you would like to share this asset with:

	Name ▲	Title	Online	Last Login	Status Message
□	GoldCP		☒	N/A	
□	test		☒	N/A	

2. Type in a message which will be presented with this asset:

Take a look at this..

⚠ [4 - Minor] redcellNetConfigBackupFailureNotification, Device Driver Error: Error pattern "^%.*" matched "% Invalid input detected at '^' marker."

3 Select a user with whom you want to share.

4 Type any message you want to include.

5 Click Share Asset.

NOTE:

Sharing can only handle one item, so it uses the first one in the selection.

The chat message to the selected user includes your text and a link that opens to display the Widgets panel for the selected item.

Editing Custom Attributes

The Edit Custom Attributes pop-up menu option is available from many portlets (Managed Resources, Port, Contact, Vendor, or Location),

Edit custom attributes as follows.

1 Right-click the items whoseattributes you want to modify.

Forexample, right-click an alarm from the Alarms portlet.

The Custom Attribute Editor displays the definitions appropriate for your selection. See Redcell > Data Configuration on page 40 for another way to get to this editor.

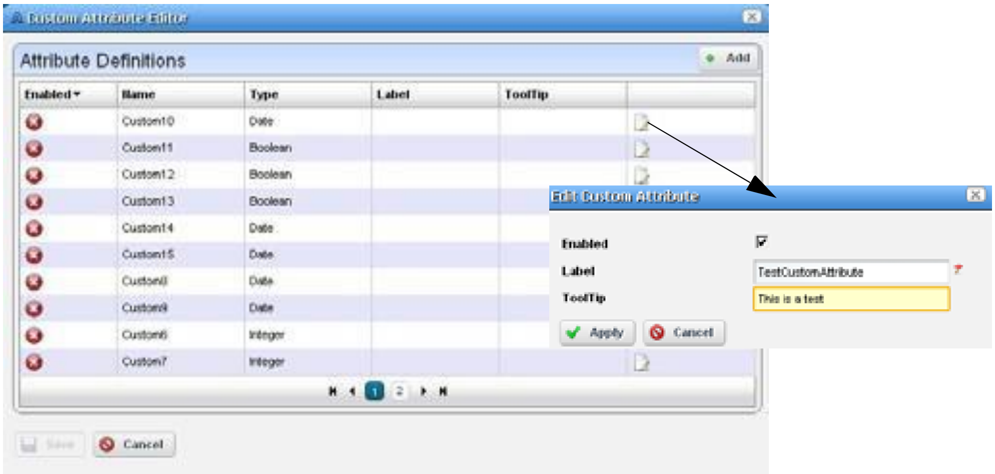

Enabled	Name	Type	Label	ToolTip
⊗	Custom10	Date		
⊗	Custom11	Boolean		
⊗	Custom12	Boolean		
⊗	Custom13	Boolean		
⊗	Custom14	Date		
⊗	Custom15	Date		
⊗	Custom8	Date		
⊗	Custom9	Date		
⊗	Custom6	Integer		
⊗	Custom7	Integer		

Enabled ☑

Label TestCustomAttribute

ToolTip This is a test

2 Click the edit tool.

3 Select Enable to activate the custom field.

4 Enter a label and optional tooltip.

The label is for the tooltip and is what you see in the portlets appropriate for the entity type you have selected. The tooltip that appears when you hover the cursor over the custom field.

NOTE:

Tooltips do not always function where the custom attribute appears in the Webclient. However, even If it does not appear, other views, Web services, or reports may use it.

1 Click Apply.

2 Click Save to preserve any changes you have made.

3 Verifythat the Custom Attributes panel nowexists.

a. Right-click a resource.

b. Select Edit.

c. See if the Custom Attributes tab exists in the Extended Details.

d. a resource and look in the Extended Details/Custom Attributes panel to see them.

NOTE:

If you want to enter the longitude and latitude for your Cruz installation, go to Control Panel's Redcell > Application Settings to enter the information as a default location.

Audit Trail/Job Status

When you execute an action, such as resyncing network resources, a Job Status window displays the message between the Cruz product and the devices the action addresses.

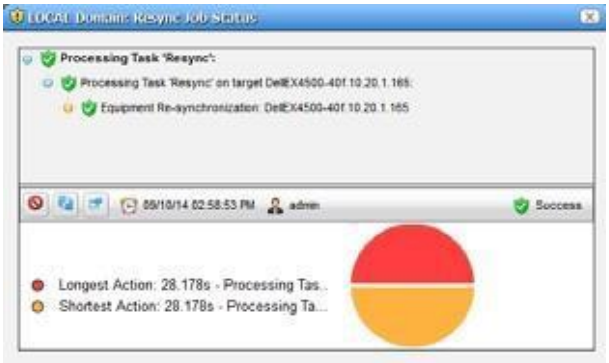

LOCAL Domain: Resync Job Status

- Processing Task 'Resync':
  - Processing Task 'Resync' on target DellEX4500-401 10.20.1.165:
    - Equipment Re-synchronization: DellEX4500-401 10.20.1.165

08/10/14 02:58:53 PM admin Success

● Longest Action: 28.178s - Processing Tas...
● Shortest Action: 28.178s - Processing Ta...

To see the details of any message, click on it, and those details appear below the toolbar. If you click on a summary message (not a “leaf” on the tree), a graph appears displaying the duration for its component messages. Hover your cursor over each portion of the graph for more details.

The time for messages and logged in user initiating the action appear on the bar between the upper and lower screen, and an icon summarizing the action appears on its right. Click the second icon from the left to configure the amount of detail displayed in audit messages. Click the first (Refresh) icon to re-display messages if you re-configure the types displayed.

To review the audit trail for recently completed processing, open the My Alerts tab in the lower left corner of the portal, and click the magnifying glass to the right of the message.

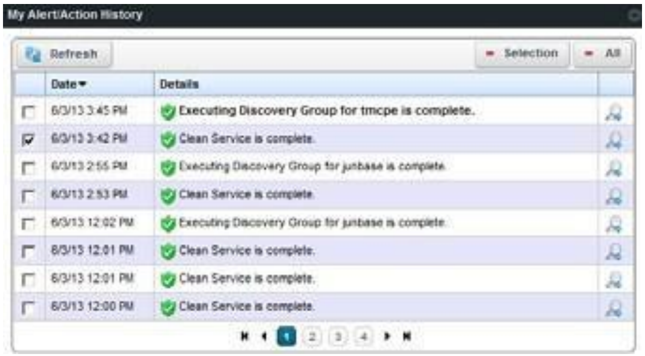

My Alert/Action History

	Date	Details	
□	6/3/13 3:45 PM	✓ Executing Discovery Group for tmcpe is complete.	
☑	6/3/13 3:42 PM	✓ Clean Service is complete.	
□	6/3/13 2:55 PM	✓ Executing Discovery Group for junbase is complete.	
□	6/3/13 2:53 PM	✓ Clean Service is complete.	
□	6/3/13 12:02 PM	✓ Executing Discovery Group for junbase is complete.	
□	6/3/13 12:01 PM	✓ Clean Service is complete.	
□	6/3/13 12:01 PM	✓ Clean Service is complete.	
□	6/3/13 12:00 PM	✓ Clean Service is complete.	

Some audit trails display as many as three tabs for the Input (the command variables sent to the device), the Job Viewer with the message traffic to the device, and finally the Results ofsending the messages to a device. This lists devices on the left and message traffic for a selected device on the right.

NOTE:

By default, the Job Viewer window conceals info-level messages. To see them, click the icon next to the Refreshicontoopenthemessagelevelselectorandchecktheinfocirclelevelofreporting,thenclick Refresh to see those blue circles.

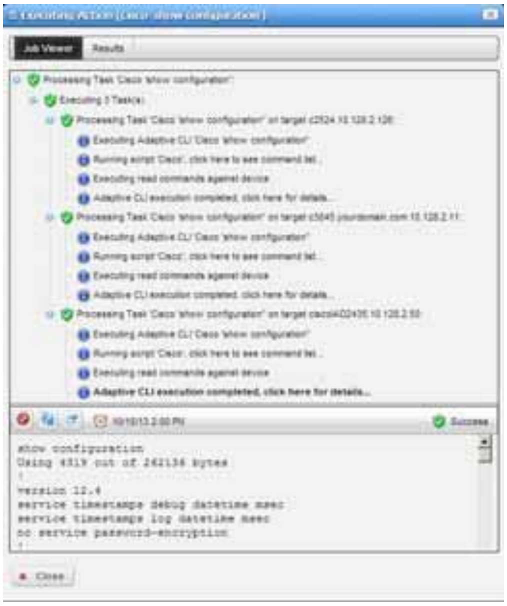

Processing Task 'Cisco 'show configuration''
Executing 3 Task(s)
Processing Task 'Cisco 'show configuration'' on target c2524.10.128.2.128
Executing Adaptive CLI 'Cisco 'show configuration''
Running script 'Cisco', click here to see command list...
Executing read commands against device
Adaptive CLI execution completed, click here for details...
Processing Task 'Cisco 'show configuration'' on target c3845.yourdomain.com 10.128.2.11
Executing Adaptive CLI 'Cisco 'show configuration''
Running script 'Cisco', click here to see command list...
Executing read commands against device
Adaptive CLI execution completed, click here for details...
Processing Task 'Cisco 'show configuration'' on target ciscoIAD2435.10.128.2.50
Executing Adaptive CLI 'Cisco 'show configuration''
Running script 'Cisco', click here to see command list...
Executing read commands against device
Adaptive CLI execution completed, click here for details...

```
show configuration
Using 4319 out of 262136 bytes
!
version 12.4
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
```

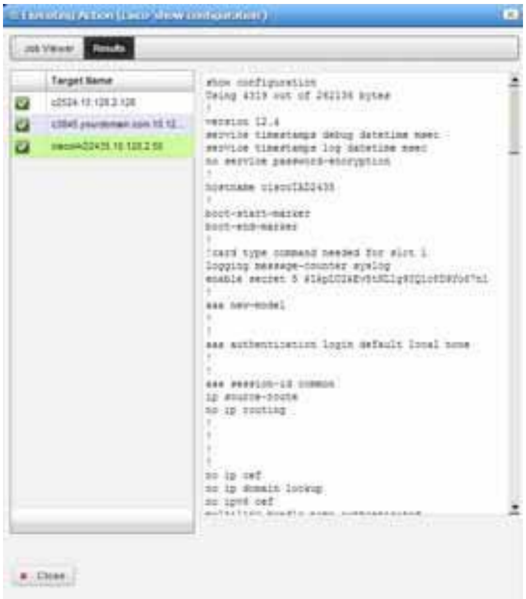

	Target Name
☑	c2524.10.128.2.128
☑	c3845.yourdomain.com 10.12...
☑	ciscoIAD2435.10.128.2.50

```
show configuration
Using 4319 out of 262136 bytes
!
version 12.4
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname ciscoIAD2435
!
boot-start-marker
boot-end-marker
!
!card type command needed for slot 1
logging message-counter syslog
enable secret 5 [illegible]
!
aaa new-model
!
!
aaa authentication login default local none
!
!
aaa session-id common
ip source-route
no ip routing
!
!
!
!
no ip cef
no ip domain lookup
no ipv6 cef
```

Close the audit trail viewer any time, and the action continues in the background. The audit trail is archived in the Audit Trail Portlet .

Cancel option, when displayed, stops some, but not all jobs in progress. The underlying feature (Discovery, Resync, Backup, and so on) described in the audit trail is responsible for gracefully stopping the execution flow, ensuring that the system and the database is left in agood state; not all features can do this. For performance reasons, it checks for cancellation at appropriate spots in the transaction where it is easy and safe to exit the execution flow. This means that even if the type of job supports cancellation, it may not cancel the current execution. If you press cancel while in the middle ofa multiple device resync, the Cruz application does not stop the resync or that device but instead bypasses the resync of subsequent devices.

Cancellation does not roll back work that has already been completed. So, if you are executing an Adaptive CLI action against 10 devices and you cancel the job after the third device is configured the Cruz application does not try to roll back the work that has already occurred against the first three, it does, however, stop executing against the remainingseven.

You can modify the Job Viewer appearance.

Modifying Job Viewer’s Appearance

Modify the Job Viewer appearance as follows.

1 Select Go to > Control Panel.

The Control panel is displayed.

2 Select Redcell > Application Settings .

The General application settings are displayed.

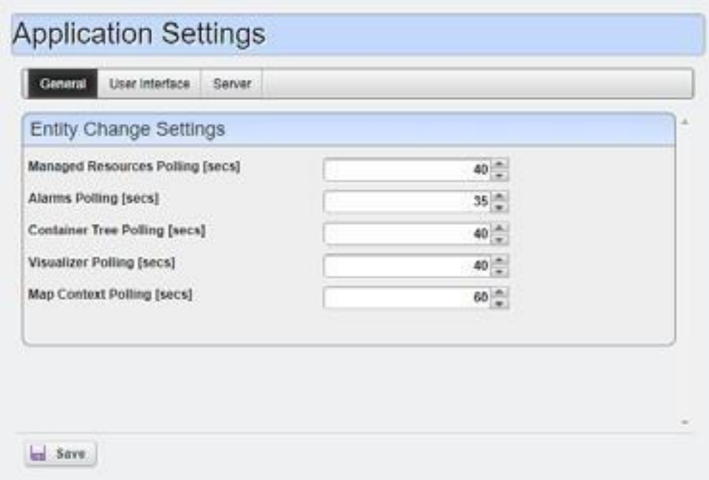

Entity Change Settings

Managed Resources Polling [secs] 40

Alarms Polling [secs] 35

Container Tree Polling [secs] 40

Visualizer Polling [secs] 40

Map Context Polling [secs] 60

3 Select the User Interface tab.

4 Go to the Job Viewer options.

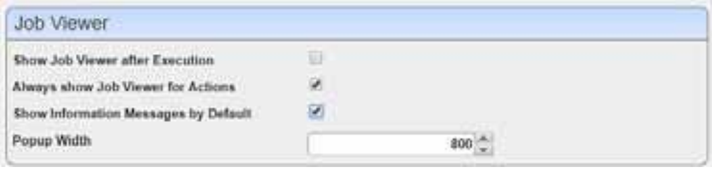

Job Viewer

Show Job Viewer after Execution □

Always show Job Viewer for Actions ☑

Show Information Messages by Default ☑

Popup Width 800

5 Select any of the following options that apply:

- Show Job Viewer after Execution
- Always show Job Viewer for Actions
- Show Information Messages by Default

6 Change the pop-up width if needed.

Audit Trail Portlet

The Audit Trailsummary portlet contains an archive of the Audit Trail/Job Status message traffic between the Cruz application and monitored devices, as wellas the Cruz reaction to failed message transmission.

By default, this portlet is available by selecting Settings > Audit from the navigation bar.

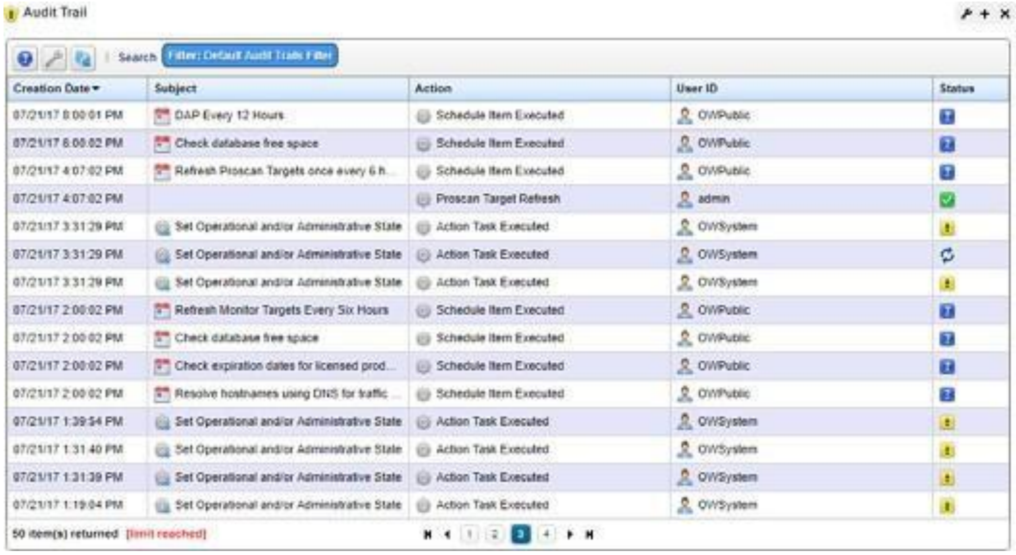

Creation Date ▾	Subject	Action	User ID	Status
07/21/17 8:00:01 PM	DAP Every 12 Hours	Schedule Item Executed	OWPublic	ℹ
07/21/17 6:00:02 PM	Check database free space	Schedule Item Executed	OWPublic	ℹ
07/21/17 4:07:02 PM	Refresh Proscan Targets once every 6 h...	Schedule Item Executed	OWPublic	ℹ
07/21/17 4:07:02 PM		Proscan Target Refresh	admin	✓
07/21/17 3:31:29 PM	Set Operational and/or Administrative State	Action Task Executed	OWSystem	!
07/21/17 3:31:29 PM	Set Operational and/or Administrative State	Action Task Executed	OWSystem	⟳
07/21/17 3:31:29 PM	Set Operational and/or Administrative State	Action Task Executed	OWSystem	!
07/21/17 2:00:02 PM	Refresh Monitor Targets Every Six Hours	Schedule Item Executed	OWPublic	ℹ
07/21/17 2:00:02 PM	Check database free space	Schedule Item Executed	OWPublic	ℹ
07/21/17 2:00:02 PM	Check expiration dates for licensed prod...	Schedule Item Executed	OWPublic	ℹ
07/21/17 2:00:02 PM	Resolve hostnames using DNS for traffic ...	Schedule Item Executed	OWPublic	ℹ
07/21/17 1:39:54 PM	Set Operational and/or Administrative State	Action Task Executed	OWSystem	!
07/21/17 1:31:40 PM	Set Operational and/or Administrative State	Action Task Executed	OWSystem	!
07/21/17 1:31:39 PM	Set Operational and/or Administrative State	Action Task Executed	OWSystem	!
07/21/17 1:19:04 PM	Set Operational and/or Administrative State	Action Task Executed	OWSystem	!

50 item(s) returned [limit reached]

The Creation Date, Subject, Action (the summary message of the audit trail), User ID (the login ID of the user whose actions resulted in this trail), and the status message appears when you hover over the Status field.

Right-click an item to View Job status, Delete a message, manage its Aging Policy, Viewas PDF, or Share with User. See Implementing DAP on page 80 for more about such policies.

The Audit Trail/Job Status portlet displays additional information.

Expanded Audit Trail

When you click the plus (+) in the upper right corner of the summary screen, the expanded portlet appears. Click the *Settings* button to configure the columns that appear in this screen and their order. Filter the screen appearance using the *Advanced Filter* capabilities at its top.

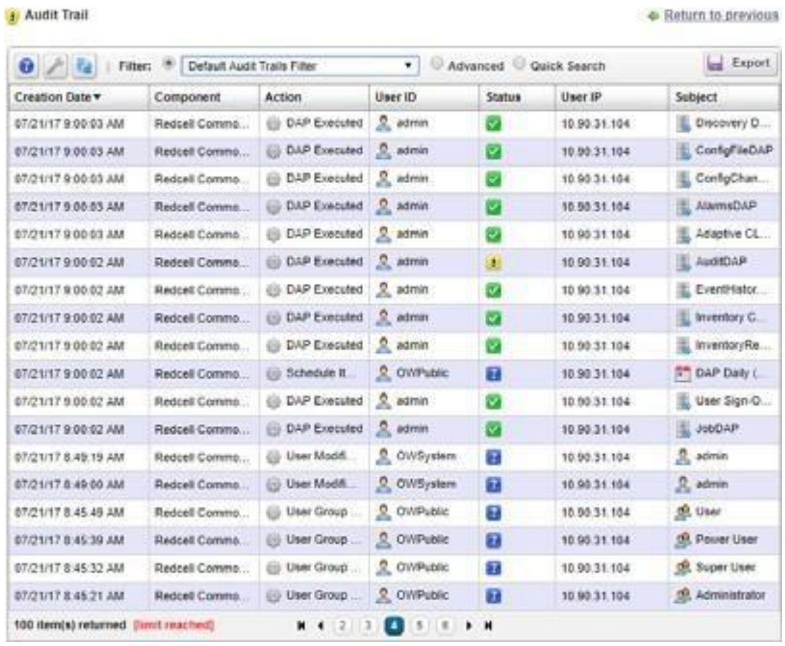
Audit Trail
Return to previous
Filter: ● Default Audit Trails Filter
○ Advanced ○ Quick Search
Export

Creation Date ▾	Component	Action	User ID	Status	User IP	Subject
07/21/17 9:00:03 AM	Redcell Commo...	DAP Executed	admin	✓	10.90.31.104	Discovery D...
07/21/17 9:00:03 AM	Redcell Commo...	DAP Executed	admin	✓	10.90.31.104	ConfigFileDAP
07/21/17 9:00:03 AM	Redcell Commo...	DAP Executed	admin	✓	10.90.31.104	ConfigChan...
07/21/17 9:00:03 AM	Redcell Commo...	DAP Executed	admin	✓	10.90.31.104	AlarmsDAP
07/21/17 9:00:03 AM	Redcell Commo...	DAP Executed	admin	✓	10.90.31.104	Adaptive CL...
07/21/17 9:00:02 AM	Redcell Commo...	DAP Executed	admin	!	10.90.31.104	AuditDAP
07/21/17 9:00:02 AM	Redcell Commo...	DAP Executed	admin	✓	10.90.31.104	EventHistor...
07/21/17 9:00:02 AM	Redcell Commo...	DAP Executed	admin	✓	10.90.31.104	Inventory C...
07/21/17 9:00:02 AM	Redcell Commo...	DAP Executed	admin	✓	10.90.31.104	InventoryRe...
07/21/17 9:00:02 AM	Redcell Commo...	Schedule It...	OWPublic	?	10.90.31.104	DAP Daily (...
07/21/17 9:00:02 AM	Redcell Commo...	DAP Executed	admin	✓	10.90.31.104	User Sign-O...
07/21/17 9:00:02 AM	Redcell Commo...	DAP Executed	admin	✓	10.90.31.104	JobDAP
07/21/17 8:49:19 AM	Redcell Commo...	User Modifi...	OWSystem	?	10.90.31.104	admin
07/21/17 8:49:00 AM	Redcell Commo...	User Modifi...	OWSystem	?	10.90.31.104	admin
07/21/17 8:45:49 AM	Redcell Commo...	User Group ...	OWPublic	?	10.90.31.104	User
07/21/17 8:45:39 AM	Redcell Commo...	User Group ...	OWPublic	?	10.90.31.104	Power User
07/21/17 8:45:32 AM	Redcell Commo...	User Group ...	OWPublic	?	10.90.31.104	Super User
07/21/17 8:45:21 AM	Redcell Commo...	User Group ...	OWPublic	?	10.90.31.104	Administrator

100 item(s) returned (limit reached)

In addition to the summary screen’scolumns, the following columns are available in expanded view:

- **User IP** the Cruz application creates the Audit Entry for IP Address of the related user. If it cannot acquire the user’s IP Address or if the audit entry occurred because of a Scheduled or System event, then the IP address is for the related Application Server.
- **Subject** the equipment at the origin of the message traffic with Cruz.

Right-click an item to provides the same options as the minimized view (View Job, Delete a message, manage its Aging Policy, View as PDF, or Share with User).

Job Status

The Job Status window displays the audit trail messages in tree form. Access this window by right-clicking an audit trail item and then selecting View Job. To see the contents of an individual message that appears in the upper panel, select it and view its contents in the bottom panel. The divider has Refresh double-arrow, and screen/arrow icons in the left corner, and an icon indicating the status of the job on the right. Click Refresh to clear an old message so you can view a new one.

Click the refresh button to check (info, warning, error) filters that limit the types of visible messages. Notice that when you select a message, its date and time appears to the right followed by the status.

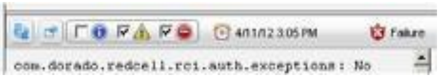

```
com.dorado.redcell.rci.auth.exceptions: No
```

Schedules

Use the Schedule options to schedule an action. These options display in either a window’s panel or a window of itsown depending on how you accessed them.

Access the Schedule options from a portlet that ordinarily executes schedules using the Schedule pop-up menu option, clicking the Schedule button from a window, or selecting the Schedule tab from a window. For example, right-click a discovery profile and then select Schedule from the Discovery Profiles portlet. Alternatively, right-click in the Schedules or Configuration Management Schedules portlets, select New > actionType > Schedule.

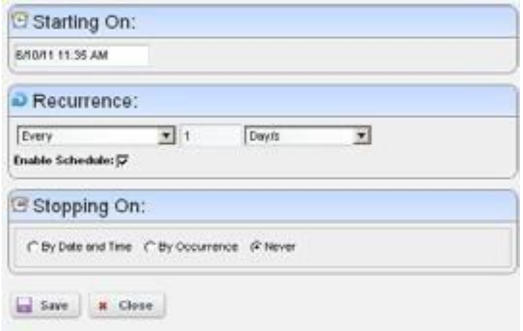

Starting On:

6/10/11 11:35 AM

Recurrence:

Every 1 Day/s

Enable Schedule: ☑

Stopping On:

○ By Date and Time ○ By Occurrence ● Never

Once you save the schedule, the action (for ex ample Discovery Profile) also appears in the Configuration Management Schedule portlet as a scheduleditem.

If you have the Cruz Change Management/Proscan capabilities installed, use Schedules to initiate the Change Determination process. See Change Determination Process on page 506. It is disabled by default.

The following descriptions provide more information about the Schedule fields and options.

Field/Option	Description
Starting On	Sets the date and time to execute the selected action.
Recurrence	Specifies how often to execute the selection action. Select one of the following recurrence types: • **Every**, the number, and unit of measure (Minutes, Hours, Days, Weekdays, Weekend Days, Weeks, Months, Years) • **Increment** (by minutes) and then specify the number of minutes • **Only at Startup** • **Only Once**
Enable Schedule	Activates the schedule when selected.
Stopping On	Sets whether to stop by date and time, occurrence, or never. If you select By Date and Time, the information displays in the *m/d/yy h:mm a* format. If you specify By Occurrence, specify the number of occurrences.

Once you schedule actions, you can view and modify them using the Schedules Portlet on page 155.

Scheduling Actions

Schedule an action rather than execute it immediately, for example from Managed Resource portlet, follow these steps:

1 Select the action in the right-click menu. For example: device Backup.

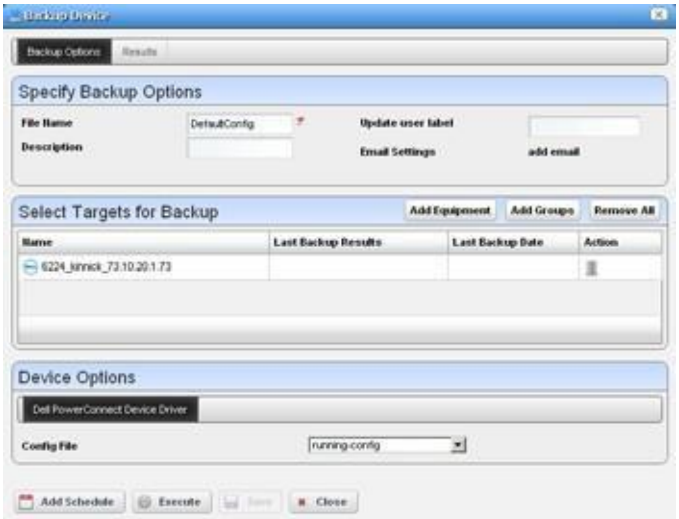

File Name: DefaultConfig

Description:

Update user label:

Email Settings: add email

Name	Last Backup Results	Last Backup Date	Action
6224_kinnick_73.10.20.1.73			

Config File: running-config

2 Click Add Schedule instead of Execute.

The schedule panel appears.

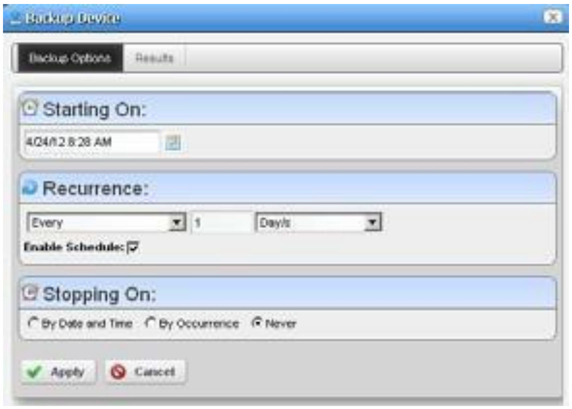

Starting On: 4/24/12 8:28 AM

Recurrence: Every 1 Day/s

Enable Schedule: ☑

Stopping On: ○ By Date and Time ○ By Occurrence ● Never

Configure the start time and date, recurrence, and stop parameters in this screen. The *Results* tab displays an audit trail when the action executes.

3 Click *Apply*.

The previous panel returns, the *Add Schedule* button now appearing as *Edit Schedule.*

4 Click *Save*.

Cruz creates a scheduled item around the activity and its data. A row also appears in the screen described in Schedules Portlet on page 145 for this schedule.

When you have scheduled something from the *Add Schedule* button, clicking *Apply* in the schedule panel returns you to the previous screen.

If you click *Execute* in that previous screen, the action begins, and the audit trail panel appears, displaying the running job for the activity. If you have attached a Schedule, Cruz also saves the activity as a scheduled item in the Schedules Portlet.

Schedules Portlet

Use the Schedules portlet to viewand modify schedules. If you have Cruz’s Change Management/ Proscancapabilities installed, you can use Schedule s to initiate the Change Determination process. See Change Determination Process on page 506. It is disabled bydefault.

This portlet is intended for users who are interested in scheduling tasks, such as Scheduling Actions on page 587.

Access this portlet byselecting Settings > Schedule s from the navigation bar. This portlet has both a summary view and an expanded view. Each view could display different columns and has the same pop-up menu options available.

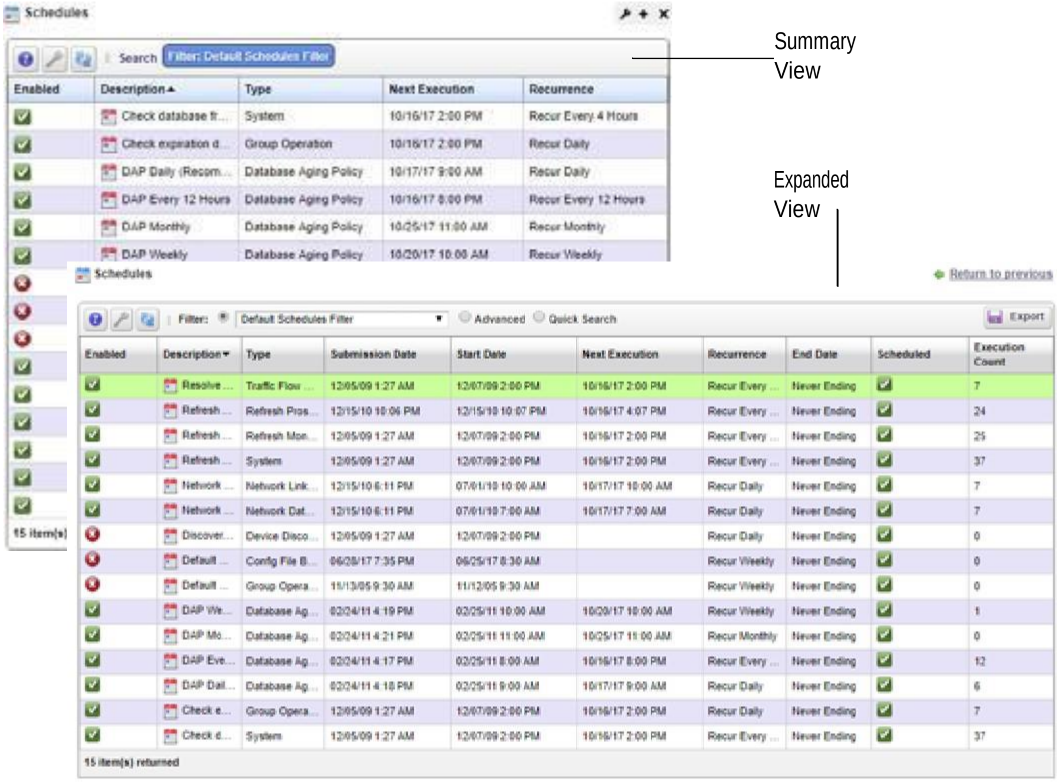

Enabled	Description	Type	Next Execution	Recurrence
✓	Check database fr...	System	10/16/17 2:00 PM	Recur Every 4 Hours
✓	Check expiration d...	Group Operation	10/16/17 2:00 PM	Recur Daily
✓	DAP Daily (Recom...	Database Aging Policy	10/17/17 9:00 AM	Recur Daily
✓	DAP Every 12 Hours	Database Aging Policy	10/16/17 8:00 PM	Recur Every 12 Hours
✓	DAP Monthly	Database Aging Policy	10/25/17 11:00 AM	Recur Monthly
✓	DAP Weekly	Database Aging Policy	10/20/17 10:00 AM	Recur Weekly

15 item(s)

Summary View

Expanded View

Enabled	Description	Type	Submission Date	Start Date	Next Execution	Recurrence	End Date	Scheduled	Execution Count
✓	Resolve ...	Traffic Flow ...	12/05/09 1:27 AM	12/07/09 2:00 PM	10/16/17 2:00 PM	Recur Every ...	Never Ending	✓	7
✓	Refresh ...	Refresh Pros...	12/15/10 10:06 PM	12/15/10 10:07 PM	10/16/17 4:07 PM	Recur Every ...	Never Ending	✓	24
✓	Refresh ...	Refresh Mon...	12/05/09 1:27 AM	12/07/09 2:00 PM	10/16/17 2:00 PM	Recur Every ...	Never Ending	✓	25
✓	Refresh ...	System	12/05/09 1:27 AM	12/07/09 2:00 PM	10/16/17 2:00 PM	Recur Every ...	Never Ending	✓	37
✓	Network ...	Network Link...	12/15/10 6:11 PM	07/01/10 10:00 AM	10/17/17 10:00 AM	Recur Daily	Never Ending	✓	7
✓	Network ...	Network Dat...	12/15/10 6:11 PM	07/01/10 7:00 AM	10/17/17 7:00 AM	Recur Daily	Never Ending	✓	7
✗	Discover...	Device Disco...	12/05/09 1:27 AM	12/07/09 2:00 PM		Recur Daily	Never Ending	✓	0
✗	Default ...	Config File B...	06/28/17 7:35 PM	06/25/17 8:30 AM		Recur Weekly	Never Ending	✓	0
✗	Default ...	Group Opera...	11/13/05 9:30 AM	11/12/05 9:30 AM		Recur Weekly	Never Ending	✓	0
✓	DAP We...	Database Ag...	02/24/11 4:19 PM	02/25/11 10:00 AM	10/20/17 10:00 AM	Recur Weekly	Never Ending	✓	1
✓	DAP Mo...	Database Ag...	02/24/11 4:21 PM	02/25/11 11:00 AM	10/25/17 11:00 AM	Recur Monthly	Never Ending	✓	0
✓	DAP Eve...	Database Ag...	02/24/11 4:17 PM	02/25/11 8:00 AM	10/16/17 8:00 PM	Recur Every ...	Never Ending	✓	12
✓	DAP Dai...	Database Ag...	02/24/11 4:18 PM	02/25/11 9:00 AM	10/17/17 9:00 AM	Recur Daily	Never Ending	✓	6
✓	Check e...	Group Opera...	12/05/09 1:27 AM	12/07/09 2:00 PM	10/16/17 2:00 PM	Recur Daily	Never Ending	✓	7
✓	Check d...	System	12/05/09 1:27 AM	12/07/09 2:00 PM	10/16/17 2:00 PM	Recur Every ...	Never Ending	✓	37

15 item(s) returned

Columns

Other thanthegeneralnavigationandconfigurationoptions, the Schedulesportlets(summaryand expanded views) include the following columns. The columns displayed by default are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Enabled	An indicator that shows whether the schedule is enabled (check mark) or disabled (X). This field displays on the summary and expanded views by default.
Description	A detailed description of the scheduled action, such as Network Data Collection, Default Device Config Backup, and so on. This field displays on the summary and expanded views by default.
Type	The type of action scheduled, such as Traffic Flow Analyzer, Refresh Proscan Targets or Monitor Targets, System, and so on. This field displays on the summary and expanded views by default.
Submission Date	The date and time that this schedule was submitted. This field displays on the expanded view by default.
Start Date	The date and time to start the scheduled action execution. This field displays on the expanded view by default.
Next Execution	The next date and time that the schedule will execute. This field displays on the summary and expanded views by default.
Recurrence	The frequency in which to execute the scheduled action, such as each weekday, every three months, only at startup, only once, and so on. This field displays on the summary and expanded views by default.
End Date	The date and time to end the scheduled action execution or the number of occurrences before ending this scheduled action. This field displays on the expanded view by default.
Scheduled	An indicator that shows whether action will execute on the next start date (check mark) or not (X). The scheduled action does not execute on the next schedule start date if it has exceeded the execution count or the specified end date has passed. This field displays on the expanded view by default.
Execution Count	The number of times the schedule will execute. This field displays on the expanded view by default.
Domain ID	The identifier for the resource domain.
Run Status	Indicate the schedules current state, such as waiting.

Pop-Up Menu

The Schedules pop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
New	Provides access to the following menu options that allow you to schedule a variety of actions: • Action (see Scheduling Actions on page 587 for the details) • Al arm Suppression (see Managed Resources on page 177 fordetails) • Config File Backup on page 443 • Config File Restore on page444 • Database Aging Policy (see Implementing DAP on page 80 for more about DAP) • OS Image Deploy on page 445 Based on the action selected, the subsequent windows allow you to configure actions, their targets, and the order in which they execute. **Note:** You can also schedule new actions from the portlet that ordinarily executes them, for example Discovering Resources on page 211.
Edit	Opens the Editing Schedule window, where you modify the activity’s schedule parameters. To edit an existing schedule for an already scheduled action like a Discovery Profile, just right-click the item in its portlet and select Schedule. This displays the schedule information for the discovery profile and lets you make modifications. **Note:** You can also modify a schedule from the portlet that ordinarily executes it. For example, Discovering Resources on page 211.
Delete	Removes the selected scheduled item.
Disable Schedule	Appears on an already enabled scheduled item.
Execute	Runs the scheduled action. If the scheduled item is an activity-based or discovery-profile-based scheduled item, an audit viewer displays progress (see Audit Trail/Job Status on page 148). For other types of scheduled actions, the following message is displayed: `The scheduled item(s) has been sent to the application server for immediate execution.` Monitor its status from the Audit Trail portlet.
Share with User	Opens the Share with User window where you select the colleague you want to share the selected rules with and then type your message.

3

Resource Management

This section describes the user interface components used to manage resources and then provides some resource management tasks. If you already have a good understanding of the portlets and editors, go directly to the tasks you want to perform.

Optional applications and device drivers may increase the basic functionality described here; your steps may differ slightly from the example steps provided.

Resource Management Portlets and Editors

The Resource management portlets let you manage devicesyou have discoveredorcreated on your network and view device-specific information, both general (name, type, location, contact) and technical (vendor, subcomponents, and so on). This section provides a detailed description of the following resource management portlets and editors:

- Authentications
- Discovery
- Managed Resource Groups
- Managed Resources
- Links
- Search by IP or Mac Address
- Connected Devices
- Equipment Details
- Direct Access
- Resource Management
- Interfaces
- Cards

Authentications

Use the Authentications portlet to see access credentials that let you discover and manage devices on your network. Authentications portlets include a summary and an expanded authentications portlet.

By default, this portlet is available by selecting Settings > Audit from the navigation bar.

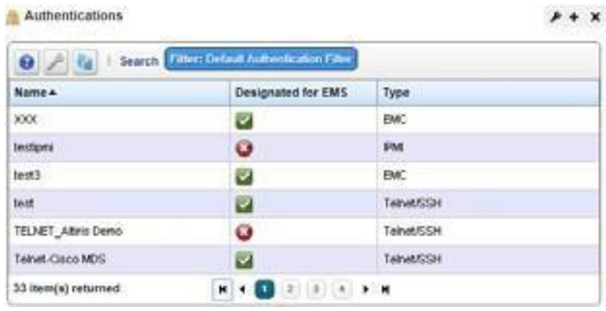

Name ▲	Designated for EMS	Type
XXX	✓	EMC
testipmi	✗	IPMI
test3	✓	EMC
test	✓	Telnet/SSH
TELNET_Altiris Demo	✗	Telnet/SSH
Telnet-Cisco MDS	✓	Telnet/SSH

33 item(s) returned

The Name column lists identifiers for sets of credentials, Designated for EMS means the credentials are accessible byall users, and Type indicates the protocol for that authentication.

NOTE:

If you have multitenancy installed, you can also elect to display a sites column that designates which tenant site owns an authentication.

The following pop-up menu are available.

New/Edit — Opens Authentication Editor , where you can create a new authentication or edit the selected authentication. You cannot change the Authentication Type when you edit an existing authentication.

Details—Displaysareference tree, associated equipment, and the configuration created or altered in Authentication Editor .

Audit — Opens an audit trail viewer for the selected authentication.

Delete—Deletes the selected authentication. If it is in use, an error message appears saying that deletion is not allowed.

Import/Export—Export the selected config file to disk or import it fromdisk. Youcan also import/ export a selected configuration file.

Provides the following actions when available for the selected image:

- Import retrieves a file containing XML compliance descriptions. Some imports can come from a URL.
- Export Selection exports the selected description to an XML file.
- Export All exports all descriptions to an XML file.

Click Download Export File to specify where to save the file.

The Import/Export option is useful as a backup or to share descriptor s with other projects.

Youmust import data into the correctportlet. Forexample, youcannotimport event data into the Actions portlet.

If one type of data depends on another, you must import the other data before importing the data that depends on it.

Share with User—Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Authentication Editor

Youcan right-click and select *New* or *Open* to create or modify credentials for your system. Youcan also *Delete* and *Share with User* from that right-click menu.

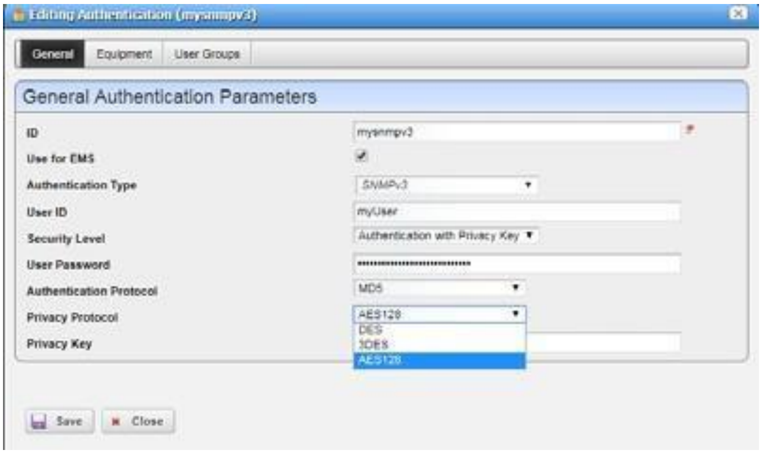
Editing Authentication (mysnmpv3)
General
Equipment
User Groups
General Authentication Parameters
DES
3DES
AES128
Save
Close

ID	mysnmpv3
Use for EMS	☑
Authentication Type	SNMPv3
User ID	myUser
Security Level	Authentication with Privacy Key
User Password	••••••••••••••••••••••••••
Authentication Protocol	MD5
Privacy Protocol	AES128
Privacy Key	

The fields that appear in this editor vary, depending on the type of authentication. The *ID* (name) for the authentication is mandatory. If you *Add* an existing authentication, for example to Discovery, you can also configure the Management Interface Parameters like *Timeout*, *Retries*, and *Port* used. If you have an authentication that works for multiple protocols (for example SSH or Telnet), you can also select the *Protocol Type*.

NOTE:

Discovery can fail because of network latency/timeout issues. Increasing the timeout or retries for Cruz authentications can circumvent that.

△ CAUTION:

If you do not get access to the deepest level of authentication—for example the "enable" user—you cannot access all of Cruz's functionality. Also: many devices require more than one authentication—for example SNMP and Telnet/SSH (including that enable authentication).

When attempting to access a device configured with SNMP v3, if you see an error message like `unable to read device serial number for selected credential`, discovery fails. Thisindicates the SNMP v3 credential is faulty. One common problem: SMNP v3 credentials must be at least eight characters long. Correct it, and discovery and other access should be available. Cruz's SNMP v3 authentications support the following:

- No Auth - No Priv
- Auth with MD5 and SHA digests - No Priv
- Auth with MD5 and SHA digests - Priv with DES, 3DES or AES128 encryption

△ CAUTION:

The standard for SNMP v3 passwords is eight characters or larger. Some devices may accept shorter passwords, but Cruz requires eight characters or longer. Also: Traps from an SNMP v3-accessed device do not appear when you change an SNMP v3 login or password on the device or Cruz, unless you resync the device, or until an SNMP monitor polls the device. Finally: Rebooting an SNMPv3 device may change the number for the SNMP v3 engine. You must restart application server to pickup this change.

Use the *Equipment* and *User Groups* tabs to associate the authentication you configure here to devices or groups of users.

Expanded Authentications Portlet

The *Settings* button in the expanded Authentication portlet lets you configure column appearance (see Modifying Column Settings on page 144). This offers the same column setup as the summary screen.

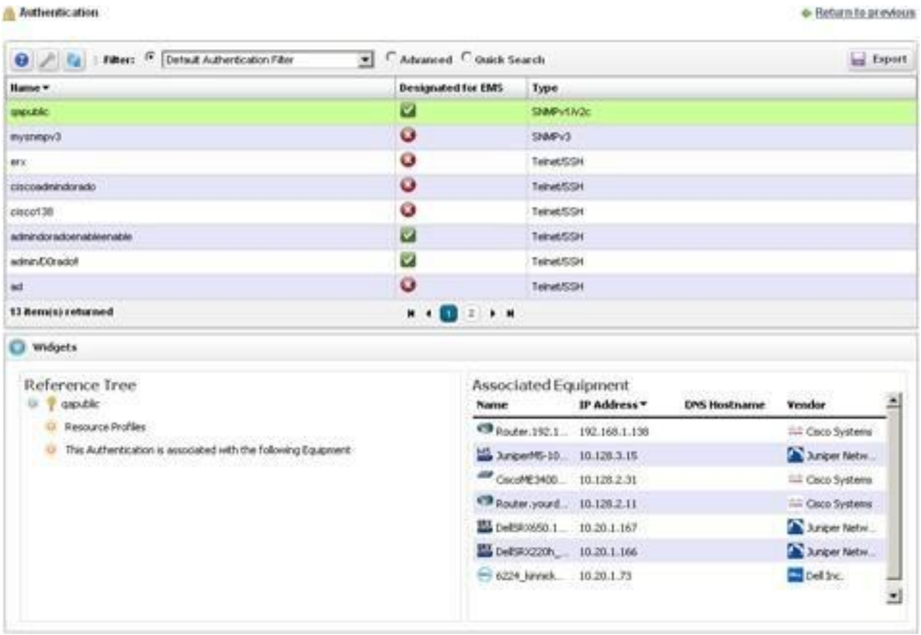

Name	Designated for EMS	Type
qapublic	☑	SNMPv1/v2c
mysnmpv3	☒	SNMPv3
erx	☒	Telnet/SSH
ciscoadmindorado	☒	Telnet/SSH
cisco138	☒	Telnet/SSH
admindoradoenableenable	☑	Telnet/SSH
admin/D0rado!	☑	Telnet/SSH
ad	☒	Telnet/SSH

13 item(s) returned

Reference Tree

- qapublic
  - Resource Profiles
  - This Authentication is associated with the following Equipment

Associated Equipment

Name	IP Address	DNS Hostname	Vendor
Router.192.1...	192.168.1.138		Cisco Systems
JuniperM5-10...	10.128.3.15		Juniper Netw...
CiscoME3400...	10.128.2.31		Cisco Systems
Router.yourd...	10.128.2.11		Cisco Systems
DellSRX650.1...	10.20.1.167		Juniper Netw...
DellSRX220h_...	10.20.1.166		Juniper Netw...
6224_kinnick...	10.20.1.73		Dell Inc.

When you select a listed authentication the *Reference Tree* Snap Panel displays a tree of that authentication's connections to Discovery profiles and equipment.

Discovery

The Cruz application has a Discovery Profiles portlet and Discovery Profile Editor. This section describes both.

Discovery Profiles

Discoveryprofiles configure equipment discoveryfor Cruz. These profilesconfigure equipment discovery for the Cruz application.

Access this portlet by selecting Resources > Discovery from the navigation bar.

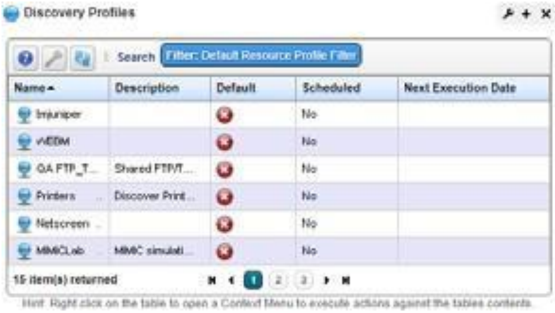

Name	Description	Default	Scheduled	Next Execution Date
Injuniper		☒	No	
vNEDM		☒	No	
QA FTP_T...	Shared FTP/T...	☒	No	
Printers ...	Discover Print...	☒	No	
Netscreen ...		☒	No	
MIMICLab ...	MIMIC simulati...	☒	No	

15 item(s) returned

Hint: Right click on the table to open a Context Menu to execute actions against the tables contents.

The summary view displays the *Name, Description, Default* (the green check indicates the default profile), whether the profile is *Scheduled* and *Next Execution Date* for scheduled discovery.

The Expanded portlet adds a Reference Treesnap panel that displays a tree ofassociations between selected profiles and authentication and tasks that they execute. See Discovery on page 163 for more about this portlet.

Youcan import discovery profiles to target multi-tenant domains with a command line importer. Thecommand is `importprofiles`andisin the `owareapps/redcell/bin directory.` This command takes the import file name an argument. The required domains should be available in the Cruz before import occurs. Before importin g discovery profiles to domains, any referenced authentications should be available in the domains or should be imported first by using the `importauths` command (the same way you would import discovery files). In other words, you should either manually create authentications for domains or import authentication files using importauths command before importing discovery files to those domains. Example XML files (with the `<customer>` tag for domains) are in `owareapps\redcell\db.`

NOTE:

> The date format follows the operating system's conventions for the location and language selected. Restarting the system changes system menus to the new language. If you want to revert back to the original language in Linux, you may also need to update the cache file under `/var/cache/gdm.`

When Cruz discovers unknown devices, it examines the RFC1213 MIB for hints of the device's capabilities, determining if it looks similar to a layer 3 router or a layer 2 switch. Since some device can do both, Cruz classifies such ambiguous devices as routers. See Base Driver on page 107 for more about generic discovery capabilities.

When you right-click a profile, the following menu options appear:

New — Opens Discovery Profile Editor in new profile mode. (see General on page 166)

If you have the multitenancy option installed, you can limit a discoveryprofile to a tenant site or have it discover devices within the entire system. To create discovery profiles for a site, select the site with *New (Specified Site)* right-click menu item for the Discovery Profile portlet.

Edit — Opens Discovery Profile Editor.

Copy— Opens Discovery Profile Editor, and renames the selected profile as "CopyOf[Original Name]". Rename this copy appropriately beforeproceeding.

Execute—Executes a discovery profile. This also produces an audit trail (see Audit Trail/Job Status on page 148). A message appears indicating the success or failure of discovery execution.

Discovery execution continues in the background even when you close the audit trail/jobs screen, but the message indicating success/failure still appears when the discoveryprocessis done.

Inspect—This validates that the device responds to ping, the profile's credentials, and that the device is licensed for discovery. See Inspection on page171.

Quick Discovery—Opensdiscoverywizarddisplayingnetworkand authentications, but without the Actions and Inspection panels. Not only does the Actions not appear, Quick Discovery does not executeany Actions. It is important to note that Quick Discovery settings, including Authentications and temporary device mappings, are not persisted. Device authentication mappings are only persisted in saved Discovery profiles. Click the Execute button once you open this screen to quickly discover equipment. See Network on page 168 for more about the screen this displays. Quick discovery is not available on tenant sites in a Multitenant system.

NOTE:

When using Quick Discovery, a "Quick Discovery" profile is created in Discovery Profiles. Quick discovery is intended tominimize userinputs by eliminating some data entry. For example, there is no description fieldandprofilewillhaveablankdescriptionfield.Userscaneditthequickdiscoveryprofileandseta description,buttheprofilewillbeoverwritten withtheexecutionoftheQuick Discoveryprocess.Itis importanttonotethat QuickDiscoverysettings, including Authenticationsandtemporarydevice mappings, are not persisted. Device authentication mappings are only persisted in saved Discovery profiles

Schedule—Opens schedule editor where youcancreateand/or modify theschedule fora discovery profile’s execution.

Audit — Displays audit trails for the selected profile. See Audit Trail/Job Status on page 148.

Delete—Deletes a discovery profile. After confirming that is what you want to do, a notification message appears when deletion is completed on the application server.

Import/Export—Export the selected config file to disk or import it fromdisk. Youcan also import/ export a selected configuration file.

Provides thefollowingactions when availablefor theselectedimage:

- Import retrieves a file containing XML compliance descriptions. Some imports can come from a URL.
- Export Selection exports the selected description to an XML file.
- Export All exports all descriptions to an XML file.

Click Download Export File to specify where to save the file.

The Import/Export option is useful as a backup or to share descriptor s with other projects.

Youmust import data into the correctportlet. Forexample, youcannotimport event data into the Actions portlet.

If one type of data depends on another, you must import the other data before importing the data that depends on it.

Share with User—Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

NOTE:

Cruz discovers Aruba Accesspoints through the controllers to which they connect, discovery does not find stand-alone accesspoints.

Discovery Profile Editor

The Discovery Profile editor lets you create or modify profiles by setting the following parameters and options:

- General
- Network
- Actions
- Inspection
- Results

General

The Discovery Profile editor opens to the General.

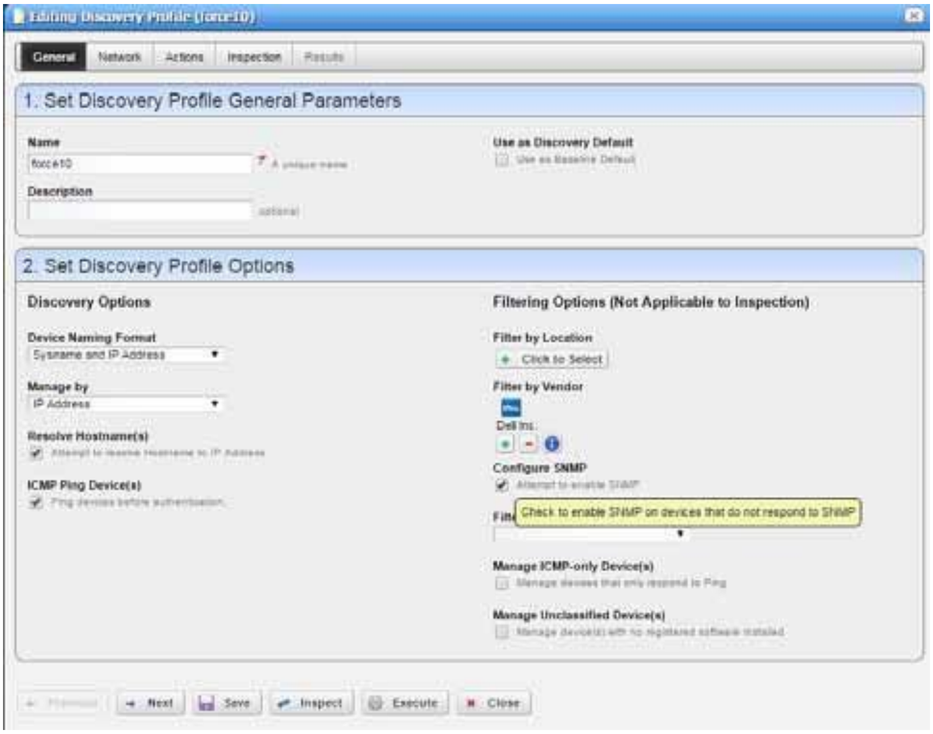

1. Set Discovery Profile General Parameters

Name
force10
A unique name

Description
optional

Use as Discovery Default
□ Use as Baseline Default

2. Set Discovery Profile Options

Discovery Options

Device Naming Format
Sysname and IP Address ▾

Manage by
IP Address ▾

Resolve Hostname(s)
☑ Attempt to resolve Hostname to IP Address

ICMP Ping Device(s)
☑ Ping devices before authentication.

Filtering Options (Not Applicable to Inspection)

Filter by Location
+ Click to Select

Filter by Vendor
Dell Inc.

Configure SNMP
☑ Attempt to enable SNMP

Check to enable SNMP on devices that do not respond to SNMP

Manage ICMP-only Device(s)
□ Manage devices that only respond to Ping

Manage Unclassified Device(s)
□ Manage device(s) with no registered software installed

The General panel includes settings for:

- **General Parameters,** where youspecify the name, description, and whether this profile is the discovery default.
- **Profile Options,**where you specify the Device Naming Format (how the device appears in lists, once discovered), whether to Manage by IP addressor hostname, and whether to Resolve Hostnames, ICMP Ping Devices, Manage ICMP-only Devices, or Manage Unclassified Devices. This last option determines whether the Cruz application attempts to manage devices that have no Cruz device driver installed. If your system's license permits it, such management may be possible, but more limited than for devices with drivers installed.

  If your license limits the number of devices you manage, discovering such "generic" devices may count against that limit. See Base Driver on page 107 for more.

  The Filters (by *Location*, *Vendor*, or *Device Type*) let you narrow the list of devices discovered by the selected item(s). As the screen says, this filtering will not have any impact on the processing that occurs during the Inspection step.

  The Attempt to enable SNMP option is visible if a vendor has been selected. If you select this option and a device is not reached during the SNMP discovery process, the discovery engine attempts to enable SNMP on the device using the SNMP credentials provided in the authentication section of the discovery profile. In order to enable SNMP on the device the device must be reachable via CLI with the CLI credentials provided. If the device is successfully enabled for SNMP, the SNMP discovery process will be reattempted.

NOTE:

Fieldslike*Location*querythedatabaseforcurrentinformation,soeventhoughitsfieldmayappear empty, Locations may exist. Click the Search button to the right of this field to populate it. Keeping such fields empty until you use them enhances performance.

The buttons at the bottom of the Profile Editor let you navigate through this series of panels. *Previous/Next* movebackand forthbetweenscreens, *Save* lets you preserve whatever stage you have configured, and close the editor, *Inspect* moves directly to the Inspection screen (described below), and *Execute* triggers the discovery profile and opens the Results panel, displaying message traffic between Cruz and the device(s). Click the “X” in the top right corner of these screens to close them without saving.

If you discover devices without retrieving their hostnames, and need that hostname later, you can re-run discovery after checking *Resolve DNS Hostnames*. This fetches the DNS hostname and resyncs the device.

NOTE:

If there is a need to sync up discovered equipment names to a new format, or to keep the current format insync with thedevice- ieas sysname changes, You can enter one of the property options below into the installed.propertiesfilein `/<installdir>/owareapps/lib/installed.properties.`

Save the file, restart the applcation server and resync devices to update the equipment name.

com.dorado.devicedriver.base.updateName=sysname_ip

com.dorado.devicedriver.base.updateName= hostname_ip

com.dorado.devicedriver.base.updateName=sysname

com.dorado.devicedriver.base.updateName= hostname

com.dorado.devicedriver.base.updateName=ip

com.dorado.devicedriver.base.updateName=false

This last setting will disable update of the equipment name on resync. This property will control whether resync updates contact and location.

com.dorado.devicedriver.base.updateCo ntactAndLocation=true|false

true will refresh on the contact and location on resync. false will not update information on resync.

Network

The Network Panel collects the network (IP range, hosts, and so on) and the authentication information for the discovery profile. After you click *Next,* the *Network* panel appears.

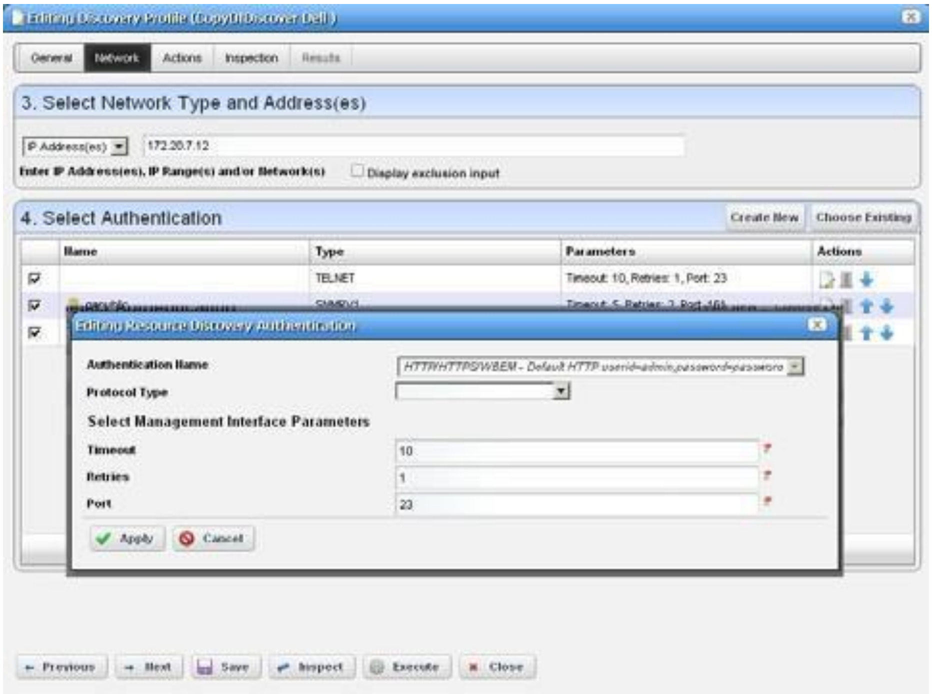

3. Select Network Type and Address(es)

IP Address(es): 172.20.7.12

Enter IP Address(es), IP Range(s) and/or Network(s) □ Display exclusion input

4. Select Authentication

	Name	Type	Parameters	Actions
☑		TELNET	Timeout: 10, Retries: 1, Port: 23	
☑	public	SNMPV1	Timeout: 5, Retries: 2, Port: 161	
☑				

Authentication Name: HTTP/HTTPS/WBEM - Default HTTP userid=admin,password=passwora

Protocol Type:

Select Management Interface Parameters

Timeout: 10

Retries: 1

Port: 23

The Network panel includes settings for:

- **Network Type and Addresses**, where you select the type of entry in the pick list (IP Address(es), CIDR Address, Hostname, SNMP Broadcast, Subnet*)*.
  The tooltips tell what valid entries look like.

NOTE:

Cruz now discovers all IP addresses in a specified range, regardless of the specified base IP address is (middle, starting IP, or last in the range). IP addresses outside of range will not be discovered. You can use the CIDR specification of the network to discover rather that the subnet ID.

You can exclude IP addresses, or ranges of IP addresses if you check the *Display exclusion input* checkbox and input the addresses you want excluded as you did for those you entered in the *Address(es) for Discovery* field. Such exclusions only apply to the profile where you enter them. To exclude an address or range, use the `com.dorado.redcell.discovery.exclude` property. Examples of how to enter such exclusions appear in the `redcell.properties` file under `owareapps\redcell\lib`. As always, best practice, if you want such properties to persist is to put the property in `owareapps\installprops\lib\installed.properties.`

- **Authentication**, where you can create new, or add existing authentications. See Authentications on page 160 for the way to create such authentications outside the discovery process.

△ CAUTION:

If adevice orits driver requires twoauthentications andyouonly enter one, it maynotappearin inventory after discovery. To correct this, enter both authentications in the Discovery Profile or in Quick Discovery. If you discover a device partially with only one authentication—typically the SNMP community—you can re-discover with the correct authentications later or *Edit* the resource to add that correct authentication *and* the management interface for it.

Note that authentications appear with Edit/Delete icons and Up/Down arrows on their right. The Edit icon opens the authentication editor. Click the arrows to arrange the order in which the application tries credentials (top first). Ordering only applies when two credentials are of the same type.

If you have imported a discovery profile without importing or creating the authentications it uses, editing its authentications is not possible. If you cannot import authentications or have not created them when you do attempt to edit them, the easiest solution is to delete the un-imported un-created authentication the profile refers to and create a new one.

If two similar authentications include one with a **deeper**, enable login, and a “shallower” one without that additional login, arrange to try the deeper login first. If the device rejects it, discovery still tries the shallower one later.

Actions

When you click Next, the Actions panel appears.

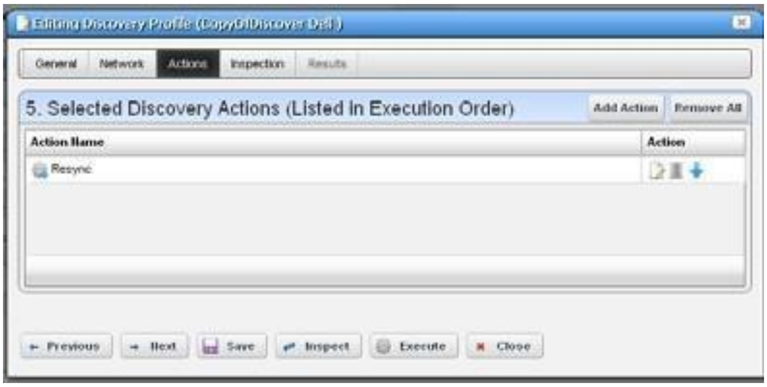

Editing Discovery Profile (CopyOfDiscover Dell)

5. Selected Discovery Actions (Listed In Execution Order)

Action Name	Action
Resync	

You can accept the default actions that appear here (like Resync, adding the device to a global Scheduled Resync, link discovery, and so on) by clicking Next to the Inspection portion of discovery. (See also Configuring Resync on page 104).

Alternatively, you can click **Add Action**. This opens a screen with a list of available actions. Click *Apply* to select an action to add to the list for this profile.

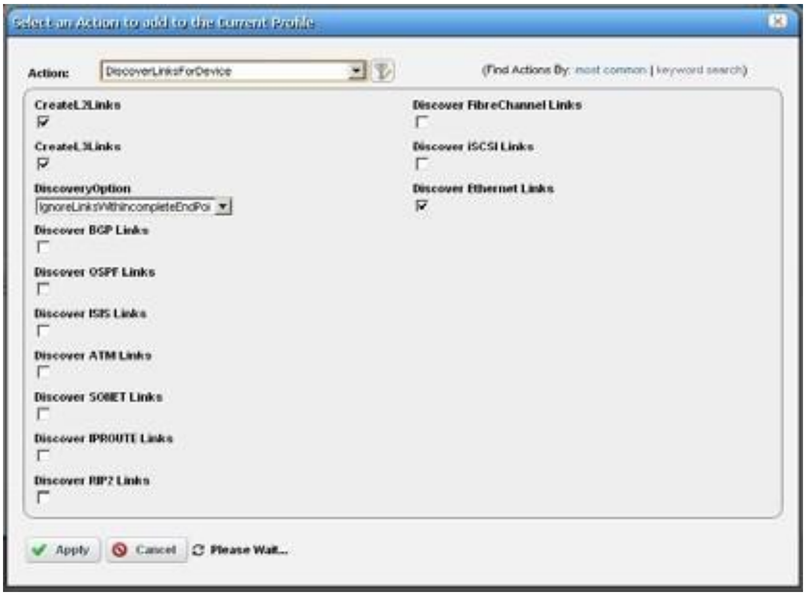
Select an Action to add to the Current Profile
(Find Actions By: most common | keyword search)
Apply
Cancel
Please Wait...

Action: DiscoverLinksForDevice

CreateL2Links ☑

CreateL3Links ☑

DiscoveryOption IgnoreLinksWithIncompleteEndPoi

Discover BGP Links □

Discover OSPF Links □

Discover ISIS Links □

Discover ATM Links □

Discover SONET Links □

Discover IPROUTE Links □

Discover RIP2 Links □

Discover FibreChannel Links □

Discover iSCSI Links □

Discover Ethernet Links ☑

Notice that the default is for the *most common* actions, but you can also click *keyword search* to display a search field instead ofa pick list with the most common actions. The search results appear in the pick list. When you select an item, if it has parameters, they appear listed below that item. Select the options or from the pick list to configure these parameters, then click *Apply* to select this action as part of the profile. The screen appearance changes depending on the selected action.

The **Edit, Delete, Move** icons appear to the right of each action. If you Edit a profile with parameters, you can change them. The screen looks like the one that appears when you Add actions. Deleting actions removes them from the list, and the Move arrows help arrange the order in which actions appear listed and are executed. The list of actions the profile executes goes from top-to-bottom.

Inspection

Using the Inspection panel is optional. If you want to execute the profile after entering the required information on the General and Network panels, you can skip this step, and just click Execute. The Inspectionpanelletsyou previewthe discoveryprofile'sactionsand accessto devices. Ifyou clicked *Next* rather than *Inspect* at the bottom of the previous screen, click *Start Inspection* in the top right corner of this screen to begin the inspection process that validates the device's credentials.

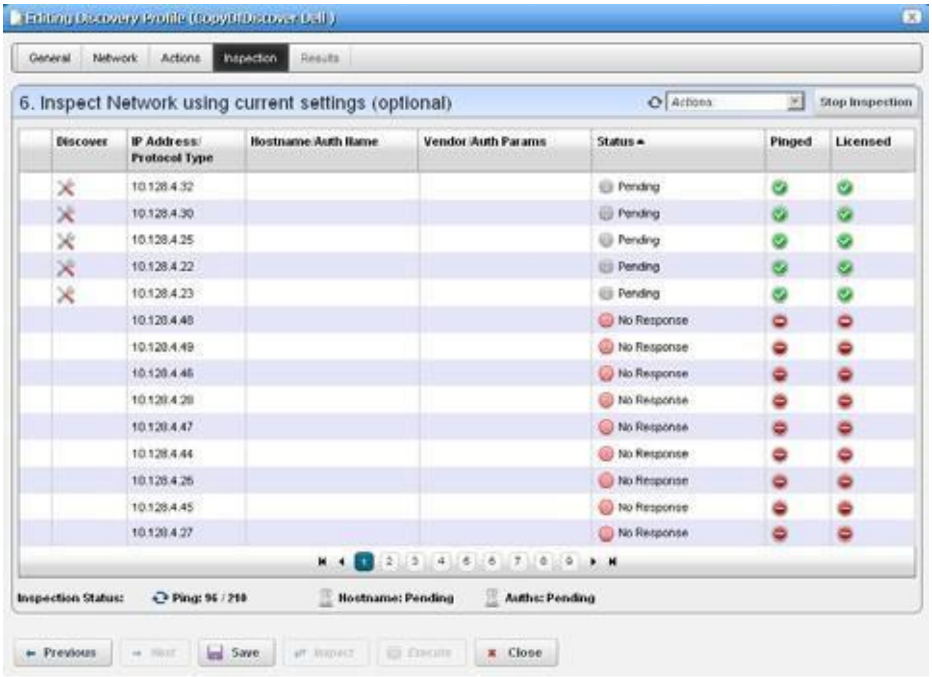

Editing Discovery Profile (CopyOfDiscover Dell)

6. Inspect Network using current settings (optional)

	Discover	IP Address/ Protocol Type	Hostname/Auth Name	Vendor/Auth Params	Status ▲	Pinged	Licensed
	✕	10.128.4.32			Pending	✓	✓
	✕	10.128.4.30			Pending	✓	✓
	✕	10.128.4.25			Pending	✓	✓
	✕	10.128.4.22			Pending	✓	✓
	✕	10.128.4.23			Pending	✓	✓
		10.128.4.48			No Response	⊖	⊖
		10.128.4.49			No Response	⊖	⊖
		10.128.4.46			No Response	⊖	⊖
		10.128.4.28			No Response	⊖	⊖
		10.128.4.47			No Response	⊖	⊖
		10.128.4.44			No Response	⊖	⊖
		10.128.4.26			No Response	⊖	⊖
		10.128.4.45			No Response	⊖	⊖
		10.128.4.27			No Response	⊖	⊖

1 2 3 4 5 6 7 8 9

Inspection Status: Ping: 96 / 210 Hostname: Pending Auths: Pending

Notice that the Inspection Status fields indicate the success or failure of Ping, Hostname resolution, and Authentications, and the Status column displays whether a valid authentication exists, whether it has been tested, and whether the test is successful.

When authentications are unsuccessful, click the icons to their right to remove or edit them. You can also click the wrench/screwdriver “fix it” icon in the *Discover* column to open an editor where you can revise the authentications for that device.

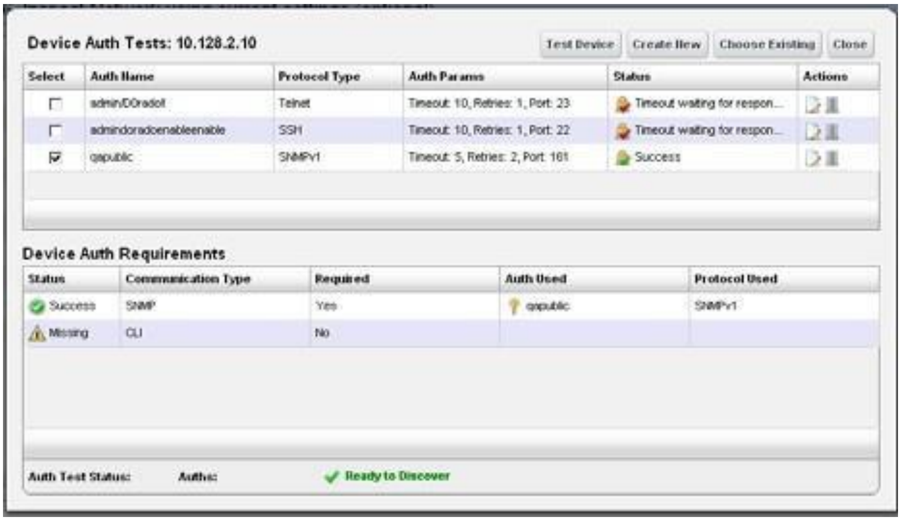

Device Auth Tests: 10.128.2.10

Select	Auth Name	Protocol Type	Auth Params	Status	Actions
☐	admin/DOradot	Telnet	Timeout: 10, Retries: 1, Port: 23	Timeout waiting for respon...	
☐	admindoradoenableenable	SSH	Timeout: 10, Retries: 1, Port: 22	Timeout waiting for respon...	
☑	qapublic	SNMPv1	Timeout: 5, Retries: 2, Port: 161	Success	

Device Auth Requirements

Status	Communication Type	Required	Auth Used	Protocol Used
✓ Success	SNMP	Yes	qapublic	SNMPv1
⚠ Missing	CLI	No		

Auth Test Status: Auths: ✓ Ready to Discover

Create New lets you create new authentications, *Choose Existing* lets you select from existing authentications, *Test Device* lets you try out the authentications you have selected, and *Close* closes this screen. Notice that you can configure new or existing authentications’ port, retry and timeout settings before you click *Apply* (or *Cancel*) in the authentication editor that appears after clicking the “Fix it” button.

Save preserves the profile. Youcan then right-click it to select *Execute.* If you select *Execute* from the profile editor, Cruz does not save the profile to execute later.

Execute begins discovery, and the message traffic between Cruz and the device appears on the *Results* panel.

Results

This is a standard Audit Trail/Job Status window displaying the message traffic. See also Audit Trail/Job Status on page 148 for more about retrieving audit trailarchives.

A message (*Discovery Profile Execute is complete*) appears in the *Messages* at the bottom left of the status bar.

Click the X in the top right corner of the discovery profile editor to close it.

Discovery Profiles Expanded

This larger view offers a *Reference Tree* snap panel where you can see the connection between a selected profile and the authentications and discovery tasks it includes.

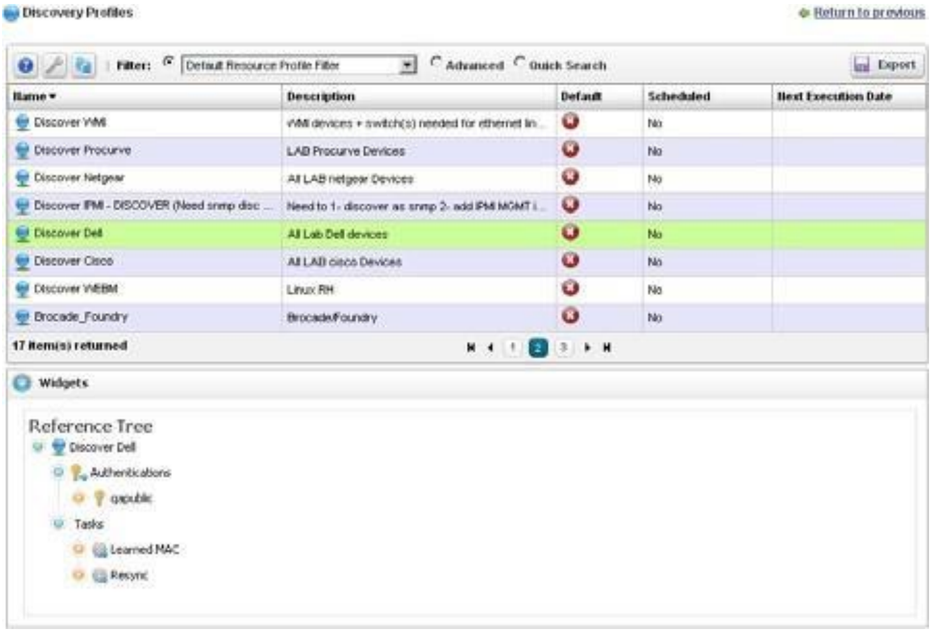
Discovery Profiles
Return to previous
Filter: ● Default Resource Profile Filter ○ Advanced ○ Quick Search
Export

Name ▾	Description	Default	Scheduled	Next Execution Date
Discover VMI	VMI devices + switch(s) needed for ethernet lin...	☒	No	
Discover Procurve	LAB Procurve Devices	☒	No	
Discover Netgear	All LAB netgear Devices	☒	No	
Discover IPMI - DISCOVER (Need snmp disc ...	Need to 1- discover as snmp 2- add IPMI MGMT i...	☒	No	
Discover Dell	All Lab Dell devices	☒	No	
Discover Cisco	All LAB cisco Devices	☒	No	
Discover WEBM	Linux RH	☒	No	
Brocade_Foundry	Brocade/Foundry	☒	No	

17 item(s) returned

Managed Resource Groups

The managed resource groups make acting on several devices at once more convenient. They also make managing groups ofdevices possible. The summary screen displays columns describing the group *Name*, *Type,* and *Icon.* Typically, a variety of Dynamic Groups come with your package. For example, discovered devices are all added to *All Devices* automatically.

By default, this portlet is available from the Settings page by selecting the Groups & Locations menu option.

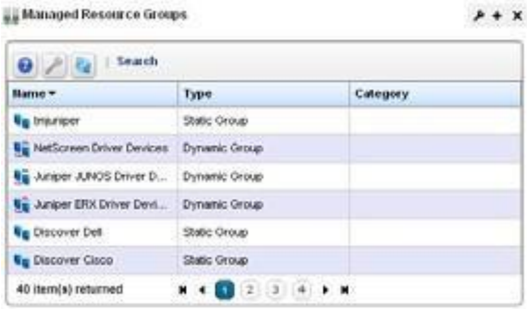

Name ▾	Type	Category
Imjuniper	Static Group	
NetScreen Driver Devices	Dynamic Group	
Juniper JUNOS Driver D...	Dynamic Group	
Juniper ERX Driver Devi...	Dynamic Group	
Discover Dell	Static Group	
Discover Cisco	Static Group	

40 item(s) returned

You can also right-click to do the following:

New—Lets you make either a Static Group (one in which you select devices) or a Dynamic Group (one in which a filter selects devices). See details of these screens below.

Edit — This opens the same editors as *New*, populated with the information for the selected group.

Edit Resources—Lets you edit resourcesassociated with the selected group like its location, contact, or whether to manage it by hostname.

Topology — Displays a topology map of the selected group. See Presentation Capabilities for more.

Actions — Select from a sub-menu of actions available for the group. This includes Adaptive CLIs. Select from a sub-menu of Adaptive CLI.

If you want the target to be a group of devices, Ctrl + click the target devices before your right-click to invoke an action, or right-click in the Groups portlet and select *Actions* there.

File Management > Backup, Restore, Deploy —Lets you callon Cruz's NetConfig configuration file backup, restore and deploy capabilities. See Backing Up Configurations on page 462 for an example of the steps this follows. See also Configuration File Compare Window on page 449 and more about deploying updates to the OS for the selected resource group. See Deploying Firmware on page 465 fordetails.

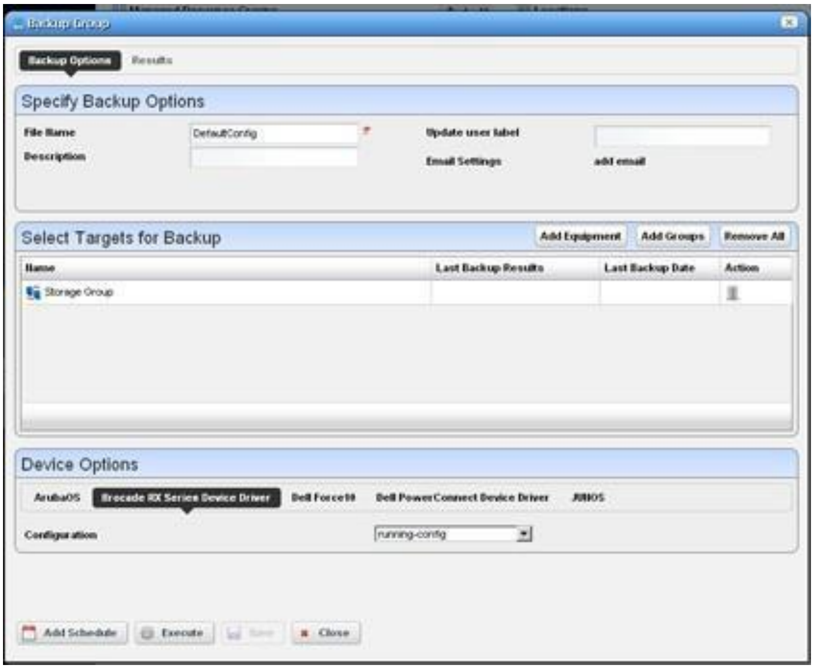

File Name: DefaultConfig

Description:

Name	Last Backup Results	Last Backup Date	Action
Storage Group			

Configuration: running-config

When you select a group backup, and the group contains devices of several types, the *Device Options* panel displays a tab for each device type. Select the backup parameters there before executing or scheduling backup.

NOTE:

Some devices merge rather than replace configurations when you select Restore. (Cisco XR, for one)

You cantune resources consumed byprocesseslikebackup orresync. See How to: Tuning Application Features' Performance Impact on page 114

Link Discovery —Discover links between members of the selected group, and others. See New Link on page 187 and Link Discovery on page 188 for details.

Resync Resources—Queries the devices in the group to update Cruz's database. Resyncing also resyncs alarms on the selected device.

Delete—Remove the selected group from inventory. The devices remain in inventory, but this removes the grouping.

Import/Export—Export the selected config file to disk or import it fromdisk. Youcan also import/ export a selected configuration file.

Provides thefollowingactions when availablefor theselectedimage:

- Import retrieves a file containing XML compliance descriptions. Some imports can come from a URL.
- Export Selection exports the selected description to an XML file.
- Export All exports all descriptions to an XML file.

Click Download Export File to specify where to save the file.

The Import/Export option is useful as a backup or to share descriptor s with other projects.

Youmust import data into the correctportlet. Forexample, youcannotimport event data into the Actions portlet.

If one type of data depends on another, you must import the other data before importing the data that depends on it.

Share with User—Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Cruz does not support static groups that include members retrieved by (dynamic) filters. Youcan configure membership with dynamic resource groups that include group memberships as filter criteria. For example, you can create a filter for members of ResourceGroupABC or members of ResourceGroupXYZ.

Expanded Managed Resource Groups

The expanded Managed Resource Groups screenlets you see the summary screen’s groups with a Reference Tree snap panel that displays a selected group’s connection to its devices and any assigned monitors.

Static Group

Selecting *Static Group* as the type to create displays a selector screen where you can *Name* and select a *Category* for the group, then search for available resources with a filter. Click *Apply Filter* after you have configured it, and a list of devices fitting its criteria appears. Select device(s) and click *Add Selected*, or simply click *Add All* to add the entire list to your static group. Notice that you can continue to re-use this filter to list devices and continue to select them.

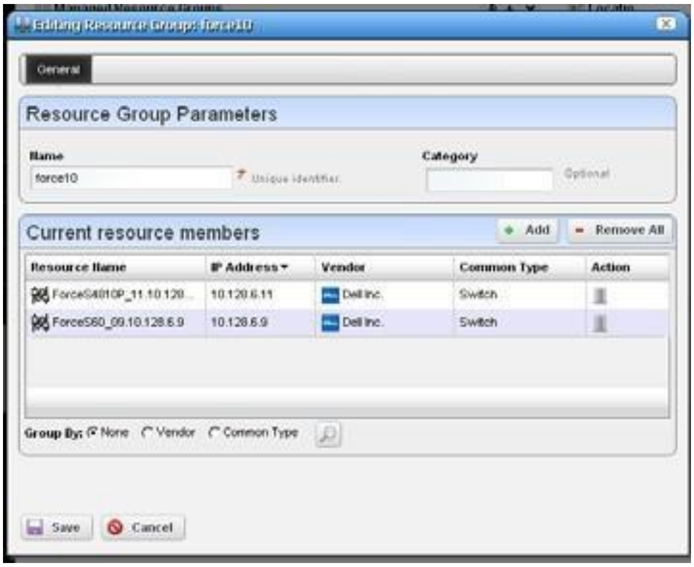

Resource Group Parameters

Name: force10 * Unique identifier.

Category: Optional

Current resource members

Resource Name	IP Address	Vendor	Common Type	Action
ForceS4810P_11.10.128...	10.128.6.11	Dell Inc.	Switch	
ForceS60_09.10.128.6.9	10.128.6.9	Dell Inc.	Switch	

Group By: ● None ○ Vendor ○ Common Type

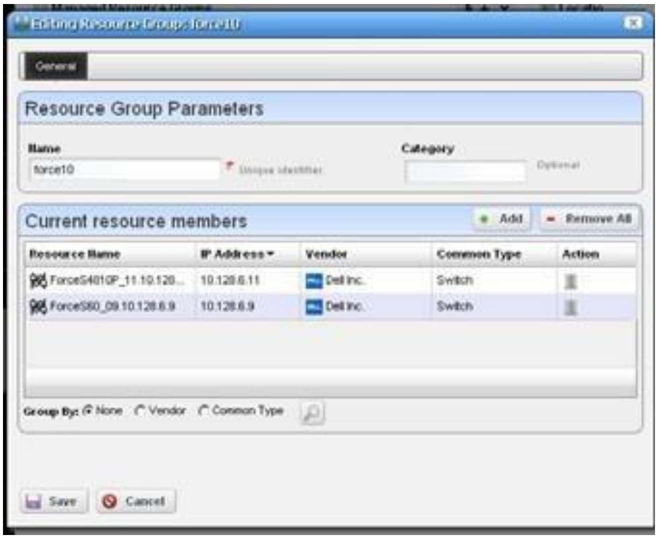

Resource Group Parameters

Name: force10 * Unique identifier.

Category: Optional

Current resource members

Resource Name	IP Address	Vendor	Common Type	Action
ForceS4810P_11.10.128...	10.128.6.11	Dell Inc.	Switch	
ForceS60_09.10.128.6.9	10.128.6.9	Dell Inc.	Switch	

Group By: ● None ○ Vendor ○ Common Type

When you select a device, it no longer appears listed. When you click *Done* the subsequent screen displays all devices you have selected. You can click *Add* on this screen to return to the previous screen (or *Remove All* to delete the listed devices from the group). At the bottom of this screen, you can also elect to group devices by *None*, *Vendor* or *Common Type* (Switch, Router, and so on). These last two create “trees” with nodes for each vendor or type. Youcan also click the magnifying glass to search through listed devices. Clicking *Remove All* removes all devices in the group.

Click *Save* to preserve the group you have configured.

Dynamic Group

In contrast to Static Groups, Dynamic Groups do not let you select individual equipment. You simply configure a filter, and Cruz creates the group on the fly. After you enter the *Name* and *Category* for the group, create the filter. To see what the group would look like, click *Preview Group*. This opens the *Preview* tab, concealing the *General* tab. To return to *General*, click that at the top of the screen. Click *Save* to preserve the group configuration or *Cancel* to exit without saving.

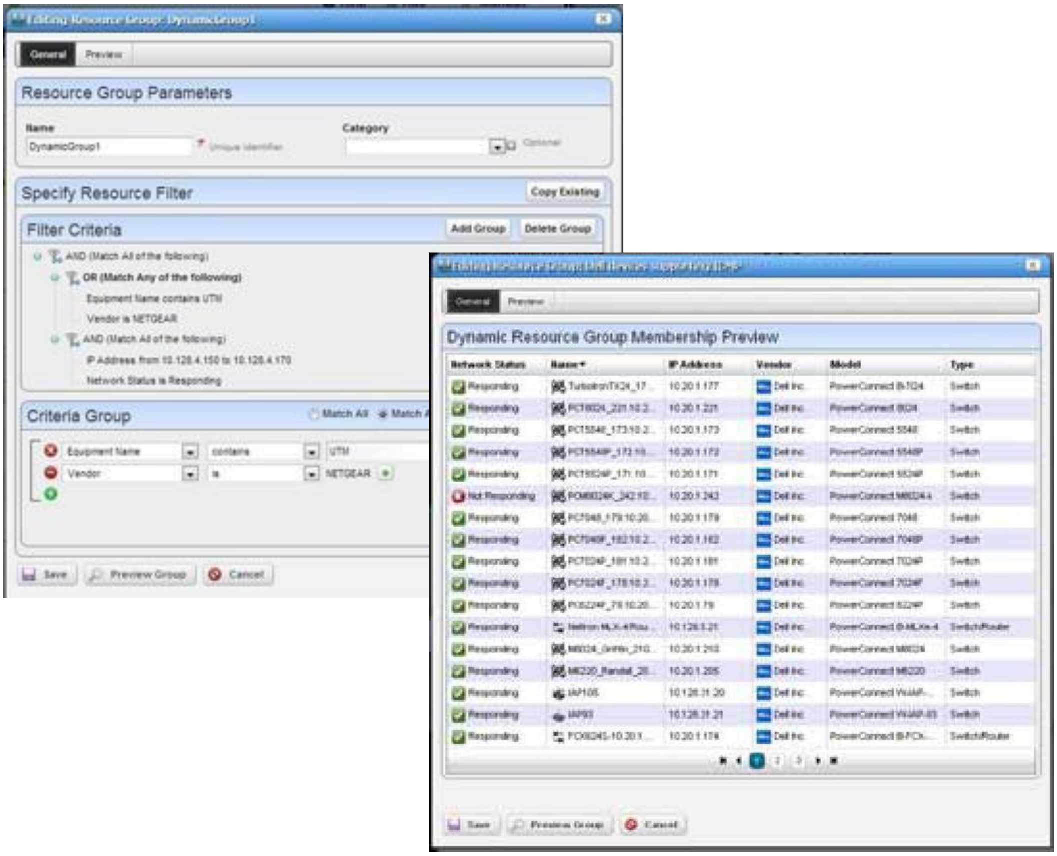

Name: DynamicGroup1 (Unique Identifier)

Category: (Optional)

- AND (Match All of the following)
  - OR (Match Any of the following)
    - Equipment Name contains UTM
    - Vendor is NETGEAR
  - AND (Match All of the following)
    - IP Address from 10.128.4.150 to 10.128.4.170
    - Network Status is Responding

Equipment Name	contains	UTM
Vendor	is	NETGEAR

Network Status	Name▾	IP Address	Vendor	Model	Type
Responding	TurbotronTX24_17...	10.20.1.177	Dell Inc.	PowerConnect B-TX24	Switch
Responding	PCT8024_221 10.2...	10.20.1.221	Dell Inc.	PowerConnect 8024	Switch
Responding	PCT5548_173 10.2...	10.20.1.173	Dell Inc.	PowerConnect 5548	Switch
Responding	PCT5548P_172 10...	10.20.1.172	Dell Inc.	PowerConnect 5548P	Switch
Responding	PCT5524P_171 10...	10.20.1.171	Dell Inc.	PowerConnect 5524P	Switch
Not Responding	PCM8024K_242 10...	10.20.1.242	Dell Inc.	PowerConnect M8024-k	Switch
Responding	PCT7048_179 10.20...	10.20.1.179	Dell Inc.	PowerConnect 7048	Switch
Responding	PCT7048P_182 10.2...	10.20.1.182	Dell Inc.	PowerConnect 7048P	Switch
Responding	PCT7024P_181 10.2...	10.20.1.181	Dell Inc.	PowerConnect 7024P	Switch
Responding	PCT7024F_178 10.2...	10.20.1.178	Dell Inc.	PowerConnect 7024F	Switch
Responding	PC6224P_78 10.20...	10.20.1.78	Dell Inc.	PowerConnect 6224P	Switch
Responding	NetIron MLX-4 Rou...	10.128.1.21	Dell Inc.	PowerConnect B-MLXe-4	Switch/Router
Responding	M8024_Griffin_210...	10.20.1.210	Dell Inc.	PowerConnect M8024	Switch
Responding	M6220_Randall_20...	10.20.1.205	Dell Inc.	PowerConnect M6220	Switch
Responding	IAP105	10.128.31.20	Dell Inc.	PowerConnect W-IAP...	Switch
Responding	IAP93	10.128.31.21	Dell Inc.	PowerConnect W-IAP-93	Switch
Responding	FCX624S-10.20.1...	10.20.1.174	Dell Inc.	PowerConnect B-FCX...	Switch/Router

Managed Resources

Use this Resource Management portlet instance to manage the configuration of devices that you discovered or created on your network. Resource Management portlets show device-specific information, both general (name, type, location, contact) and technical (vendor, subcomponents, and so on). The default layouts in the summary and expanded views are geared towards device configuration management.

This portlet is intended for anyone that wants to manage their network device’s configuration.

Access this portlet by selecting the Configuration Management page. This portlet has both a summary view and an expanded view. Each viewdisplays different columns and have the same pop-up menu options available.

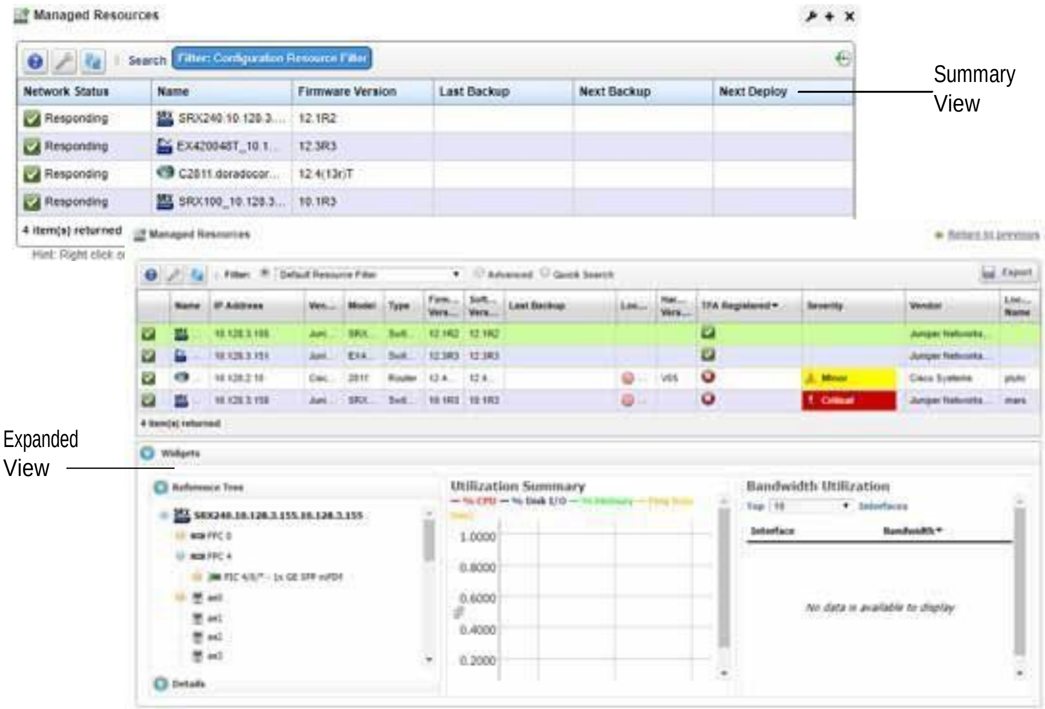

Network Status	Name	Firmware Version	Last Backup	Next Backup	Next Deploy
☑ Responding	SRX240.10.128.3...	12.1R2			
☑ Responding	EX420048T_10.1...	12.3R3			
☑ Responding	C2811.doradocor...	12.4(13r)T			
☑ Responding	SRX100_10.128.3...	10.1R3			

4 item(s) returned

	Name	IP Address	Ven...	Model	Type	Firm... Vers...	Soft... Vers...	Last Backup	Loc...	Har... Vers...	TFA Registered	Severity	Vendor	Loc... Name
☑	...	10.128.3.155	Juni...	SRX...	Swit...	12.1R2	12.1R2				☑		Juniper Networks...	
☑	...	10.128.3.151	Juni...	EX4...	Swit...	12.3R3	12.3R3				☑		Juniper Networks...	
☑	...	10.128.2.10	Cisc...	2811	Router	12.4...	12.4...		● ...	V05	●	Minor	Cisco Systems	pluto
☑	...	10.128.3.158	Juni...	SRX...	Swit...	10.1R3	10.1R3		● ...		●	Critical	Juniper Networks...	mars

4 item(s) returned

Other that the Reference Tree, Detailsproperties Utilization Summary, and Bandwidth Utilization widgets, theexpandedviewincludesa **Detailssettings** widget. The Detailssettingswidgetincludes the *system Object Id*, *Date Created* (that is, discovered), *Creator* (the user who performed discovery), *Install Date*, *Administrative State* (Locked [Device use is prohibited] Shutting Down [Only existing users can use the device] Unlocked [Normal use of device is permitted]) , *Operational State* (Disabled [Inoperable because of a fault, or resources are unavailable] Enabled [Operable and available for use] Active [Device is operable and currently in use with operating capacity available to support further services] Busy [Operable and currently in use with no operating capacity to spare]).

Cruz includes a task that updates the operational and administrative states of a port or interface when an event processing rule (EPR) responds to the events listed below. Five automation EPRs respond to these events and invoke this update task with the target from the entity associated with the event. Only subcomponents are affected. The EPRs are as follows (impacts in parenthesis follow them):

monitorTargetDown (operational state = Down)

monitorTargetUp (operational state = Up)

monitorTargetIndeterminate (operational state = Unknown)

linkDown (operational state = Down, administrative state = the value of the var bind ifAdminStatus)

*linkUp (*operation state = The value of the var bind ifOperStatus, administrative state = Up)

Columns

Other than the general navigation and configuration options, this instance of the Managed Resourcesportlets (summaryandexpandedviews) include thefollowingcolumns. Thecolumns displayed by default are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Network Status	The resource's status in the network, such as: • **Responding** means that this application gets a responsefrom the device through a network protocol. • **Not Responding** means that the device does not respond to the protocol. • **Indeterminate** means that the monitoring software has not tried to reach the device or there was some other error that prevented us from determining one of the other two statuses. The Network Status option availability depends on its response to the ICMP monitor. If you disable the monitor (for example, for performance reasons), a status may appear, but it is not meaningful. This field displays on the summary view by default.
Name	The device name. This field displays on the summary and expanded views by default.
Firmware Version	The version for the device's firmware. This field displays on the summary and expanded views by default.
Last Backup	The date and time that the device was last backed up. You can sort on this column. This field displays on the summary and expanded views by default.
Next Backup	The next date and time that the device is scheduled to back up. This field displays on the expanded view by default.
Next Deploy	The next date and time that the device is scheduled to deploy. This field displays on the expanded view by default.
IP Address	The device's IP address. This field displays on the expanded view by default.
Vendor	The vendor for this device. This field displays on the expanded view by default.
Model	The device model. This field displays on the expanded view by default.
Type	The device type, such as switch, router, and so on. This field displays on the expanded view by default.
Software Version	The version for the device's software. This field displays on the expanded view by default.
Location	The device's location shown with text and an icon. expanded view
Hardware Version	The device's version. This field displays on the expanded view by default.
TFA Registered	An indicator that traffic flow analyzer (TFA) is registered (checkmark) or not (). This field displays on the expanded view by default.

Column	Description
Alarm Severity	The highest alarm severity opened for a device and /or its subcomponents. Alarms against the device as a whole are always considered and this also includes all alarms that are against subcomponents (ports, interfaces, etc.) and whose Resource Propagation is "Impacts Top Level" or "Bot h" but it does not include alarms against subcomponents whose Resource Propagation is "Default" or "Impacts Subcomponents". This field displays on the expanded view by default.
Vendor Name	The resource's vendor name. This field displays on the expanded view by default.
Location Name	The location where the resource resides. Shows both text and icon. This field displays on the expanded view by default.
Widgets	Additional information about the selected rule, such as: • Reference Tree shows the device and connected components • Details shows information about equipment name, vendor, location, contact, icon, and last modified and discovery date • Utilization Summary graphs device utilization, such as CPU, disk I/O, memory and ping rate • Bandwidth Utilization graphs the device's bandwidth utilization. Notice that you can change the number of top interfaces graphed, when this is applicable. See also Bandwidth Calculation on page 410 The Widgets field is available only from the expanded view.
Additional IP Addresses	Any additional IP addresses assigned to the resource, other than the primary IP address. For example, a resource might have an IPv4 address and a IPv6 address. If this is the case, the primary IP address appears in the IP Address column and the additional IP addresses appear in this column.
Alarm Suppression Description	The description provided for an alarm suppression configuration.
Alarm Suppression Mode	Indicates if the device is in alarm suppression, and if so, which mode of suppression is active. The two modes of active alarm suppression are "Ad Hoc", which means that suppression will remain indefinitely until it is stopped manually, and "Scheduled", which means that suppression will automatically stop at some point in the future (see the Schedules portlet). If the device is not in alarm suppression, then a dash "-" is listed. If some form of alarm suppression is listed, the device does not receive alarms.
Asset Tag	Any asset tag identifier for the device.
Contact	Any contact associated with the resource.
Creator	The user or process (such as, a discovery rule or device driver) that created this resource.
DNS Hostname	The resource's host name.
Description	A text field for a description of the resource.
Discovery Date	The date and time the resource was discovered.
Domain ID	The identifier for the device's domain. This is populated in a multitenant environment.
Foreign ID	An optional identifier for the customer's organization.
Install Date	The date and time that the resource was installed.
Last Configuration Change	The time/date stamp for the last configuration change detected. This is driven from Automations (which are a category of Event Processing Rules) whose configured action is "Config Changed"

Column	Description
Last Modified	The date and time that the last modification occurred.
Last Status Change	The last time that the Network Status changed. For example, if the device has been not responding for 8 hours continuously (but was previously responding) then this value will give the time that that the status changed to "Not Responding"
MAC Address	The resource's Machine Address Code.
Manage By Hostname	An indicator that shows whether the device is managed by host name rather than IP address.
Management State	The management state of the resource.
Mediation Server 2 IP Address	The resource's secondary mediation server IP address (applies in a high availability system with two mediation servers).
Mediation Server IP Address	The resource's mediation server IP address.
Next Restore	The next date and time to restore a resource.
Notes	A text field for information about the resource.
Operational State	The resource's state independent of its network status. Valid values are: • Disabled indicates that the resource is inoperable because of a fault, or it is not available. • Enabled indicates that the resource is operable and available for use. • Active indicates that the resource is operable and currently in use with operating capacity available to support further services. • Busy indicates the resource is operable and currently in use with n operating capacity to spare.
Physical Index	The given index associated to the resource. This is useful when a vital device component, such as a power supply or fan has sensors attached.
RTM Category	The resource's right-to-manage category.
Serial Number	The resource's serial number.
Service Tag	The device's service tag.
SupportAssist Status	The device's SupportAssist status. Valid values are Opt-In and Opt-Out.
System Object ID	The identifier for the resource's system object.
Topology Icon Size	The equipment icon size. Modify this size by right clicking a device and then selecting Edit.

Pop-Up Menu

The Managed Resources pop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
New	You can create a new device without discovering it with this menu item. Select the device vendor, model and type in the next screen, then fill in the information about the device in the editor that appears after that selection.
Edit	Opens the Resource Editor , where you use modify a resources options and parameters.

Menu Option	Description
Details (Click +Shift)	Opens a Details screen for the selected resource. This contains information like the performance indicators, power supplies and fans, monitor status summary, reference tree, network (ports, interfaces, details, and associated links), alarms and history events, audit trail and execution history, performance data, and maintenance information.
Manage	Provides access to the following actions: • Set Management State for the selected devices to one of the following states: • Normal — The device is unconstrained by the other Administrative States. Changing from Suspended to Normal stops alarm suppression. Standard access, and inclusion in right-to-manage count. • Decommissioned — While this device is in inventory, it is not active. No device access allowed, no Monitor associations, no event processing, no Management Interfaces, no Authentication, no links, and no services are permitted. • Suspended — Suspends all device-related activities. No device access allowed, Monitoring Suspended, No event processing, Counts againstright-to-manage. • Planned — The planned (future) device. No device access allowed, no monitor associations, and no event processing. • Maintenance — Neither alarms or polling apply to the device. Does allow resync and Adaptive CLI. Standard device visibility. The Management State menu option is available if the write permissions are functional. • Set Domain Access Control Select the Multitenancy domain where the selected device is to be visible, or manageable. This only appears if you implemented the Multitenancy option. Refer to the Cruz *Installation Guide* for details about that option. • View the Maintenance Log for the selected device. See Connected Devices on page 189 for more about maintenance logging capabilities.
Topology	Opens the Topology portlet, where you define a topology. It also displays a topology map that includes the selected resources. See Presentation Capabilities on page 223 for more about these maps.
Actions	Displays a list of action that you can initiate, such as Adaptive CLI Actions and other actions specific to the selected device. The list's content depends on the device selected. Use the magnifying glass search option to narrow the list of actions. Select an action and then click Load Selected to run it manually. See Actions Portlet on page 541 for more about configuring these activities.
Change Management	Executes or schedules for execution one of the following policies based on the option selected: • Change Determination runs the selected device. (See Change Determination Process on page 506.) • Execute Proscan runs any ProScan (also known as Compliance Policy) connected to the selected device. • Execute Proscan Policy runs any Proscan. (See Compliance Policy Summary on page 482.)
Direct Access	Open an SNMP Mib Browser to the alarmed device, a CLI Terminal (Telnet window) to the alarmed device, or ICMP Ping the device alarmed. Only those available appear in the subsequent menu.

Menu Option	Description
Event Management	Suppresses or updates alarms related to the selected resource. The following event management options available are: • **Start Alarm Suppression** starts alarm suppression after you enter a description and then click Start. A Success/Failure confirmation message is displayed. • **Stop All Alarm Suppression** discontinues all alarm suppression. A confirmation message is displayed. You can also stop alarm suppression using the View Active Suppressions option. • **Schedule Alarm Suppression** opens the New Schedule *suppression Action* window, where you set the exception options and suppression targets parameters and then set the schedule options. • **View Active Suppressions** displays a list of active suppressions. A message is displayed if there are no active suppressions. You can stop selected alarm suppressions from this list. • **View Event History** shows event history for the selected resource in the Event History portlet. • **Resync Alarms** refreshes the information displayed with the latest information in the database for the selected resource or all resources. Except for View Event History, perform these actions on one or more resources.
File Management	Displays a current configuration file, compares it to previous backups, backs up, restores, imports or exports a configuration file by selecting the appropriate option. You can also deploy firmware to devices from this menu. If you go to the Configuration Files portlet, you can also edit backed up configuration files. See Configuration File Compare Window on page 449 for details.
Links	Create a new link or discover links between members of the selected group, and others. See New Link on page 187 and Link Discovery on page 188 for details.
Performance	Shows performance data for the selected interfaces in the Performance Dashboard window. Click the edit tool to modify the dashboard view properties, entities, or attributes. See Dashboard Editor on page 420 for more about the editor and its options.
Resource Groups	Adds the selected device to new Dynamic or Static groups, or to existing groups. See for Managed Resource Groups on page 173 more about this.
Resync	Queries the device again for more current information, including alarms.
SupportAssist	Changes whether the selected devices SupportAssist status is Opt-In or Opt-Out. This option is available only if you have configured the Cruz SupportAssist feature.

Menu Option	Description
Traffic Flow Analyzer	Configures the selected device to appear in the Traffic Flows displays when you select the Register option. Selecting UnRegister removes the selected device from the Traffic Flows displays. See Traffic Flow Analyzer on page 515 for more details on the analyzer. There are two Show Traffic menu options available from this sub-menu, both of which will navigate to full-screen views of the expanded Traffic Flow Portlet and will show the traffic flow data for the selected device. Show Traffic as Endpoint is available for all managed devices and this option navigates to a page showing the traffic where this device was either the sender or the receiver of the flow (as endpoints are essentially senders and receivers combined). Show Traffic as Exporter is available only for registered exporters and this option shows all traffic going through the device. You also have the option to execute or schedule execution for: • Create Configuration • Remove Configuration • Show Configuration
Delete	Removes the selected devices from inventory.
Edit Custom Attributes	Opens the Custom Attribute Editor where you define field characteristics, such as whether it is enabled, the label name, and the tooltip.
View as PDF	Creates an Acrobat PDF document containing the selected alarms' content displayed in the portlet.
Share with User	Opens the Share with User window where you select the colleague you want to share the selected alarms with and then type your message.

Resource Editor

The Resource editor allows you to modify the following options and parameters:

- Resource Editor: General
- Resource Editor: Authentication
- Resource Editor: Management Interface
- Resource Editor: Custom Attributes—This tab appears only if you have configured custom attributes. See Redcell > Data Configuration on page 40 for more about them.

Click *Save* to preserve any changes made in these screens to Cruz's database, or *Close* to abandon any changes made in editor screens. Unless the device is a printer, changes to these screens typically make database changes, not changes on the device.

NOTE:

You can edit fields like Notes and Description insubcomponent cards byright clicking them in the resource tree.

Resource Editor: General

This screen may vary for different kinds of devices. Its *General Details* panel displays the *Name, Description*, *Vendor, Location, Contact,* and *Equipment Icon* for the selected device.

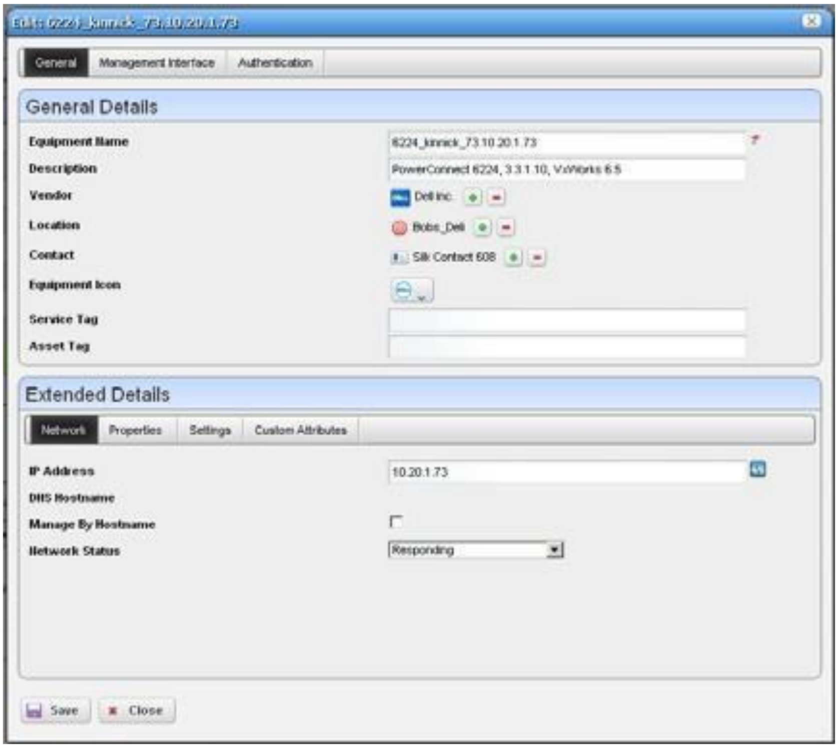

Edit: 6224_kinnick_73.10.20.1.73

General Details

Equipment Name	6224_kinnick_73.10.20.1.73
Description	PowerConnect 6224, 3.3.1.10, VxWorks 6.5
Vendor	Dell Inc.
Location	Bobs_Deli
Contact	Silk Contact 608
Equipment Icon	
Service Tag	
Asset Tag	

Extended Details

IP Address	10.20.1.73
DNS Hostname	
Manage By Hostname	□
Network Status	Responding

The *Extended Details* panel includes *Network, Properties* and *Settings* tabs. These let you view or alter things like *IP Address, DNS Hostname, Manage by Hostname, Network Status, Model* and *Equipment Type, Serial Number, Software Version Firmware* and *Hardware* versions. The *Settings* tab lists the *System Object ID, Date created* (the date this managed device entered the database), *Creator* (the user who discovered or created the device), *Install Date, Management State* (see also the *Manage > Administrative State* menu below) , *Operational State, Topology Icon Size (eXtra Large - Small),* and any *Notes* about the device.

NOTE:

ChangingfieldsintheEditorscreenslikeNetwork Status,AdministrativeState,Operational State(and MAC address for ports) do not change thedevice; they change only the Cruz database. You can alter thesefields totakenotes orsetaspirational values,but nochangegoestothedevice,andresync eradicates changes made if the device has conflicting values.

Resource Editor: Management Interface

This lists the management interfaces for the selected device, including the *IP Address*, *Port*, *Retries*, and *Timeout*.

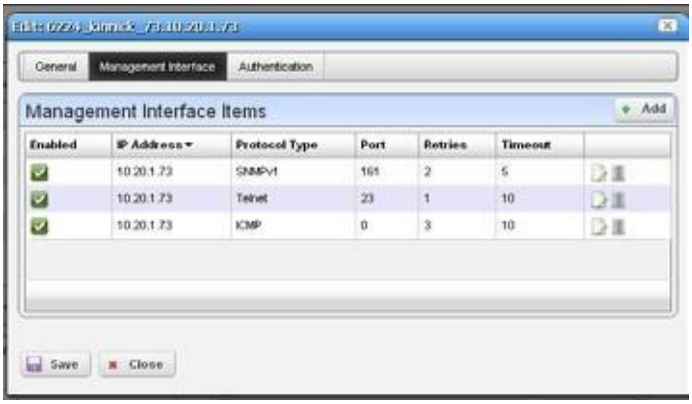

Enabled	IP Address	Protocol Type	Port	Retries	Timeout
☑	10.20.1.73	SNMPv1	161	2	5
☑	10.20.1.73	Telnet	23	1	10
☑	10.20.1.73	ICMP	0	3	10

You can *Add* interfaces with the button in the upper right corner, delete them with the icon to the right of the listed interface.

NOTE:

If an operation produces an error saying the device lacks authentications, if none exists that corresponds to the authentication type, make sure that you add a management interface as well as authentication to remedy that problem.

Resource Editor: Authentication

This lists the authentications for the selected device. You can *Add* authentications with the button in the upper right corner, delete them with the icon to the right of the listed authentication. These originate in the portlet described in Authentications on page 160.

Details—Displays several tabs with detailed resource information. A reminder of the selected device's name appears above the tab bar. See Connected Devices on page 189 for more information about this screen.

NOTE:

Not all devices support the options listed below for the *Manage* menu.

Links

The links portlet displays discovered or created links in your system. If information is truncated, hover the cursor over a column to see the contents of that column as a tooltip. The expanded portlet displays link connections in a Reference Tree snap panel and includes columns for the interface and device for each end of the link.

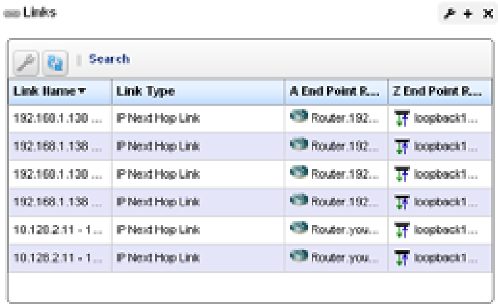

Link Name	Link Type	A End Point R...	Z End Point R...
192.168.1.130 ...	IP Next Hop Link	Router.192...	loopback1...
192.168.1.138 ...	IP Next Hop Link	Router.192...	loopback1...
192.168.1.130 ...	IP Next Hop Link	Router.192...	loopback1...
192.168.1.138 ...	IP Next Hop Link	Router.192...	loopback1...
10.128.2.11 - 1...	IP Next Hop Link	Router.you...	loopback1...
10.128.2.11 - 1...	IP Next Hop Link	Router.you...	loopback1...

The expanded portlet displays link connections in a Reference Treesnap panel and includes columns for the interface and device for each end of thelink.

By right-clicking, you can create a *New* link, *Edit* an existing, selected one, or *Discover* links for specified devices. See New Link and *Link Discovery* on page 188 for more about creating, editing and discovering links.

New Link

When you create a new link or edit an existing one, the *Link Details* screen appears where you can configure the link.

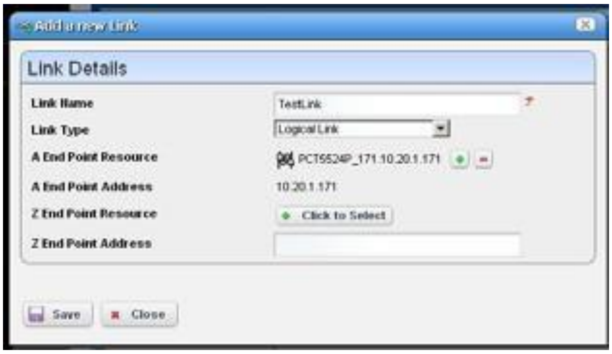

Link Name	TestLink
Link Type	Logical Link
A End Point Resource	PCT5524P_171.10.20.1.171
A End Point Address	10.20.1.171
Z End Point Resource	
Z End Point Address	

This screen has the following fields:

Link Name — A text identifier for the link.

Link Type — Select the type of link from the pick list.

A End Point Resource/Address — Click the plus (+) to select a resource for one end of the link. When you right-click a selected resource, it automatically appears here. Click the minus (-) to remove it.

Z End Point Resource/Address — Click the plus (+) to select a resource for one end of the link. When you have selected two resources, they automatically appear as A and Z endpoints.

Link Discovery

This is an automated network link discovery feature that you can initiate from individual devices in the Resource Editor portlet, or with the *Link Discovery* button on the home screen. The following device features that provide link information:

- IEEE Link Layer Discovery Protocol (LLDP) support
- Cisco Discovery Protocol (CDP) support

Links discovered can also appear in the screen described in Links in a Topology on page 247.

Although Cruz automates link discovery, you must enable the sourcesfor link discovery information on the devices where you do such discovery.

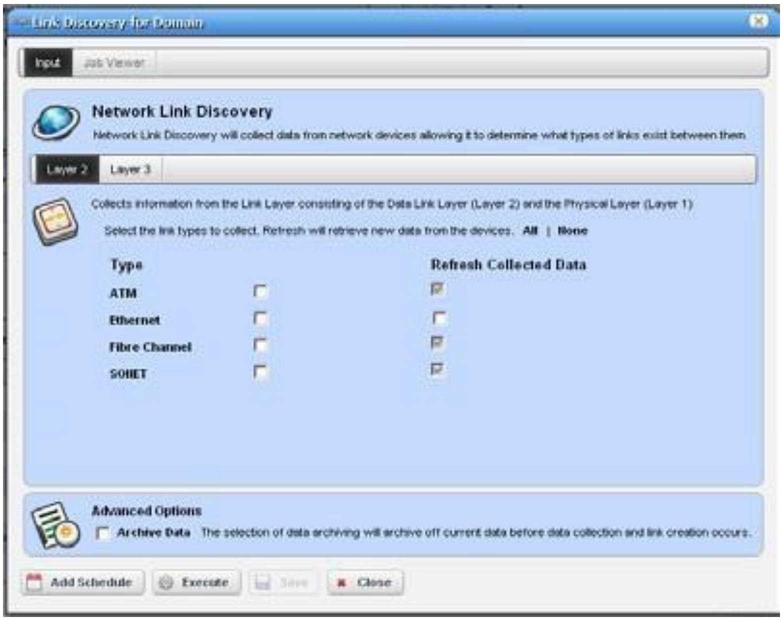

Network Link Discovery

Network Link Discovery will collect data from network devices allowing it to determine what types of links exist between them

Collects information from the Link Layer consisting of the Data Link Layer (Layer 2) and the Physical Layer (Layer 1)

Select the link types to collect. Refresh will retrieve new data from the devices. All | None

Type		Refresh Collected Data
ATM	□	☑
Ethernet	□	□
Fibre Channel	□	☑
SONET	□	☑

Advanced Options

□ Archive Data The selection of data archiving will archive off current data before data collection and link creation occurs.

When you elect to discover links from a right-click menu, the *Network Link Discovery* screen appears. Check the type of links you want to discover or from which you want to refresh collected data. Other options available on this screen include thefollowing:

- **Layer 2/Layer 3** option allows you to select the layer for which you want to discover links. Depending on the layer selected, the available types appear as options below this tab selection.

NOTE:

Click All/None to select all or none of the displayed types for each layer. Remember, selecting more link types consumes more time and processing power.

- **Archive Data** is an advanced option. Selecting this option archives current data before collecting information about and discovering links.

NOTE:

Links with incomplete endpoint information are not discovered.

Click Add Schedule to schedule link discovery or Execute to run it now (and confirm you are willing to wait for results in a subsequent screen). The Job Viewer tab in the link discovery screen displays the message traffic between Cruz and the device(s). See Audit Trail Portlet on page 150 for more about Job Viewer screens.

Search by IP or Mac Address

The Search by IP or Mac Address portlet lets you find Managed Equipment, Ports and Interfaces for the IP or MAC address entered.

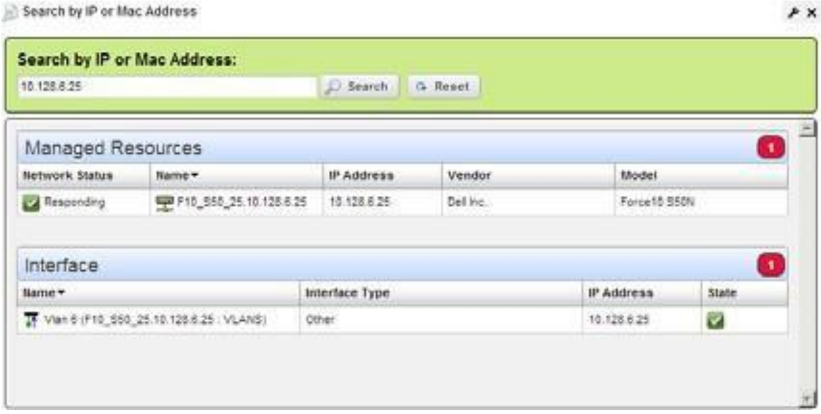

10.128.6.25

Managed Resources 1

Network Status	Name	IP Address	Vendor	Model
✓ Responding	F10_S50_25.10.128.6.25	10.128.6.25	Dell Inc.	Force10 S50N

Interface 1

Name	Interface Type	IP Address	State
Vlan 6 (F10_S50_25.10.128.6.25 : VLAN6)	Other	10.128.6.25	✓

The same right-click menus as appear in Resource Editor, Ports or Interfaces appears in the search results. The display confines those results to what is found; if only ports satisfy the search criteria, then Managed Equipment and Interface do not appear. A count of found items appears in the upper right corner of each panel.

NOTE:

Search by IP or Mac Address portlet supports search by description on both ports and interfaces

Connected Devices

Connected Device(s) is an enhanced version of the Find Physical Connection for IP or MAC Address feature. Connected Device(s) can assist you in locating the port that a device you are searching for is physically connected to. Additionally, using Connected Device(s), you can determine everything that is connected to aspecified managed device - whether those somethings are managed or unmanaged.

Cruz's Connected Device(s) feature works by utilizing LLDP, CDP, EDP, Bridging, and ARP data that is collected during Network Data Collection. At the very end of the Network Data Collection action, using that data, a list of connected devices is compiled and maintained for each device being managed by Cruz.

NOTE:

Connected Device(s) relies on Network Data Collection related data being up to date. It is automatically scheduled to be run daily, however - depending on your network situation - you may wish to schedule Network Data Collection to happen more frequently.

Additionally, when using Connected Device(s), if you have reason to believe that your Network Data Collection data may be out of date, it is suggested that you run Network Data Collection on any devices that you believe may havechanged.

NOTE:

In order for the Connected Device(s) feature to have maximum success at finding devices that may be connected to your managed equipment, it is important to ensure that all devices in your network have all configuration necessary to enable the following applicable protocols/network technologies: LLDP, CDP, EDP, ARP, Bridging Tables. The necessary configuration varies from manufacturer to manufacturer, so please consult with the applicable guides for the devices you are working with.

Access this portlet from the Resources page or the selected connected device's Details portlet.

Searching and Filtering

Data displayed in the Connected Device(s) portlet can be searched and filtered on in two ways.

- The Quick Search feature can be used from the minimized portlet viewand allows the user to search the Remote Device IP, Remote Device ID, Remote Port ID, and VLAN fields.

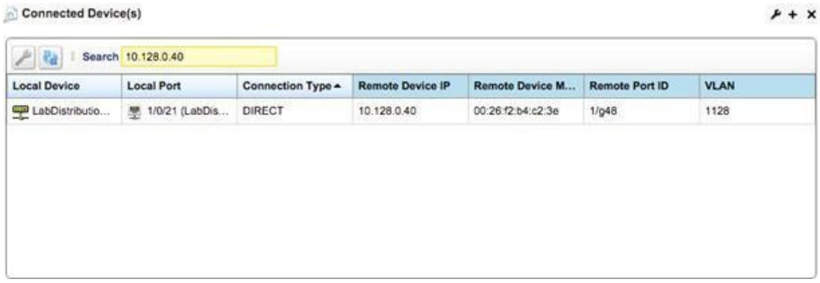

Search: 10.128.0.40

Local Device	Local Port	Connection Type ▲	Remote Device IP	Remote Device M...	Remote Port ID	VLAN
LabDistributio...	1/0/21 (LabDis...	DIRECT	10.128.0.40	00:26:f2:b4:c2:3e	1/g48	1128

- Advanced filtering can be done from the portlet's maximized view and allows any field to be filtered on.

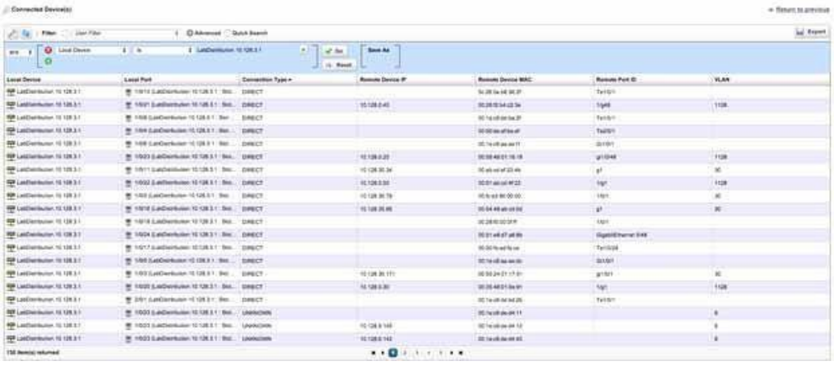

Fields

The following fields are available for Connected Devices:

- Local Device - The endpoint that is being managed by Cruz.
- Local Port - The port that the Local Device knows about the connected devicethrough.
- Connection Type - This will either be DIRECT or UNKNOWN. If the type is DIRECT, the connection was discovered via LLDP, CDP, or EDP, and a direct connection is known to exist. In the case that the type is UNKNOWN, the connection was discovered via bridging data, and the connection may be either indirect, or direct.

- Remote Device IP - The IP Address of the device connected to Local Device.
- Remote Device ID - The Remote ID of the device connected to the Local Device. The value contained within this field can vary depending upon how LLDP, CDP, EDP, etc. is configured on the devices in your network (For example, it could be the MAC address of the remote chassis, or it could be a descriptive string).
- Remote Port ID - The Remote Port ID of the port that the remote device is connecting to the Local Device via.
- VLAN - The VLAN through which the Local Device knows about the Remote Device.

NOTE:

For a given connection, certain fields may or may not be populated, depending on which data (LLDP, CDP, EDP, Bridging, ARP) that connection was present in.

Equipment Details

This screen displays the details for a selected resource. You can see it by selecting *Details* in the right-click menu for the Resource Editor portlet. You can also install an Equipment Details portlet on a page and use the Hierarchical View portlet to select individual devices that appear in it. In that case, you must select an individual device before it displaysdata.

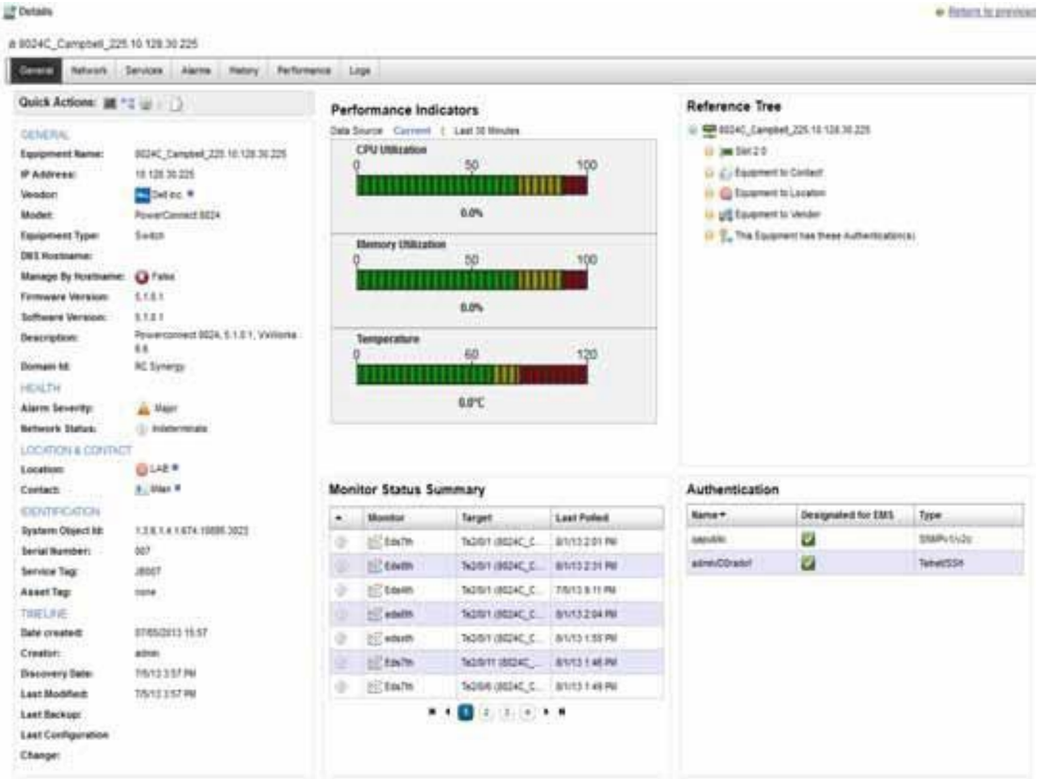

8024C_Campbell_225.10.128.30.225

GENERAL

Equipment Name: 8024C_Campbell_225.10.128.30.225
IP Address: 10.128.30.225
Vendor: Dell inc.
Model: PowerConnect 8024
Equipment Type: Switch
DNS Hostname:
Manage By Hostname: False
Firmware Version: 5.1.0.1
Software Version: 5.1.0.1
Description: Powerconnect 8024, 5.1.0.1, VxWorks 6.6
Domain Id: RC Synergy

HEALTH

Alarm Severity: Major
Network Status: Indeterminate

LOCATION & CONTACT

Location: LAB
Contact: Man

IDENTIFICATION

System Object Id: 1.3.6.1.4.1.674.10895.3023
Serial Number: 007
Service Tag: J8007
Asset Tag: none

TIMELINE

Date created: 07/05/2013 15:57
Creator: admin
Discovery Date: 7/5/13 3:57 PM
Last Modified: 7/5/13 3:57 PM
Last Backup:
Last Configuration Change:

Monitor Status Summary

▲	Monitor	Target	Last Polled
	Eds7th	Te2/0/1 (8024C_C...	8/1/13 2:01 PM
	Eds8th	Te2/0/1 (8024C_C...	8/1/13 2:31 PM
	Eds4th	Te2/0/1 (8024C_C...	7/8/13 8:11 PM
	eds8th	Te2/0/1 (8024C_C...	8/1/13 2:04 PM
	edsxth	Te2/0/1 (8024C_C...	8/1/13 1:55 PM
	Eds7th	Te2/0/11 (8024C_...	8/1/13 1:46 PM
	Eds7th	Te2/0/6 (8024C_C...	8/1/13 1:49 PM

Authentication

Name ▾	Designated for EMS	Type
public	☑	SNMPv1/v2c
admin/DOradof	☑	Telnet/SSH

Detailsscreensareavailableforavarietyofthingsbesidesequipment, too. Herearesome highlights of the Equipment Details screen (and others):

The *Quick Actions* panel in the General tab also displays icons that activate direct access or the resource editor.

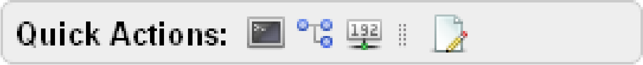

Direct accessincludes Terminal, MIBBrowser, Ping (ICMP) or HTTP/HTTPS).

Click the tab name to see the following:

General–In addition to Quick Action icons, this displays details about the selected equipment, including its Domain ID. This screen also includes performance indicators to report on the device's CPU, memory and disk utilization (flash memory) both currently and for the last 30 minutes (click the links above the panel), a Monitor Status Summary, and Reference Tree, and a list of the Authentications connected to the device. If disk utilization is less than one percent, an indication that the device is still active may appear in thatgraph.

Permissionscontrol thevisibilityofsome attributes. Information like IOS/Firmware/Software versionsappears to usersonlyif they have READ permission for these attributes. To configure attributes as not viewable—one example: management IP address—you must define the attribute set and add it to `owareapps\installprops\lib\installed.properties` as in the following example:

```
restricted.attribute.names.set1=RedCell.Config.EquipmentManager_IPAddress
```

Foreach restricted attribute set, you can define multiple attributes with comma delimited DSI names.

After defining such attributes, in the permission manager, uncheck the READ permission on roles that you do not wish to have access to `Restricted Attribute Access 1`. Youcan conceal up to five such attributes. Attribute names are similar to the Using Email Action Variables on page 310, at least when fully qualified with the prefix `Redcell.Config.EquipmentManager`. Contact technical support for identifiers not listed there.

Network–This screen lists the Portsand Interfaces for the selected device (some devices have one, but not the other), VLANs and links associated with the device.

Alarms – Displays the alarms and events associated with the selected device.

History–Includes audit trails connected to the device, and any backed up configurations. Right-click to view or otherwise act on these.

Performance – This screen contains two links at its top. One displays a performance dashboard (template) related to the selected device. See Show Performance Templates on page 425 for how to configure these. The other displays any configured *Top Talkers* for the device. See Top N [Assets] on page413.

You can also configure the interval displayed by clicking the clock icon at the top of this screen.

Logs – Displays maintenance logs connected to the device to users with permissions to see this tab. Right-click to create or edit these.

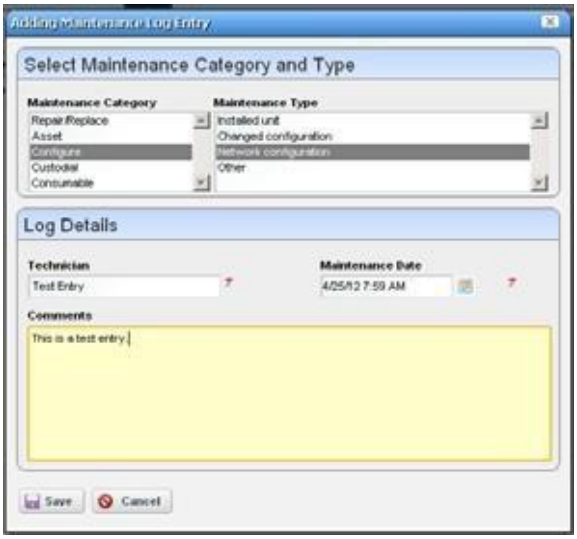

Adding Maintenance Log Entry

Select Maintenance Category and Type

Maintenance Category: Repair/Replace, Asset, Configure (selected), Custodial, Consumable

Maintenance Type: Installed unit, Changed configuration, Network configuration (selected), Other

Log Details

Technician: Test Entry

Maintenance Date: 4/25/12 7:59 AM

Comments: This is a test entry.

Notice that you can right-click listed interfaces, configuration files, and so on to perform more actions, or to see additional Detail screens.

You can also right-click and then select *Details* to view more information about some subcomponents like Interfaces and Ports. The Details portlet includes a *Reference Tree* (similar to the Widget). You can even right-click nodes in that reference tree to drill down to additional details.

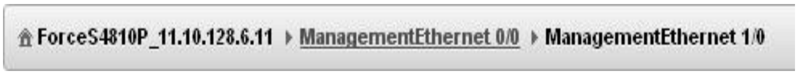

ForceS4810P_11.10.128.6.11 ▸ ManagementEthernet 0/0 ▸ ManagementEthernet 1/0

NOTE:

Notice the breadcrumb trail at the top of the Equipment Detail panel tracks the levels through which you drill down. You can click a level that appears in this trail to return to a previous screen. If you click *Return to previous* in the upper right corner of the screen, you will return to the original screen from which you selected the basicequipment.

Also: Some fields may appear truncated on screen, but you can hover the cursor over the truncated field, so the text appears as a tooltip or drill down to see the detail.

Some devices populate the ports panel, but not the interfaces panel, which is empty for such devices. Interfaces may appear for Dell EMC Networking FTOS, Cisco or Juniper devices. You may also discover such devices as type: Unknown (see Base Driver on page 107). Force 10 devices interfaces details can display Port Channels (LAGs), VLANs (SVIs) and Loopbacks.

If the Ports portlet is on the same page as the Resource Editor portlet, selecting a device in Managed Resources makes its ports appear in the Port s portlet. The display can also get out of sync but clicking the browser’s *Refresh* restores the correspondence between a selected device and the ports displayed. To resync a port, resync the device that contains it.

Field Definitions

The meanings of most fields that appear in details screens are self-evident. Here is a little more information about some of them:

Operational State — One of following possible values describing the availability of the resource.

Disabled —Inoperable because of a fault, or resources are unavailable.

Enabled — Operable and available for use.

Active — Device is operable and currently in use with operating capacity available to support further services.

Busy — Operable and currently in use with no operating capacity to spare.

Administrative State — One of the following values:

Locked — Device use is prohibited.

Shutting Down — Only existing users can use the device.

Unlocked — Normal use of device is permitted.

Network Status — The status of the resource in the network. For example: *Responding* means this application can, via some network protocol, get the device to respond. *Not Responding* means the device does not respond to the protocol. *Indeterminate* means the monitoring software has not tried to reach the device or there was some other error which prevented us from determining one of the other two statuses.

The appearance of *Network Status* depends on the default ICMP monitor (see Resource Monitors on page 361. If you exclude this equipment from the monitor or disable it (for example, for performance reasons) then a status may appear, but it is not meaningful.

You can now use monitors other than ICMP to determine this status. See Monitoring Network Availability on page 365.

NOTE:

The Alarms Details panel now lets you correlate parent/child alarm pairs.

Direct Access

Direct access provides less mediated access to the device in the following ways:

- MIB Browser
- Terminal
- Ping (ICMP)
- HTTP/HTTPS

The following sections describe those direct options in more detail.

NOTE:

Best practice is to avoid special characters, particularly # and > (command line prompts) in device banners so terminal access is unambiguous.

MIB Browser

As part of the *Direct Access* menu, the *MIB Browser* lets you examine SNMP data available about devices.

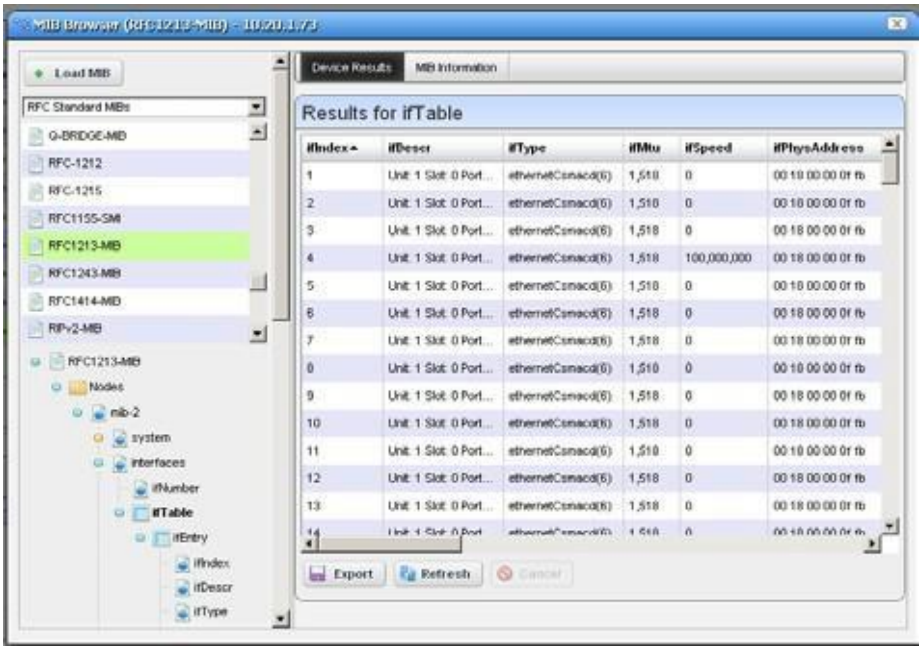

Results for ifTable

ifIndex ▲	ifDescr	ifType	ifMtu	ifSpeed	ifPhysAddress
1	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
2	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
3	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
4	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	100,000,000	00 18 00 00 0f fb
5	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
6	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
7	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
8	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
9	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
10	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
11	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
12	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
13	Unit: 1 Slot: 0 Port...	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb
14	Unit: 1 Slot: 0 Port	ethernetCsmacd(6)	1,518	0	00 18 00 00 0f fb

The screen that opens when you select this option displays MIBs available in Cruz in a tree on the left. Notice that a pick list at the top of the left column narrows what appears in the tree. A progress bar at the bottom of this screen indicates a query for the selected information is in progress.

Click *Load MIB* at the top left corner of the screen to load a new MIB. A file selection dialog opens after you click *Load MIB*. Click the *Refresh* button at the bottom of the browser to re-query the device for new information. Click the *Export* button at the bottom of the browser to export the screen contents to a spreadsheet (Excel-format) file.

Use the *Load MIB* button in the upper right corner, or the menu described in Using Extended Event Definitions on page 333 for loading new MIBs.

Select a MIB and expand it to see the contents for a selected node appear on the right. In addition to the *Device Results* tab, which displays what the currently selected device uses from the MIB, the *MIB Information* tab displays the parameters available for the selectednode.

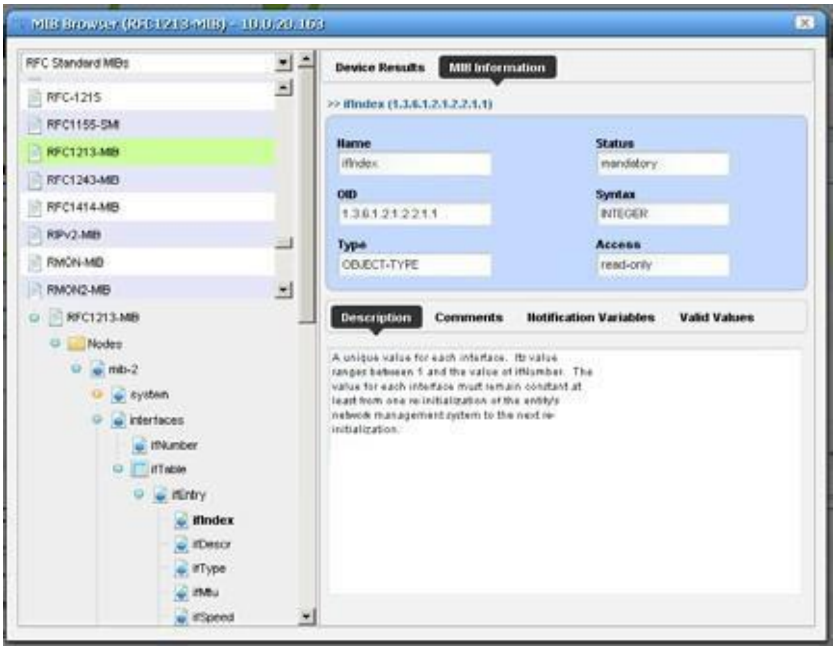

>> ifIndex (1.3.6.1.2.1.2.2.1.1)

Name: ifIndex
Status: mandatory
OID: 1.3.6.1.2.1.2.2.1.1
Syntax: INTEGER
Type: OBJECT-TYPE
Access: read-only

A unique value for each interface. Its value ranges between 1 and the value of ifNumber. The value for each interface must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization.

Noticethat the Description, Comm ents, Notification Variables, and Valid Valuestabsappearat the bottom of this screen.

Terminal

This opens a terminal shell connected to the selected device.

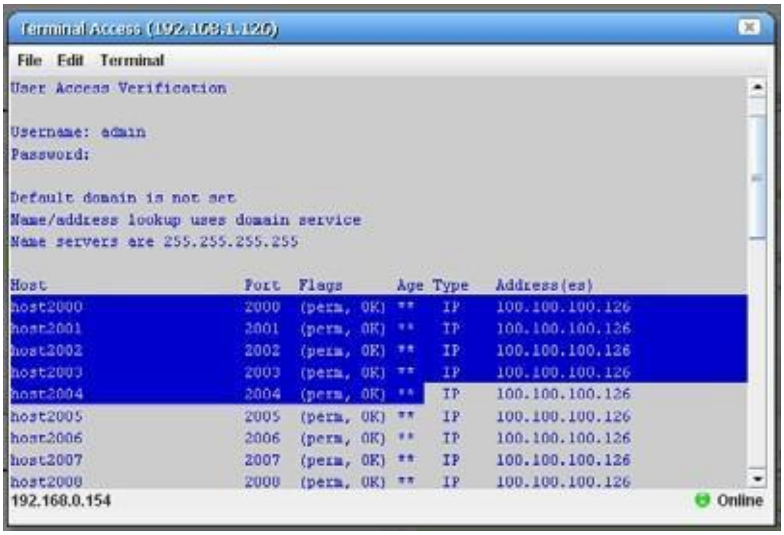

```
User Access Verification

Username: admin
Password:

Default domain is not set
Name/address lookup uses domain service
Name servers are 255.255.255.255

Host                      Port  Flags      Age Type   Address(es)
host2000                  2000  (perm, OK) **   IP    100.100.100.126
host2001                  2001  (perm, OK) **   IP    100.100.100.126
host2002                  2002  (perm, OK) **   IP    100.100.100.126
host2003                  2003  (perm, OK) **   IP    100.100.100.126
host2004                  2004  (perm, OK) **   IP    100.100.100.126
host2005                  2005  (perm, OK) **   IP    100.100.100.126
host2006                  2006  (perm, OK) **   IP    100.100.100.126
host2007                  2007  (perm, OK) **   IP    100.100.100.126
host2008                  2008  (perm, OK) **   IP    100.100.100.126
192.168.0.154
```

A green icon in the lower right corner indicates the device is online, while the IP address of the device appears in title bar. The IP address of Cruz’s server also appears in the lower left corner, when the connection is active.

The following menus appear for your terminal session:

File — This menu lets you *Connect* or *Disconnect* to the device.

Edit—This menu lets you *Copy* or *Paste* text within the terminal session. Click and drag to select text.

Terminal — This menu lets you set *Foreground* and *Background* colors, as well as configuring the *Font* and *Buffer* sizes. *Reset Terminal* restores the defaults.

Terminalisan applet that requiresa Java Runtime Environment beinstalled and associated to the browser as a plug-in on the client machine. See also Setting Java Security for Terminal Access on page 214 below.

NOTE:

You can cut and paste from the Direct Access terminal.

Telnet sessions are synchronous. You cannot interrupt a command in progress with another command you send, unless you haveenabled something that periodicallyprompts foradditional commands (for example enabling line continuation prompts).

Ping (ICMP)

Select this option from the Direct Access menu to initiate ICMP ping, and to display a list of the selected device's ping responses.

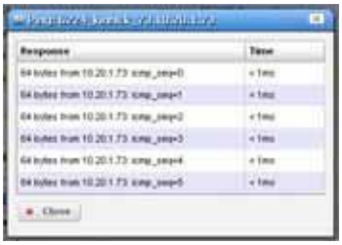

Response	Time
64 bytes from 10.20.1.73: icmp_seq=0	<1ms
64 bytes from 10.20.1.73: icmp_seq=1	<1ms
64 bytes from 10.20.1.73: icmp_seq=2	<1ms
64 bytes from 10.20.1.73: icmp_seq=3	<1ms
64 bytes from 10.20.1.73: icmp_seq=4	<1ms
64 bytes from 10.20.1.73: icmp_seq=5	<1ms

Alternatively, an error message can appear describing the device's lack of response.

When ping responds in less than one millisecond, results appear in a table with <1ms entries.

HTTP/HTTPS

Selecting this menu item opens the default browser, providing access to the selected device.

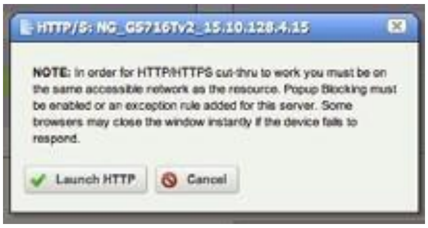

HTTP/S: NG_GS716Tv2_15.10.128.4.15

NOTE: In order for HTTP/HTTPS cut-thru to work you must be on the same accessible network as the resource. Popup Blocking must be enabled or an exception rule added for this server. Some browsers may close the window instantly if the device fails to respond.

An intervening dialog appearsadvising you about the required network conditions for a successful connection.

Secure Connections to Cruz

Typically, multi-server installations with load balancers are where Cruz users need secure connections. Consult the *Installation Guide* for instructions about how to configure HTTPS to connect to a load balancer in such an installation.

Ports

The Ports summary portlet displays discovered device ports. Note that in this context a "Port" is a typeofsubcomponent that is essentially a physical interface of the device. The Interfaces portlet is tied to a different distinct but related type of subcomponent that some device vendors might call a "sub-interface" or a "logical interface".

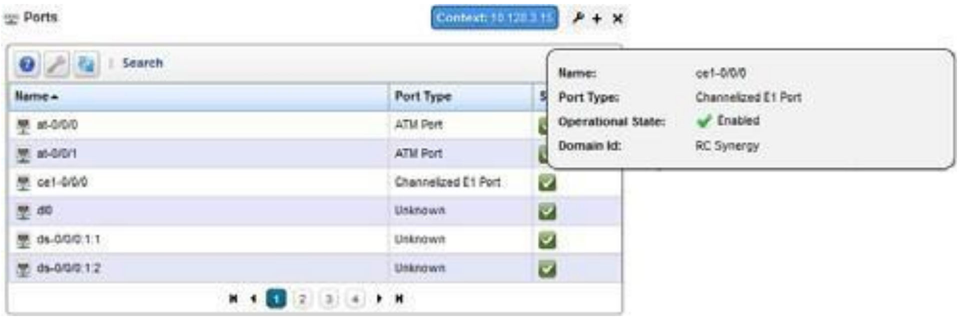

Name ▴	Port Type	S
at-0/0/0	ATM Port	✓
at-0/0/1	ATM Port	✓
ce1-0/0/0	Channelized E1 Port	✓
dl0	Unknown	✓
ds-0/0/0:1:1	Unknown	✓
ds-0/0/0:1:2	Unknown	✓

Name: ce1-0/0/0
Port Type: Channelized E1 Port
Operational State: ✓ Enabled
Domain Id: RC Synergy

This displays a list of ports, with columns for *Port Icon*, *Equipment Name*, *Name*, *Type* and *Encapsulation.* Hover the cursor over the *State* column, and a popup appears to display the port's *Name*, *Type* and *Operational State* information. Right clicking offers a subset of the actions listed in Resource Editor on page 184. You can also create links. See Links on page 187. See Port Editor on page 202 for details of the editor produced by right clicking to select *Edit* specifically for ports.

If the Ports summary portlet appears on the same page as the Managed Resources portlet, then a selection made in Managed Resources makes the Ports portlet display only ports for the selected resource. This"filter" through Managed Resources disablesfilters configured through the settings menu. See Understanding Hierarchical View on page 248 for more about thisfeature.

The Port Details portlets display all the port's settings that have been retrieved, including a Reference Treeoflogical interfaces below the port, a Learned MAC Address panel, Alarms related to the port, and other Detail.

NOTE:

Switch Mode is only populated for devices that support the Bridge-MIBs. The default value is unknown.

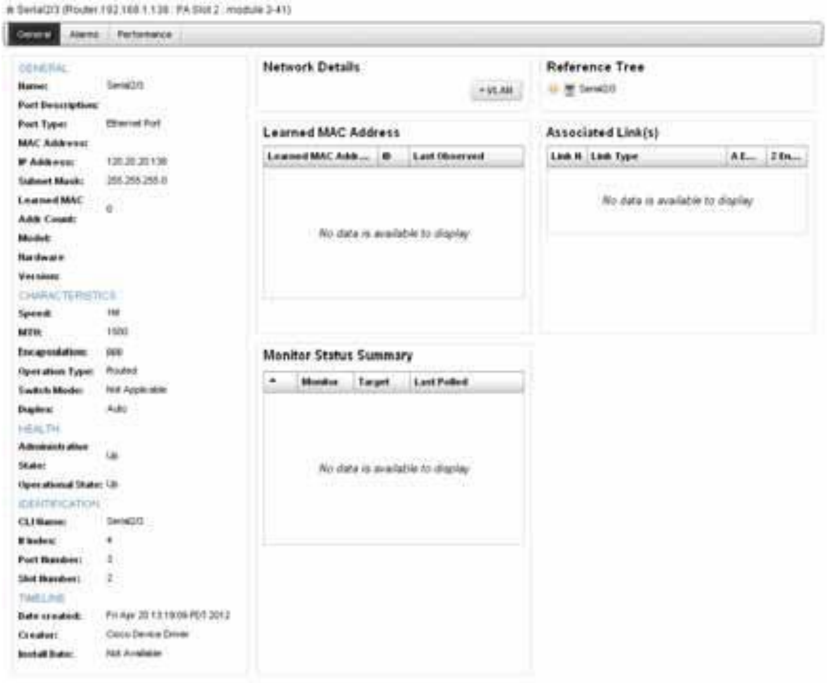

Serial2/3 (Router 192.168.1.138 : PA Slot 2 : module 3-41)

GENERAL

Name: Serial2/3

Port Description:

Port Type: Ethernet Port

MAC Address:

IP Address: 120.20.20.138

Subnet Mask: 255.255.255.0

Learned MAC Addr Count: 0

Model:

Hardware:

Versions:

CHARACTERISTICS

Speed: 1M

MTU: 1500

Encapsulation: ppp

Operation Type: Routed

Switch Mode: Not Applicable

Duplex: Auto

HEALTH

Administrative State: Up

Operational State: Up

IDENTIFICATION

CLI Name: Serial2/3

If Index: 4

Port Number: 3

Slot Number: 2

TIMELINE

Date created: Fri Apr 20 13:19:09 PDT 2012

Creator: Cisco Device Driver

Install Date: Not Available

Network Details

Reference Tree

Serial2/3

Learned MAC Address

Learned MAC Addr...	ID	Last Observed
No data is available to display		

Associated Link(s)

Link #	Link Type	A E...	Z En...
No data is available to display			

Monitor Status Summary

▲	Monitor	Target	Last Polled
No data is available to display			

NOTE:

Switch Mode is only populated for devices that support the Bridge-MIBs. The default value is unknown.

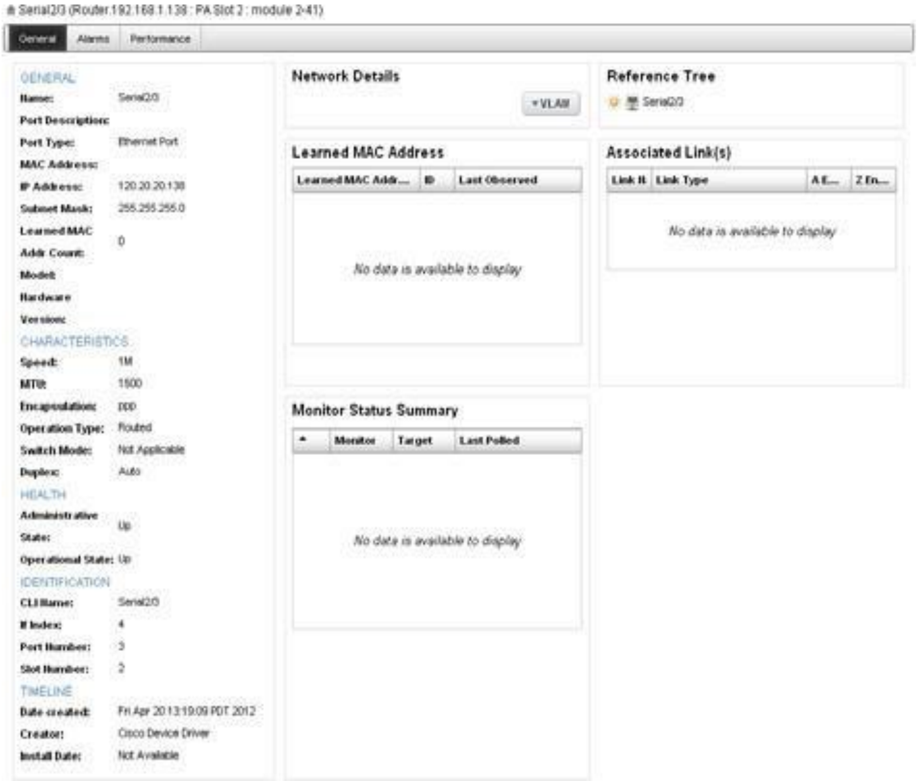

Serial2/3 (Router.192.168.1.138 : PA Slot 2 : module 2-41)

GENERAL

Name: Serial2/3

Port Description:

Port Type: Ethernet Port

MAC Address:

IP Address: 120.20.20.138

Subnet Mask: 255.255.255.0

Learned MAC Addr Count: 0

Model:

Hardware

Version:

CHARACTERISTICS

Speed: 1M

MTU: 1500

Encapsulation: ppp

Operation Type: Routed

Switch Mode: Not Applicable

Duplex: Auto

HEALTH

Administrative State: Up

Operational State: Up

IDENTIFICATION

CLI Name: Serial2/3

If Index: 4

Port Number: 3

Slot Number: 2

TIMELINE

Date created: Fri Apr 20 13:19:09 PDT 2012

Creator: Cisco Device Driver

Install Date: Not Available

Network Details

Reference Tree

Serial2/3

Learned MAC Address

Learned MAC Addr...	ID	Last Observed

No data is available to display

Associated Link(s)

Link H	Link Type	A E...	Z En...

No data is available to display

Monitor Status Summary

▲	Monitor	Target	Last Polled

No data is available to display

This screen displays the following tabs, accessed by clicking their name in the top of the screen. Just above their names, a reminder appears of the name of the selected port.

General — In this tab, fields appear describing attributes for the selected port. For example *Date Created* (typically, this is the date discovered).

Alarms — This tab displays alarms and the Event History connected to the selected port. See Alarms Portlet on page 276 and Event History on page 286 for more about that information.

Performance — Displays monitor information, if available, related to the selected port.

See also Connected Devices on page 189 and Managed Resources on page 177 for an explanation of some of these fields.

Ports Expanded

Clicking the plus (+) in the upper right corner of the summary screen displays this expanded view of available ports.

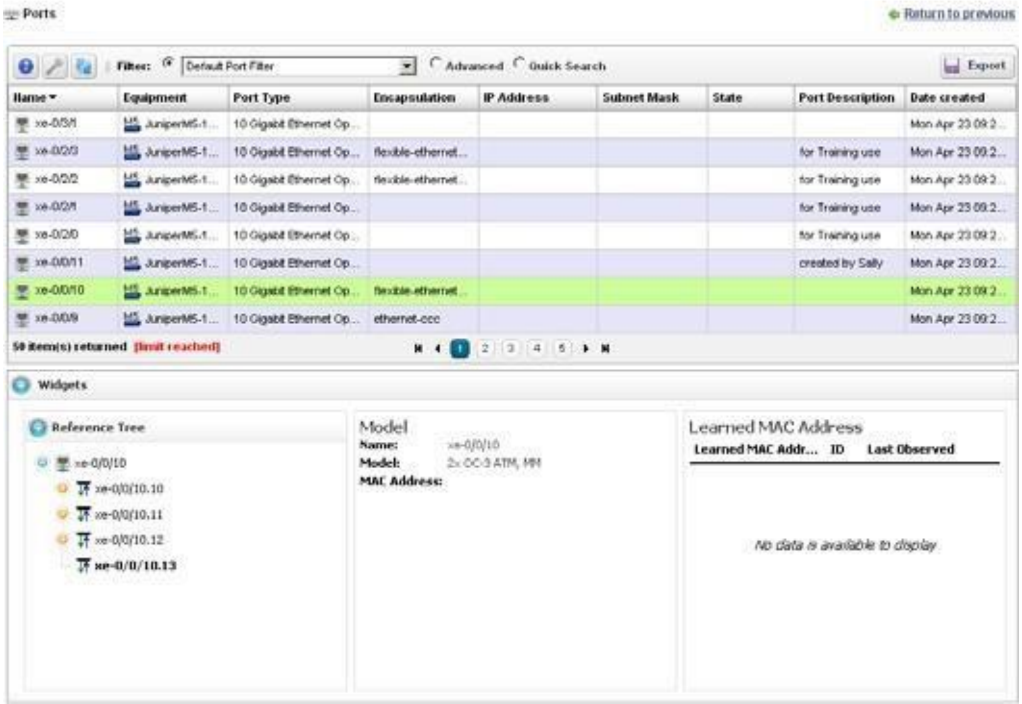

Filter: ● Default Port Filter ○ Advanced ○ Quick Search

Name	Equipment	Port Type	Encapsulation	IP Address	Subnet Mask	State	Port Description	Date created
xe-0/3/1	JuniperM5-1...	10 Gigabit Ethernet Op...						Mon Apr 23 09:2...
xe-0/2/3	JuniperM5-1...	10 Gigabit Ethernet Op...	flexible-ethernet...				for Training use	Mon Apr 23 09:2...
xe-0/2/2	JuniperM5-1...	10 Gigabit Ethernet Op...	flexible-ethernet...				for Training use	Mon Apr 23 09:2...
xe-0/2/1	JuniperM5-1...	10 Gigabit Ethernet Op...					for Training use	Mon Apr 23 09:2...
xe-0/2/0	JuniperM5-1...	10 Gigabit Ethernet Op...					for Training use	Mon Apr 23 09:2...
xe-0/0/11	JuniperM5-1...	10 Gigabit Ethernet Op...					created by Sally	Mon Apr 23 09:2...
xe-0/0/10	JuniperM5-1...	10 Gigabit Ethernet Op...	flexible-ethernet...					Mon Apr 23 09:2...
xe-0/0/9	JuniperM5-1...	10 Gigabit Ethernet Op...	ethernet-ccc					Mon Apr 23 09:2...

50 item(s) returned [limit reached] 1 2 3 4 5

Reference Tree

- xe-0/0/10
  - xe-0/0/10.10
  - xe-0/0/10.11
  - xe-0/0/10.12
  - **xe-0/0/10.13**

Model
Name: xe-0/0/10
Model: 2x OC-3 ATM, MM
MAC Address:

Learned MAC Address

Learned MAC Addr...	ID	Last Observed

No data is available to display

The Settings button lets you configure columns that appear and their order. The available columns for this view include many related to the attributes that appear in the Port Details. This screen also includes a Reference Tree displaying a tree of the selected port’s relationship to logical interfaces and monitors.

This screen has columns like those described in Alarms Portlet on page 276. Configure these as visible or hidden by clicking *Settings*. The following are some additional columns available.

Port Editor

When you right-click a port and select *Edit* this screen appears.

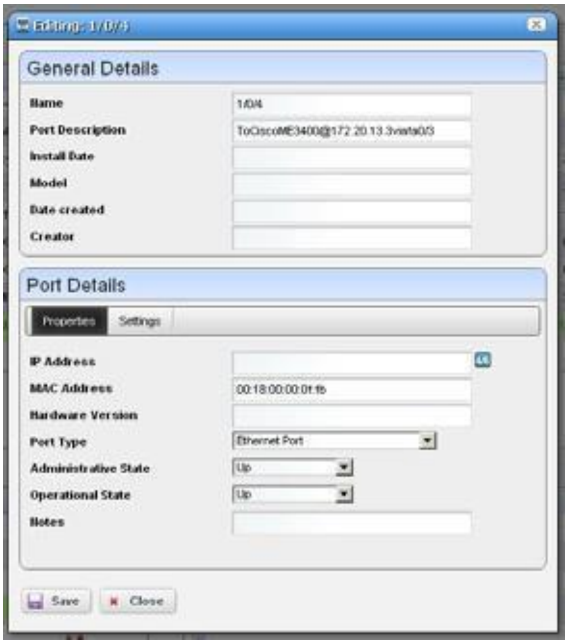

Editing: 1/0/4

General Details

Name	1/0/4
Port Description	ToCiscoME3400@172.20.13.3vista0/3
Install Date	
Model	
Date created	
Creator	

Port Details

IP Address	
MAC Address	00:18:00:00:01:fb
Hardware Version	
Port Type	Ethernet Port
Administrative State	Up
Operational State	Up
Notes	

It has the following fields:

General Details

Name — An identifier for the port.

Port Description — A text description for the port.

Install Date — The date this port was installed.

Model — The port’s model.

Date created — The date this port was created.

Port Details - Properties

IP Address — The IP address for the port.

MAC Address — The port’s Media Access Control (MAC) address.

Hardware Version — The port’s hardware version.

Port Type — The type of port.

Administrative State — One of three descriptive values. The options are:

Locked — Device use is prohibited.

Shutting Down — Only existing users can use the device.

Unlocked — Normal use of device is permitted.

Operational State — One of the following values:

Down — Inoperable because of a fault, or resources are unavailable.

Dormant — The port is dormant.

Not Present — The port is absent.

Up — Operable and available for use.

Unknown —Status is unknown.

Testing —Status is testing.

Notes — Any notes recorded about the device.

Port Details - Settings

Encapsulation — An identifier for the port.

MTU — The size of the maximum transmission unit.

Speed — The port's speed.

Subnet Mask — Any subnet mask associated with the port.

In Use — Checked if the port is in use.

IF Index — The port's SNMP If Index number.

NOTE:

Changing fields in the Editor screens like Network Status, Administrative State, Operational State and MAC address do not change the device; they change only the Cruz database. You can alter these fields to take notes or set aspirational values, but no change goes to the device, and resync eradicates changes made if the device has conflicting values.

Interfaces

Use the Interfaces portlet to monitor and manage your discovered interface resources. This portlet, like the Ports portlet, displays subcomponents of discovered resources. Unlike the Ports portlet, it does not display the Widgets panel in its expanded form, just more columns.

Access this portlet from the Resources page. This portlet has both a summary view and an expanded view. Each view could display different columns and could have the same pop-up menu options available.

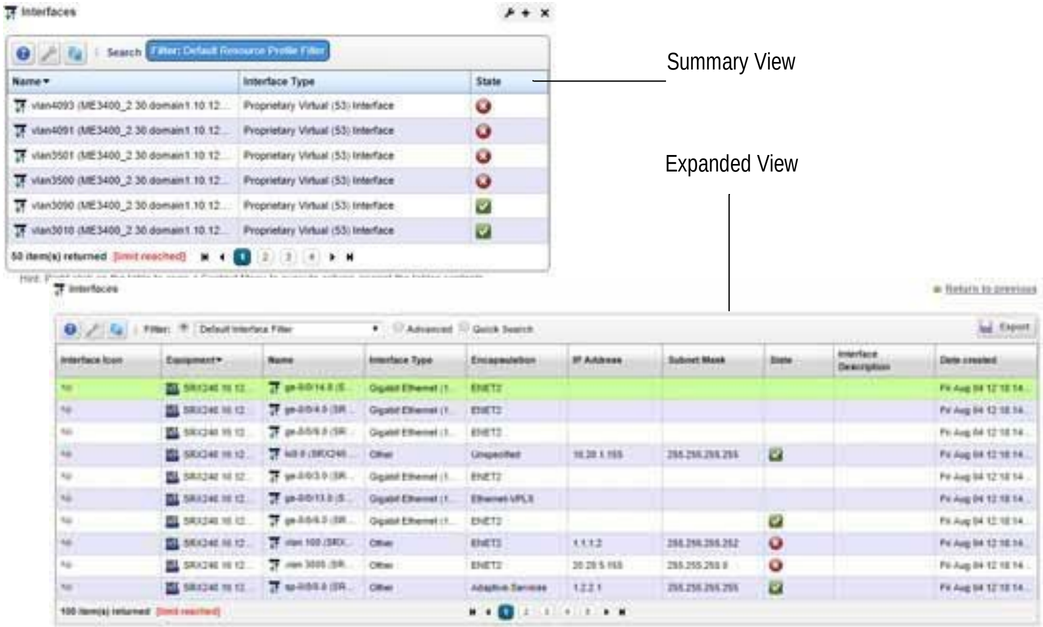

Name	Interface Type	State
vlan4093 (ME3400_2.30.domain1.10.12...	Proprietary Virtual (53) Interface	✗
vlan4091 (ME3400_2.30.domain1.10.12...	Proprietary Virtual (53) Interface	✗
vlan3501 (ME3400_2.30.domain1.10.12...	Proprietary Virtual (53) Interface	✗
vlan3500 (ME3400_2.30.domain1.10.12...	Proprietary Virtual (53) Interface	✗
vlan3090 (ME3400_2.30.domain1.10.12...	Proprietary Virtual (53) Interface	✓
vlan3010 (ME3400_2.30.domain1.10.12...	Proprietary Virtual (53) Interface	✓

50 item(s) returned [limit reached] 1 2 3 4

Interface Icon	Equipment	Name	Interface Type	Encapsulation	IP Address	Subnet Mask	State	Interface Description	Date created
N/A	SRX240 10.12...	ge-0/0/14.0 (S...	Gigabit Ethernet (1...	ENET2					Fri Aug 04 12:18:14...
N/A	SRX240 10.12...	ge-0/0/4.0 (SR...	Gigabit Ethernet (1...	ENET2					Fri Aug 04 12:18:14...
N/A	SRX240 10.12...	ge-0/0/9.0 (SR...	Gigabit Ethernet (1...	ENET2					Fri Aug 04 12:18:14...
N/A	SRX240 10.12...	lo0.0 (SRX240...	Other	Unspecified	10.20.1.155	255.255.255.255	✓		Fri Aug 04 12:18:14...
N/A	SRX240 10.12...	ge-0/0/3.0 (SR...	Gigabit Ethernet (1...	ENET2					Fri Aug 04 12:18:14...
N/A	SRX240 10.12...	ge-0/0/13.0 (S...	Gigabit Ethernet (1...	Ethernet-VPLS					Fri Aug 04 12:18:14...
N/A	SRX240 10.12...	ge-0/0/5.0 (SR...	Gigabit Ethernet (1...	ENET2			✓		Fri Aug 04 12:18:14...
N/A	SRX240 10.12...	vlan.100 (SRX...	Other	ENET2	1.1.1.2	255.255.255.252	✗		Fri Aug 04 12:18:14...
N/A	SRX240 10.12...	vlan.3005 (SR...	Other	ENET2	20.20.5.155	255.255.255.0	✗		Fri Aug 04 12:18:14...
N/A	SRX240 10.12...	sp-0/0/0.0 (SR...	Other	Adaptive Services	1.2.2.1	255.255.255.255	✓		Fri Aug 04 12:18:14...

100 item(s) returned [limit reached] 1 2 3 4 5

Columns

Other than the generalnavigationandconfiguration options, the Interfaces portlets (summaryand expanded views) include the following columns. The columns displayed by default are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Name	An interface identifier, such as the CLI Name. This field displays on the summary and expanded views by default.
Interface Type	The resource interface type, such as IP interface, Proprietary Virtual interface, Primary ISN interface, and so on. This field displays on the summary and expanded views by default.
State	An indicator that shows whether the operation state is functioning (checkmark) or not (X). This field displays on the summary and expanded views by default.
Equipment	The name of the managed resource (top-level device) in which the interface is contained. This field displays on the expanded view by default.
Encapsulation	An identifier for the port. This field displays on the expanded view by default.
IP Address	The interface’s **primary** address within the network. See Additional IP Addresses if you interface has more than one address. This field displays on the expanded view by default.

Column	Description
Subnet Mask	The network segment address to which an interface belongs. This field displays on the expanded view by default.
Interface Description	A detailed description for the interface. This field displays on the expanded view by default.
Date Created	The date this resource entered the database. This field displays on the expanded view by default.
Additional IP Addresses	Any additional IP addresses assigned to the interface, there than the primary IP address. For example, an interface might have an IPv4 address and an IPv6 address. If this is the case, the primary IP address appears in the IP Address column and the additional IP addresses appear in this column.
Administrative State	The interface's current administrative state, such as down, testing, unknown, or up.
Alarm Severity	The highest severity of any existing alarms against this resource.
Alarm Suppression Description	The description provided for an alarm suppression configuration.
Alarm Suppression Mode	Indicates if the device is in alarm suppression, and if so, which mode of suppression is active. The two modes of active alarm suppression are "Ad Hoc", which means that suppression will remain indefinitely until it is stopped manually, and "Scheduled", which means that suppression will automatically stop at some point in the future (see the Schedules portlet). If the device is not in alarm suppression, then a dash "-" is listed. If some form of alarm suppression is listed, the device does not receive alarms.
CLI Name	The command line resource name.
Hierarchical View Index	The hierarchical view index number.
Creator	The name of the person or process (such as a discovery rule) that created this interface.
Domain ID	The identifier for the interface domain.
Egress Bandwidth	The interface's outgoing bandwidth (bps).
Egress/Ingress Bandwidth Calculation	The type of calculation used to determine bandwidth: • CONFIGURED means that the Cruz system has a QoS policy configured against the interface. • TRUNK AGGREGATION means that the port is in trunking mode and calculates it bandwidth from its access ports. • INTERFACE AGGREGATION means that the total bandwidth of the parent is the sum of the bandwidth of its children if the interface has sub-interfaces. • ASSOCIATION means that the bandwidth is taken from the linked port if the port **is currently** UNCONFIGURED and has a physical and has a physical link to another port that **is not** UNCONFIGURED. • UNCONFIGURED means that the bandwidth is set to the port's IfSpeed. This is the default setting. See Bandwidth Calculation on page 410 for more details.
Foreign ID	An optional identifier for the customer's organization.
IfIndex	The SNMP interface index number used to identify a particular end point (port or interface).
Ingress Bandwidth	The interface's incoming bandwidth (bps).
Install Date	The time and date the interface was installed.
Interface Number	The number used to create a logical or virtual interface.

Column	Description
MTU	The maximum transmission unit (packet) size (256-9192).
Notes	A text field for information about the interface.
Port Number	A number assigned to user sessions and server applications in an IP network.
Rack Number	The number that identifies on which rack the component (card, port, or interface) is installed. This is useful when you have larger chassis that have internal racks.
VRF Name	An identifier for the VRF.
Slot Number	An interface attribute set by the Cruz product.
Speed	The port/interface speed.
Subslot Number	The number that identifies which slot on the card that the service module is installed. Cards have the potential of supporting child modules, configurable service modules installed on the card.

Pop-Up Menu

The Interfaces pop-up menu provides access to th e following options. Right-click a row to access these options.

Menu Option	Description
Edit	Opens the Editing *interfaceName* window, where you view or modify the resources general details or interface properties.
Details (Click +Shift)	Displays the selected resource's general details and the alarms, event history, performance, and quality of service (QOS) details.
Topology	Displays a topology map that includes the selected resources. See Presentation Capabilities on page 223 for more about these maps.
Domain Access Control	Opens the Domain Access Control window where you assign access rights to the domain.
Actions	Executes any actions associated with the selected resource. An error is displayed if you do not have permission to execute actions or if there are no actions for the selected resource.
Event Management	Suppresses or updates alarms related to the selected resource. The following event management options available are: • **Start Alarm Suppression** starts alarm suppression after you enter a description and then click Start. A Success/Failure confirmation message is displayed. • **Stop All Alarm Suppression** discontinues all alarm suppression. A confirmation message is displayed. You can also stop alarm suppression using the View Active Suppressions option. • **Schedule Alarm Suppression** opens the New Schedule *suppressionAction* window, where you set the exception options and suppression targets parameters and then set the schedule options. • **View Active Suppressions** displays a list of active suppressions. A message is displayed if there are no active suppressions. You can stop selected alarm suppressions from this list. • **View Event History** shows event history for the selected resource in the Event History portlet. • **Resync Alarms** refreshes the information displayed with the latest information in the database for the selected resource or all resources. Except for View Event History, perform these actions on one or more interfaces.

Menu Option	Description
Links	Opens the Add a New Link window when you select New. Specify a link name, type, and A/Z end points resource/address. See Links on page 187 for more details about the link options available.
Show Performance	Shows performance data for the selected interfaces in the Performance Dashboard window. Click the edit tool to modify the dashboard view properties, entities, or attributes. See Dashboard Editor on page 420 for more about the editor and its options.
Traffic Flow Analyzer	Executes the following traffic flow analyzer actions on the selected resource: • Show Traffic • Create Configuration • Remove Configuration • Show Configuration See Traffic Flow Analyzer on page 515 for more details.
Edit Custom Attributes	Opens the Custom Attribute Editor where you define field characteristics, such as whether it is enabled, the label name, and the tooltip.
View as PDF	Creates an Acrobat PDF document containing this alarm’s contents displayed in the summary portlet.
Share with User	Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Cards

Use the Cards portlet to monitor and manage your discovered card resources.

This portlet is intended for users who are interested in managing their discovered devices.

After installation, this portlet is not from any page. However, you can add it to any page, such as the Resources page by selecting Add > Application, expanding the Resource Management list, and adding the Cards portlet. This portlet has both a summary view and an expanded view. Each view could display different columns and could have the same pop-up menu options available.

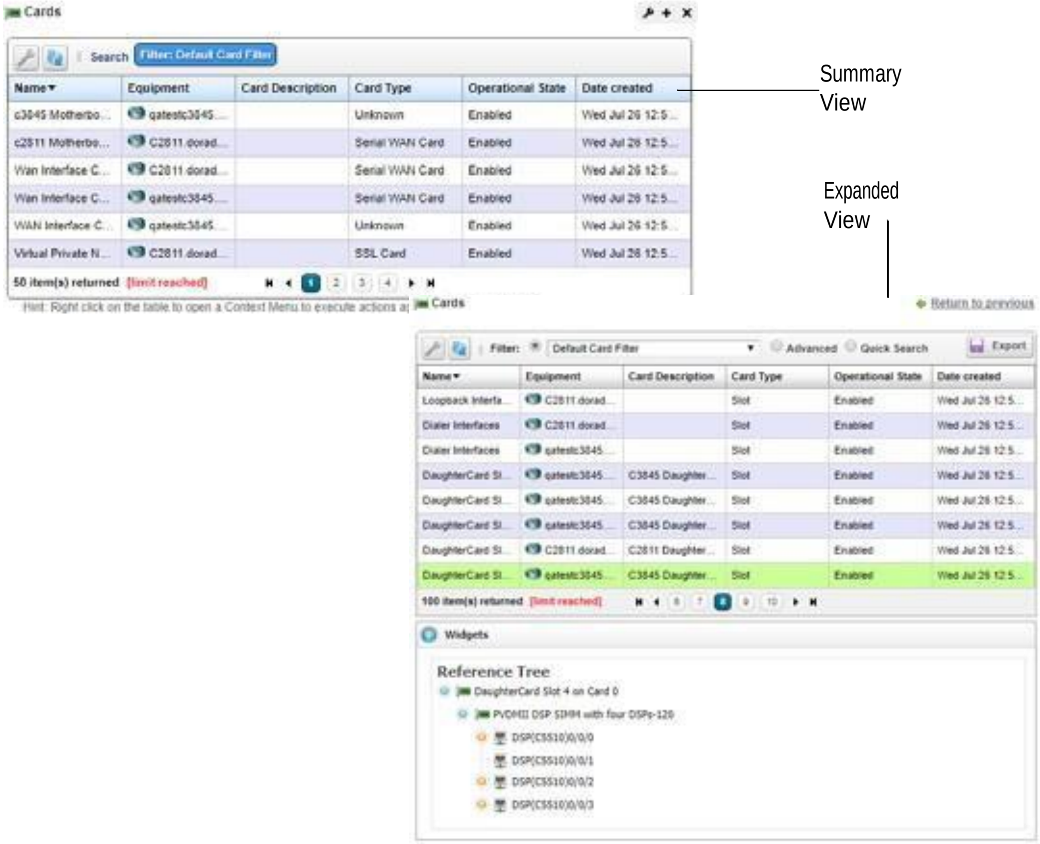

Name	Equipment	Card Description	Card Type	Operational State	Date created
c3845 Motherbo...	gatestc3845...		Unknown	Enabled	Wed Jul 26 12:5...
c2811 Motherbo...	C2811.dorad...		Serial WAN Card	Enabled	Wed Jul 26 12:5...
Wan Interface C...	C2811.dorad...		Serial WAN Card	Enabled	Wed Jul 26 12:5...
Wan Interface C...	gatestc3845...		Serial WAN Card	Enabled	Wed Jul 26 12:5...
WAN Interface C...	gatestc3845...		Unknown	Enabled	Wed Jul 26 12:5...
Virtual Private N...	C2811.dorad...		SSL Card	Enabled	Wed Jul 26 12:5...

50 item(s) returned [limit reached]

Hint: Right click on the table to open a Context Menu to execute actions a

Name	Equipment	Card Description	Card Type	Operational State	Date created
Loopback Interfa...	C2811.dorad...		Slot	Enabled	Wed Jul 26 12:5...
Dialer Interfaces	C2811.dorad...		Slot	Enabled	Wed Jul 26 12:5...
Dialer Interfaces	gatestc3845...		Slot	Enabled	Wed Jul 26 12:5...
DaughterCard Sl...	gatestc3845...	C3845 Daughter...	Slot	Enabled	Wed Jul 26 12:5...
DaughterCard Sl...	gatestc3845...	C3845 Daughter...	Slot	Enabled	Wed Jul 26 12:5...
DaughterCard Sl...	gatestc3845...	C3845 Daughter...	Slot	Enabled	Wed Jul 26 12:5...
DaughterCard Sl...	C2811.dorad...	C2811 Daughter...	Slot	Enabled	Wed Jul 26 12:5...
DaughterCard Sl...	gatestc3845...	C3845 Daughter...	Slot	Enabled	Wed Jul 26 12:5...

100 item(s) returned [limit reached]

Reference Tree

- DaughterCard Slot 4 on Card 0
  - PVDMII DSP SIMM with four DSPs-120
    - DSP(C5510)0/0/0
    - DSP(C5510)0/0/1
    - DSP(C5510)0/0/2
    - DSP(C5510)0/0/3

Columns

Other than the general navigation and configuration options, the Cards portlets (summary and expanded views) include the following columns. The columns displayed by default are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Name	A resource identifier, such as the CLI Name. This field displays on the summary and expanded views by default.
Equipment	The name of the resource, such as ISDN card, NIC, and so on. This field displays on the summary and expanded views by default.
Card Description	A description of the resource, such as the card part number. This field displays on the summary and expanded views by default.
Card Type	The type associated with this resource, such as slot, ISDN card, NIC, Other, and so on. This field displays on the summary and expanded views by default.

Column	Description
Operational State	An indicator that shows whether the resource is: • **Active** shows that the card is operable and currently in use with operating capacity available to support further services. • **Busy** shows that the card is currently in use with no operating capacity to spare. • **Disabled** shows that the card is not operable because of a fault or unavailable resources. • **Enabled** shows that the card is operational and available for use. • **Not Determined** show that the state is not determined. This field displays on the summary and expanded views by default.
Date Created	The date this resource entered the database. This field displays on the summary and expanded views by default.
Active Ports	The number of active ports related to a resource. By default, a port is considered active if the administrative state is up. You can modify this functionality to use the port's operational state to determine whether a port is active or inactive by adding the following property to your installed.properties file, saving the file, and then restarting the application server: `com.dorado.redcell.inventory.ActivePortAttribute=OW_OperationalState` This attribute only updates when the device re-syncs.
Administrative State	The interface's current administrative state, such as down, testing, unknown, or up.
Alarm Severity	The highest severity of any existing alarms against this resource.
Alarm Suppression Description	The description provided for an alarm suppression configuration.
Alarm Suppression Mode	A boolean setting (on/off). If Ad Hoc suppression is listed, the device does not receive alarms.
CLI Name	The command line resource name.
Hierarchical View Index	The Hierarchical View index number.
Creator	The name of the person or process (such as, a discovery rule or device driver) that created this interface.
Domain ID	The identifier for the resource domain.
Firmware Version	The resource's firmware version.
Foreign ID	An optional identifier for the customer's organization.
Hardware Version	The card's hardware version.
Card Icon	The icon that represents the resource.
Inactive Ports	The number of inactive ports related to a resource. By default, a port is considered inactive if the administrative state is **not** up. You can modify this functionality to use the port's operational state to determine whether a port is active or inactive by adding the following property to your installed.properties file, saving the file, and then restarting the application server: `com.dorado.redcell.inventory.ActivePortAttribute=OW_OperationalState` This attribute only updates when the device re-syncs.
Install Date	The time and date the resource was installed.

Column	Description
Model	The number that identifies the resource model.
Notes	A text field for information about the resource.
Physical Index	The given index associated to the resource. This is useful when a vital device component, such as a power supply or fan has sensors attached.
Rack Number	The number that identifies on which rack the component (card, port, or interface) is installed. This is useful when you have larger chassis that have internal racks.
Serial Number	The resource's serial number.
Slot Number	An attribute set by the Cruz product.
Software Version	The resource's software version.
Subslot Number	The number that identifies which slot on the card that the service module is installed. Cards have the potential of supporting child modules, configurable service modules installed on the card.
Total Ports	The total number of ports related to a resource. This attribute only updates when the device re-syncs.

Pop-Up Menu

The Cards pop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
Edit	Opens the Editing *cardName* window, where you view or modify the general details and card properties.
Details (Click +Shift)	Shows the general, alarms, and event history details for the selected card.
Edit Custom Attributes	Opens the Custom Attribute Editor window, where you add or modify attribute definitions.
Share with User	Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Discovering Resources

This section outlines some basic resource discovery steps. Discover resources as follows.

1 Set up Authentications for the resource you want to discover.

a. Click Settings from the navigation bar.

b. Right-click the Authentications portlet and then select New.

c. Enter a unique identification.

d. Add any associated equipment.

Cruz must be authorized to set CLI session parameters; permissions-related timeouts may occur during device accessif it is not enabled. Forexample, Cisco CLI accessrequires the `set terminal length 0` command.

2 Set up the Discovery Profiles for th e resource you want to discover.

a. Select Resources > Discovery from the navigations bar.

b. Right-click the Discovery Profiles portlet and then select New.

c. Enter a unique name for the profile.

d. Specify the discovery, network, actions, and inspection options.

e. Execute the profile.

If the discovery fails, use the Network Tools Portlet to confirm that any device missing from discovery is online. The device itself must permit the Cruz access it.

3 View the results in the Managed Resources portlet.

Modifying Discovery Profiles

Modify discovery profiles from the Discovery Profiles portlet as follows.

1 Right-click the profile you want to modify.

2 Select Edit.

The Discovery Profile editor is displayed.

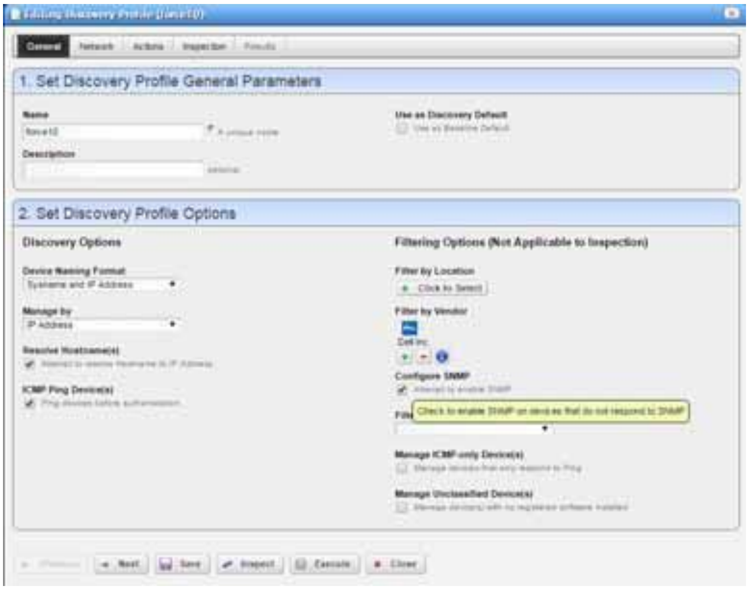

1. Set Discovery Profile General Parameters

Name: force10 * A unique name

Description: (optional)

Use as Discovery Default: □ Use as Baseline Default

2. Set Discovery Profile Options

Discovery Options

Device Naming Format: Sysname and IP Address

Manage by: IP Address

Resolve Hostname(s): ☑ Attempt to resolve Hostname to IP Address

ICMP Ping Device(s): ☑ Ping devices before authentication

Filtering Options (Not Applicable to Inspection)

Filter by Location: Click to Select

Filter by Vendor: Dell Inc.

Configure SNMP: ☑ Attempt to enable SNMP

Check to enable SNMP on devices that do not respond to SNMP

Manage ICMP-only Device(s): □ Manage devices that only respond to Ping

Manage Unclassified Device(s): □ Manage device(s) with no registered software installed

3 Modify the general parameters and profile options as needed.

4 Click the Network tab and then modify network types or authentication.

The Network panel is displayed.

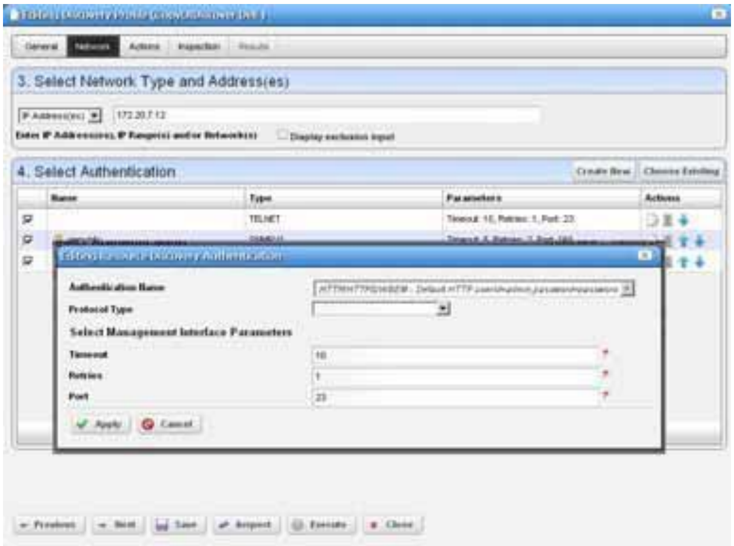

3. Select Network Type and Address(es)

IP Address(es): 172.20.7.12

Enter IP Address(es), IP Range(s) and/or Network(s) □ Display exclusion input

4. Select Authentication

	Name	Type	Parameters	Actions
☑		TELNET	Timeout: 10, Retries: 1, Port: 23	
☑		SNMPV1	Timeout: 4, Retries: 2, Port: 161	
☑				

Editing Resource Discovery Authentication

Authentication Name: HTTP/HTTPS/WSDM - Default HTTP

Protocol Type:

Select Management Interface Parameters

Timeout: 10

Retries: 1

Port: 23

5 Click the Actions tab and then modify actions as needed.

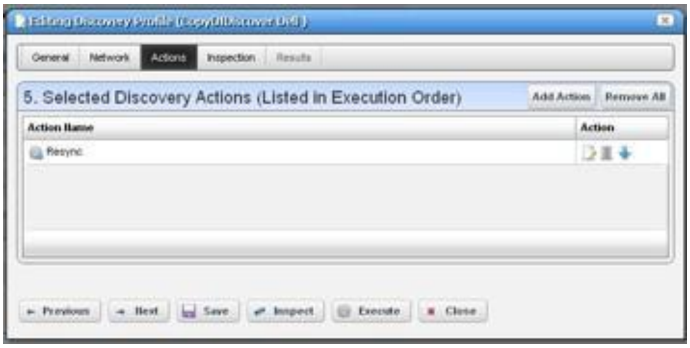

5. Selected Discovery Actions (Listed in Execution Order)

Action Name	Action
Resync	

6 Click the Inspection tab and then modify settings as needed.

7 Click Save to preserve the profile.

If you execute the profile before saving it, the Cruz system does not save the profile to execute later.

8 Click Execute.

The Results panel displays the message traffic between the Cruz system and the device. When processing completes, Success or Failed is displayed. The portal also displays a successful completion, unable to execute, or failed message.

9 Close the Discovery Profile editor.

Logging Terminal Sessions

Log terminal sessions (if you like) as follows.

1 Enable Java Console on the client.

- For Windows, do this in Control Panel.
- On Linux you must navigate to the install location of your JRE and run the Console script. Select the Advanced tab and change the Java Console setting to show the console.

  Java Control Panel > General tab > Settings displays the location where the logs are stored.

2 Open a *Direct Access > Terminal* session by right clicking a device.

The Java Console appears.

3 Configure the logging level in the *Terminal* menu of the direct access screen.

Levels, in increasing order of detail, include *None, Info, Debug,* and *Trace,* which echoes keystrokes.

Setting Java Security for Terminal Access

Newer Java distributions (7+) block websites with self-signed certificates, which interferes with Direct Access. You need to provide a security exception for the application server to avoid issues.

Set Java security for terminal access as follows.

1 Access the Java control panel by doing one of the following:
- From Windows, locate the `configure java` program and open it.
- From Linux, locate the JRE installed and associated to your browser and then run the `../jre/bin/ControlPanel` script.

2 Select the *Security* tab and then click *Edit Site List.*

3 Click Add.

4 Enter the Cruz URL (example: http://192.168.0.51:8080/).

5 Click OK and Continue.

6 Apply this change and/or click OK.

Direct access functions correctly after you make this adjustment.

Changing Customer to Port Associations

Youcan associate ordis-associate a customer with a portorinterface using the Associate Contact to Equipment or Remove Associated Contacts actions.

Change customer port associations as follows.

1 Execute the action (Associate Contact to Equipment or Remove Associated Contacts).

2 Select the port andcontact.

3 Finishexecuting the action.

Using the Add Ports Action

A Manage-type Action lets you add Juniper device ports to the database before resyncing picks them up. These can be ports created by Adaptive CLI actions that simply add ports without updating the database, or those added outside the Cruz application. Once ports are in the database, you can then use them in services in in the Service Center. The advantage of having this action available explicitly is that you do not have to wait for a resync to populate the database correctly.

△ CAUTION:
This Management Action works only for ports already added to the device. Using it to add to the Cruz inventory first then running automated or manual resync removes the port from the database.

Use the Add ports action from the Managed Resources portlet as follows.

1 Right-click a device and then select Actions.

The Actions Selectionwindow is displayed.

2 Click the Manage tab.

3 Select Add Ports to Inventory > Continue.

The Action editor displays the input parameters.

4 Enter the CLI names of one or more ports to add (separate multiple entries with commas).

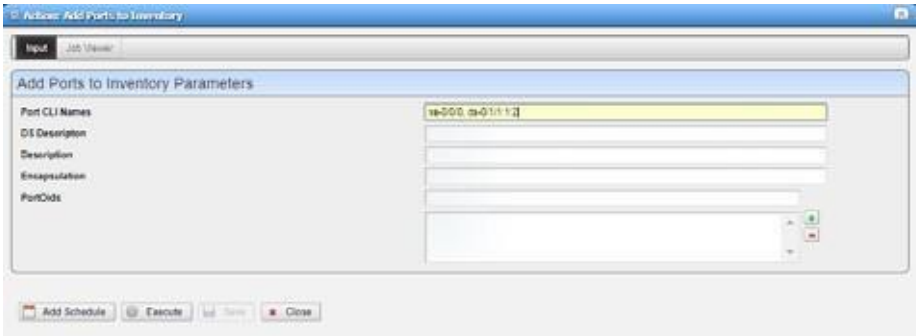

Action: Add Ports to Inventory
Input
Job Viewer
Add Ports to Inventory Parameters
Add Schedule
Execute
Save
Close

Port CLI Names: se-0/0/0, ds-0/1/1:1:2
DS Description
Description
Encapsulation
PortOids

5 Execute or schedule the Action.

If you enter something not in a format like a port (xx- #/ #/ #, or xx- #/#/#: #, and so on), the Cruz application does not attempt to add it. If you enter a port identifier for which the Cruz application cannot find, the Cruz application does not add the port. For example, the parent card (ge-1/10/20) when no card with a containerIndex 1, slotNumber 20 is in the database.

The Cruz application does not populate any of the details other than name, CLI name, slot number, and icon. The Cruz application retrieves any other attributes when a normal resync occurs.

Contacts, Locations and Vendors Portlets

This section describes the following key portlets:

- Contacts Portlet
- Contacts Editor
- Locations Portlet
- Vendors Portlet

Contacts Portlet

The Contacts portlet displays available contacts for your system. There is no expanded version of this portlet, but you can select multiple contacts (Ctrl +click).

Bydefault, this portlet is available byselecting Settings > Groups & Locations fromthe navigation bar.

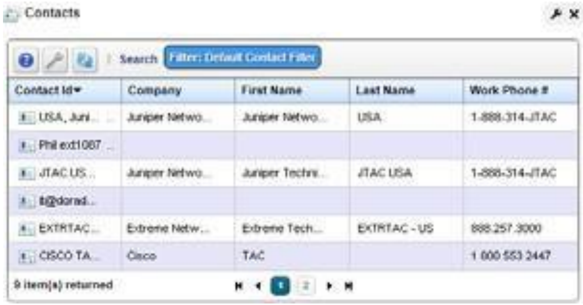

Contact Id	Company	First Name	Last Name	Work Phone #
USA, Juni...	Juniper Netwo...	Juniper Netwo...	USA	1-888-314-JTAC
Phil ext1087 ...				
JTAC US ...	Juniper Netwo...	Juniper Techni...	JTAC USA	1-888-314-JTAC
tl@dorad...				
EXTRTAC...	Extreme Netw...	Extreme Tech...	EXTRTAC - US	888.257.3000
CISCO TA...	Cisco	TAC		1 800 553 2447

9 item(s) returned

You can right-click to act on the selected contact with the following menu items.

New/Edit — Opens the Contacts Editor , where you can create new contacts or alter existing ones.

Details—Displaysa screen with contact-associated alarms, and the information entered in contacts editor .

Delete— Delete the selected item. Caution: such deletions can impact anything else referring to what you are deleting.

Because of its simplicity, this portlet does not have an expanded version, so no plus (+) appears in its upper right corner.

Cruz only retrieves Contact and Location information on initial discovery. Youcan modify these once the resource is under management, however doing so will not modifyany system information on the device.

NOTE:

Youcanimport contacts to Multitenant domainswith acommandlineimporter. Thecommandis `importcontacts`. Thisscriptisinthe `owareapps/redcell/bin`directory. Ittakesthe importfilenameasaparameter. Thedomainmustexistbeforeitcanimportcontacts, anditgeneratesan errorifthecontactspecifiesnodomain, orifthedomaindoesnotexist. Therequireddomainsshould existin the Cruz before import occurs. Example XML files (with the <customer> tag for domains) are in owareapps\redcell\db.

Contacts Editor

This editor has several panels where you can enter contact information (*Name, Address, Phone,* and so on). Click the tabs at the top of this screen to move between the panels. The *Contact ID*, a unique identifier for the contact in your system, is a required field at the top of the first page.

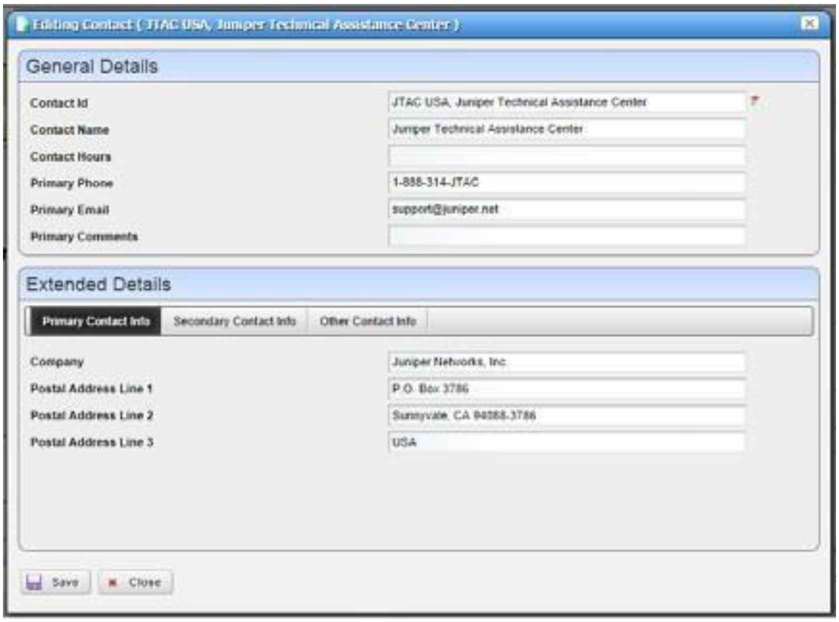
Editing Contact (JTAC USA, Juniper Technical Assistance Center)

General Details

Contact Id	JTAC USA, Juniper Technical Assistance Center
Contact Name	Juniper Technical Assistance Center
Contact Hours	
Primary Phone	1-888-314-JTAC
Primary Email	support@juniper.net
Primary Comments	

Extended Details

Company	Juniper Networks, Inc.
Postal Address Line 1	P.O. Box 3786
Postal Address Line 2	Sunnyvale, CA 94088-3786
Postal Address Line 3	USA

Click *Save* to preserve your newor modified contact information. Click *Cancel* to leave the contact unmodified.

Locations Portlet

The Locations portlet displays configured locations in your system.

By default, this portlet is available from the Settings page by selecting the Groups & Locations menu option.

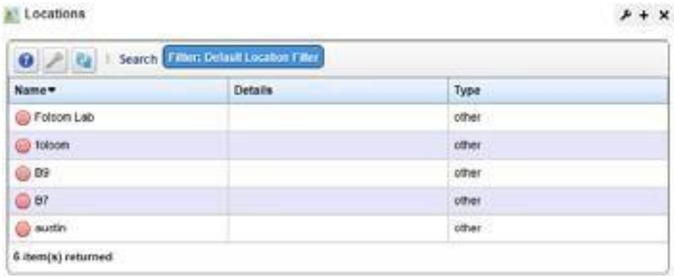

Name	Details	Type
Folsom Lab		other
folsom		other
B9		other
B7		other
austin		other

6 item(s) returned

You can right-click to create, modify or remove (*New, Open, Delete)*the selected location. See Location Editor description below for more about editing or creating locations.

If you select Topology, a map of the selected lo cation's connection to equipment appears. See Presentation Capabilities for more.

This screen has the following columns:

[Icon] — The icon for this location.

Name — The name for this location.

Details —A description for this location.

Type—A designated type for the location.

In addition to the Location port let, this section describes the:

- Location Editor
- Expanded Location Portlet
- Tag

Location Editor

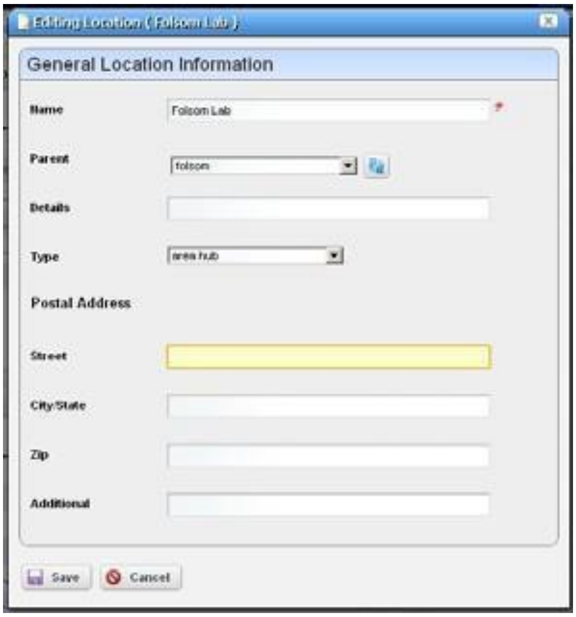

Editing Location (Folsom Lab)

General Location Information

Name: Folsom Lab *

Parent: folsom

Details:

Type: area hub

Postal Address

Street:

City/State:

Zip:

Additional:

When you click *New* or *Open*, an editor appears. The *Name* field is mandatory.

Name—A unique name for the Location. If you alter the name of an existing location already in use by existing equipment, the editor creates a new location. If you change the name of a location, this change may take a short period to percolate to all managed objects that use it. You can do this, though.

Parent —The “parent” of this location (the location to which this location is subordinate). Select a Parent Location from the pick list. The maximum number of levels supported is 15.

Details — A text description of the location.

Type—Typeoflocation, as selected from the drop-down menu. Available types are: Area Hub, Customer, National Hub, Other, Provider, Regional Hub, and State.

Postal Address — The *Street*, *City/State*, *Zip* address of the location.

Additional — Any optional notes.

Click *Save to* save the Location, or any modifications you have made.

△ CAUTION:
Deleting locations can impact anything else referring to what you are deleting.

Expanded Location Portlet

The location portlet displays a list of all locations, with Snap Panels to display a selected location’s connection to the network and details.

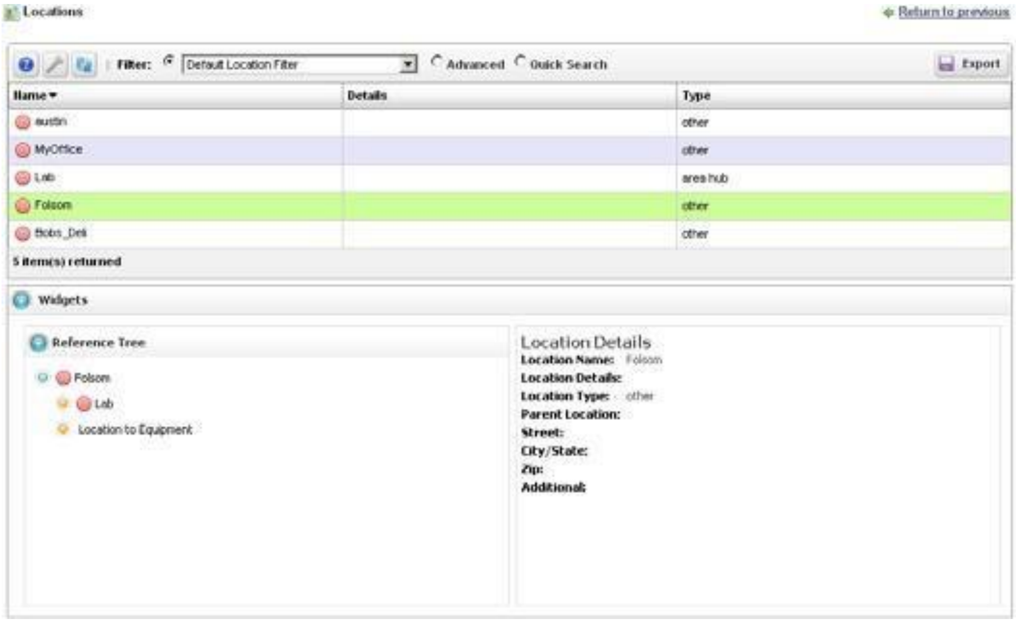

Name	Details	Type
austin		other
MyOffice		other
Lab		area hub
Folsom		other
Bobs_Deli		other

5 item(s) returned

Location Details

Location Name: Folsom

Location Details:

Location Type: other

Parent Location:

Street:

City/State:

Zip:

Additional:

The *New* menu option appears in the expanded location portlet. Click *Settings* to change the column appearance (see Modifying Column Settings on page 144). This has the same columns as Locations Portlet on page 217.

Selecting a location rowdisplays the *Reference Tree* widget, with the selected location’sconnection to hierarchical views (see Hierarchical View on page 228*)* and equipment. Click the plus (+) icons to expand the tree. The *Location Details* panel displays what has been configured in the Location Editor.

Tag

When creating a location, Cruz automatically selects the latitude and longitude for the address entered for a location. Tag a location by right clicking it in the Locations Portlet ..

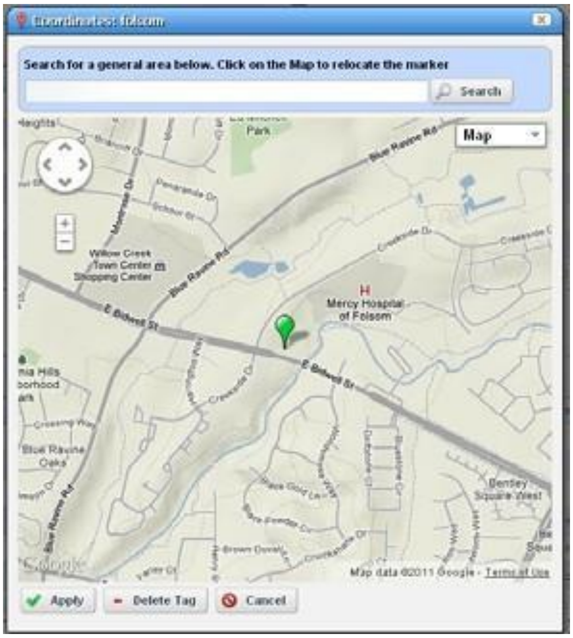

Coordinates: folsom

Search for a general area below. Click on the Map to relocate the marker

The location created by default is the address entered in the Locations Portlet editor. You can also enter the address in the Search field or click and drag the marker that appears on this screen. Click *Apply* to accept the re-location. A *Delete Tag* button appears when you have created a tag, and lets you remove it. *Cancel* closes the screen.

NOTE:

You can zoom in or out on the displayed map with the + and - buttons in the upper left corner of this screen.

Vendors Portlet

The Vendors portlet displays the available vendors for network resources. This portlet also has an Expanded Vendor Portlet .

By default, this portlet is available from the Settings page by selecting the Groups & Locations menu option.

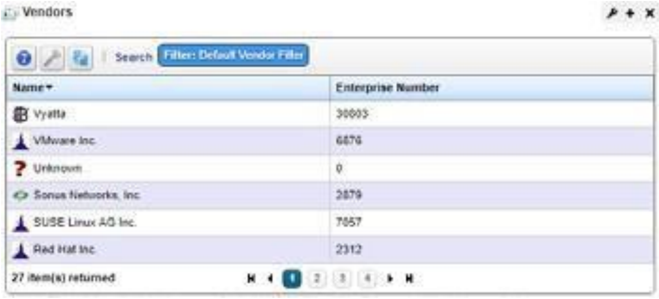

Vendors

Name	Enterprise Number
Vyatta	30803
VMware Inc.	6876
Unknown	0
Sonus Networks, Inc	2879
SUSE Linux AG Inc.	7057
Red Hat Inc	2312

27 item(s) returned

Right clicking a row lets you do the following:

New/Edit— Opens the Vendor Editor where you can configure or re-configure a vendor.

Details—Displays a panel showing the alarms, registered models, and identifiers for the selected vendor.

Topology—See a topology of the network filtered to display only the selected vendor, see Presentation Capabilities

Import/Export—Export the selected config file to disk or import it fromdisk. Youcan also import/ export a selected configuration file.

Provides thefollowingactions when availablefor theselectedimage:

- Import retrieves a file containing XML compliance descriptions. Some imports can come from a URL.
- Export Selection exports the selected description to an XML file.
- Export All exports all descriptions to an XML file.

Click Download Export File to specify where to save the file.

The Import/Export option is useful as a backup or to share descriptor s with other projects.

Youmust import data into the correctportlet. Forexample, youcannotimport event data into the Actions portlet.

If one type of data depends on another, you must import the other data before importing the data that depends on it.

Delete— Delete the selected item. *Caution:* such deletions can impact anything else referring to what you are deleting.

This screen has the following columns:

Vendor Icon — The icon for this vendor.

Enterprise Number — The enterprise number for this vendor.

Vendor Name — The name for this vendor.

Vendor Editor

The Vendor editor configures (or re-configures) vendors.

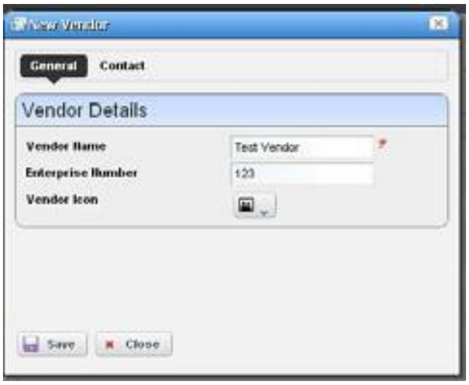

Vendor Name: Test Vendor
Enterprise Number: 123
Vendor Icon

The General panel has the following fields:

Vendor Name —A text identifier for the vendor.

Enterprise —A numeric identifier for the vendor.

Vendor Icon — Select an icon from the pick list.

From the Contacts panel, click the Add button to select from contacts in Cruz to associate with this vendor. See Contacts Portlet on page 215 and Contacts Editor on page 217 for information about configuring contacts.

Expanded Vendor Portlet

When you expand the Vendor portlet, besides sharing you can also click *Settings* to configure the columns that appear here (see Modifying Column Settings on page 144). This screen has the same columns and menu items available as the summary portlet. The Reference Tree in the Widgets panel displays a company icon for the selected vendor.

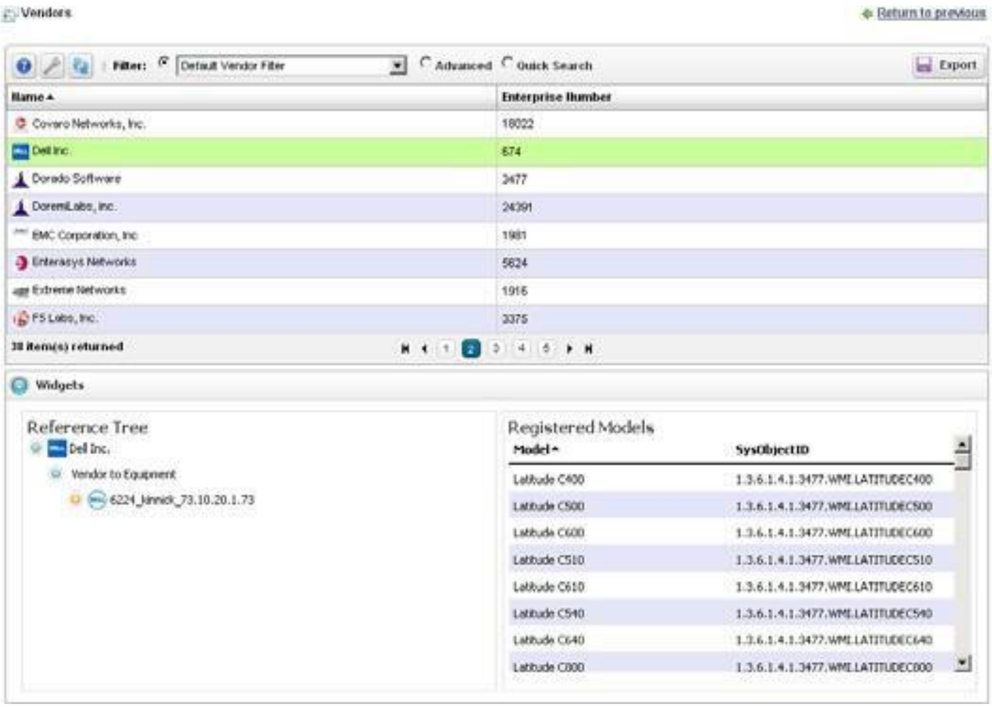

Name ▲	Enterprise Number
Covaro Networks, Inc.	18022
Dell Inc.	674
Dorado Software	3477
DoremiLabs, Inc.	24391
EMC Corporation, Inc	1981
Enterasys Networks	5624
Extreme Networks	1916
F5 Labs, Inc.	3375

38 item(s) returned

Reference Tree

- Dell Inc.
  - Vendor to Equipment
    - 6224_kinnick_73.10.20.1.73

Registered Models

Model ▲	SysObjectID
Latitude C400	1.3.6.1.4.1.3477.WMI.LATITUDEC400
Latitude C500	1.3.6.1.4.1.3477.WMI.LATITUDEC500
Latitude C600	1.3.6.1.4.1.3477.WMI.LATITUDEC600
Latitude C510	1.3.6.1.4.1.3477.WMI.LATITUDEC510
Latitude C610	1.3.6.1.4.1.3477.WMI.LATITUDEC610
Latitude C540	1.3.6.1.4.1.3477.WMI.LATITUDEC540
Latitude C640	1.3.6.1.4.1.3477.WMI.LATITUDEC640
Latitude C800	1.3.6.1.4.1.3477.WMI.LATITUDEC800

4

Presentation Capabilities

You can display devices and network arrangements in a variety of ways. This section describes the presentation capabilities user interface components and then provides some common tasks. If you already have a good understanding of the portlets and editors, go directly to the tasks you want to perform.

△ CAUTION:

If you installed a firewall on the application server, ports 80 and 8080 must both be open for topology feature to work.

Presentation Portlets and Editors

This section describes the following presentation portlets and editors used to viewyour network content graphically and to modify your graphical network view:

- Hierarchical View Manager
- Hierarchical View
- Map Context
- Maps and Hierarchical Views Together
- System Topology
- Topology Portlet
- Topology Toolbar
- Topology and View Configuration
- Alarms in Topologies
- Links in a Topology
- Topology Views

Hierarchical View Manager

Hierarchical View Manager portlet lets you create, edit and delete hierarchical tree models displayedin Hierarchical Views(described in the nextsection). These hierarchicalviews filter what appears in other portlets on the page with the Hierarchical View portlet.

The relationship to users and devices appears in Hierarchical View Manager Expanded.

Right-click to select from a menu with *New, Edit* and *Delete,* and *Refresh Members/Alarm State. Refreshing* re-queries, the database for members fitting the dynamic filter, or for new alarms for members. Selecting *New* or *Edit* displays the Hierarchical View Editor .

You can also *Tag* hierarchies, so Map views show hierarchies. See Tagon page 230 for more about how that works.

Access this portlet by selecting the Alarms/Events > Hierarchical View from the navigation bar.

Hierarchical View Manager Expanded

The expanded view displays the same information as the summary view, but displays the selected hierarchy's authorized users, creator, owner, and membership in the Reference Treesnap panel.

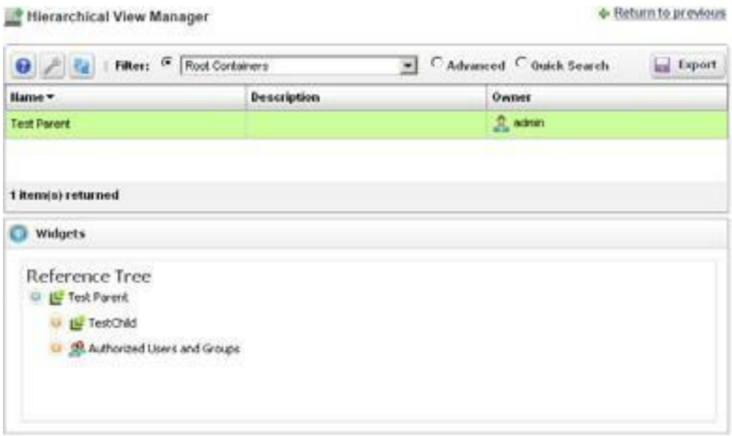
Hierarchical View Manager
Return to previous
Filter: ● Root Containers ○ Advanced ○ Quick Search
Export
Widgets

Name	Description	Owner
Test Parent		admin

1 item(s) returned

Reference Tree
- Test Parent
  - TestChild
  - Authorized Users and Groups

Multitenancy and Hierarchies

If you have the multitenancy option installed, you can limit a hierarchy’s visibility to a tenant site, or to the entire system. You can also import hierarchies to target multitenant domains with a command line importer. The command is `importcontainers` and is in the owareapps/redcell/ bin directory. This command takes the import file name an argument. The required domains should be available in the Cruz before import occurs. Example XML files (with the `<customer>` tag for domains) are in the owareapps\redcell\db directory.

Hierarchical View Editor

The Hierarchical View editor lets you create and manage hierarchies. Youcan also associate user authorizations with hierarchy models to specify which groups or users have access to contained items.

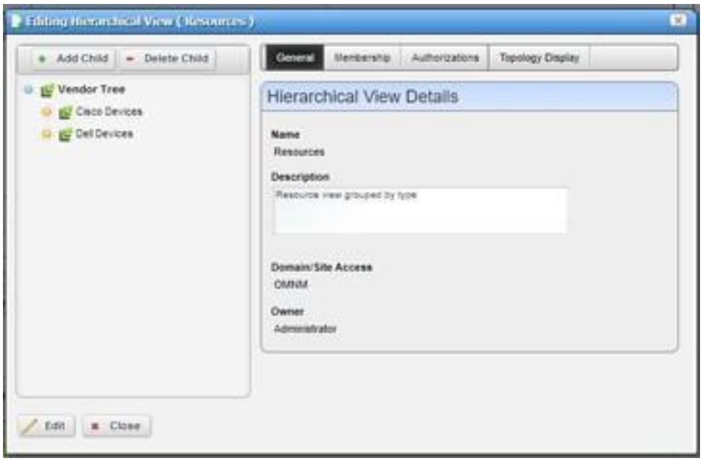

Editing Hierarchical View (Resources)

Hierarchical View Details

Name
Resources

Description
Resource view grouped by type

Domain/Site Access
OMNM

Owner
Administrator

In this editor, a tree panel on the left lets you build and navigate the hierarchical tree. Click *Add Child (*or *Delete Child)* to create (or remove) a node to/from the node you have selected in the tree. Clicking a node in the tree displays the tabbed panel on the right where you can edit it.

The Hierarchical View Details panel has the following tabs:

- General
- Membership
- Authorizations
- Hierarchical View Display

Click the labels at the top of the screen to access these. Alarm states and severities are recalculated and propagated for hierarchies as they are for System Topology on page234.

General

This panel has the following fields:

Name — The hierarchical view identifier.

Description — A text description of the hierarchical view.

Domain/Site Access—If you have the Multitenancy option installed, to expose a hierarchical view or sub-hierarchical view within a site the hierarchy must specify that site. Youcan confine a hierarchical view’s contents to what is visible on a single site. All hierarchical view contents are visible for the master site.

Owner—Select an owner for the hierarchical view. The owner of a hierarchical view can also change the ownership of the hierarchical view.

Update View Authorizations—Clicking the link here automatically includes the creating user in those authorized to see the hierarchical view. See Authorizations for more about them.

Membership

Hierarchicalview membership defines the inventory items that are in a hierarchical view. Youcan select either a *Static* membership, which cannot change, or a *Dynamic* one, based on a filter or existing groups. When the Cruz application evaluates the filter it adds the resulting items as members in the hierarchical view.

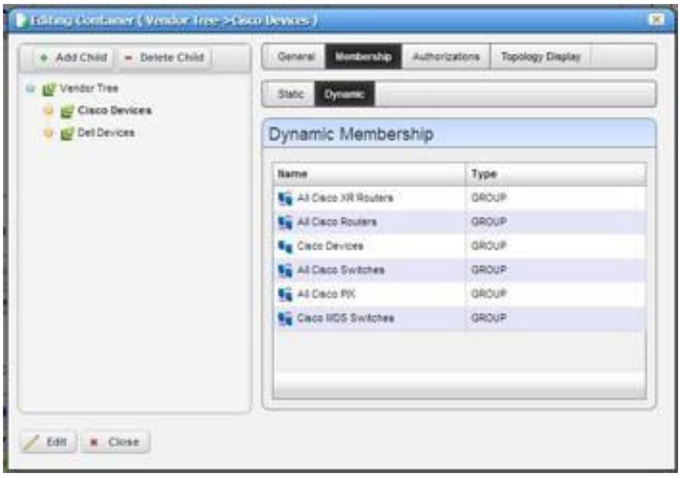

Name	Type
All Cisco XR Routers	GROUP
All Cisco Routers	GROUP
Cisco Devices	GROUP
All Cisco Switches	GROUP
All Cisco PIX	GROUP
Cisco MDS Switches	GROUP

The sub-tabs at the top of the screen let you edit these types. You can add individual items with the *Static* tab, or the results of a *Dynamic* filter with that tab. See Managed Resource Groups on page 173 for more about the specifics of editing these dynamic groups.

When you add an item or filter to your hierarchical view, notice that the subsequent screen contains a pick list *Select an entity of the following type.* The contents of that list can contain several types of managed objects, including Contact, Equipment and Subcomponent, Interface, Location, Managed Equipment, Port, Service, and Vendor. Select the type appropriate for your hierarchical view.

Click *Save* to preserve the membership you have configured. If you *Group by Entity Type* (at the bottom of the screen) rather than *None*, the list of devices appears in a tree, with each node as an entity type. Click the plus (+) to the left of the entity label to expand the tree.

Authorizations

This tab configures user or role access to the hierarchical view you are editing. By default, no authorization exists to see a hierarchical view or its contents, so you must permit specified users and roles to have access before any hierarchical views or their contents are visible in the Hierarchical View portlet.

Click *Add User* or *Add Role* to select the users or groups with permission to access the hierarchical view you are configuring. By default, hierarchical views are accessible to everyone.

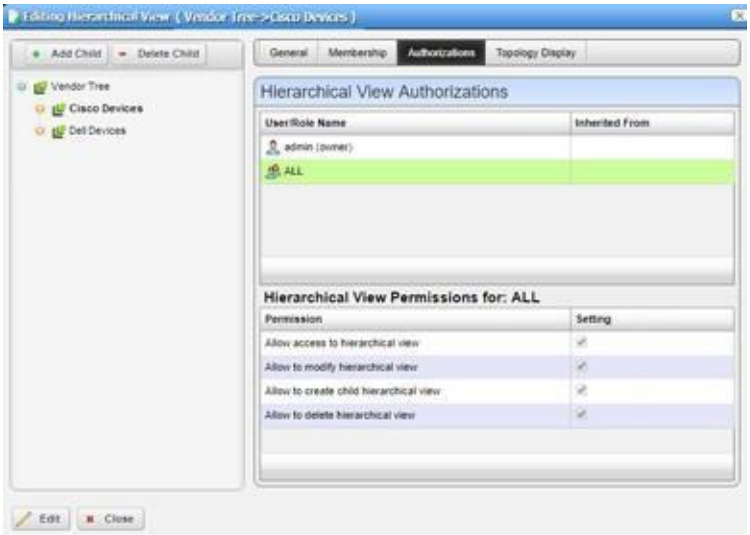

Hierarchical View Authorizations

User/Role Name	Inherited From
admin (owner)	
ALL	

Hierarchical View Permissions for: ALL

Permission	Setting
Allow access to hierarchical view	✓
Allow to modify hierarchical view	✓
Allow to create child hierarchical view	✓
Allow to delete hierarchical view	✓

Each entry in the Hierarchical View Authorizations list specifies the name of the user or role, and whether the entry is inherited or not. A child hierarchical view by default inherits the authorizations from parent hierarchy, no explicit authorizations for child hierarchical views are necessary. Edit any authorizations in the parent.

When editing a child hierarchical view, click a listed authorized user or role and its permissions appear in the panel at the bottom of this screen.

Clicking *Save* preservesanyalterations you have made. Confirm the hierarchical viewisconfigured as you like by examining it in a Hierarchical View portlet.

Hierarchical View Display

This tab configures how the hierarchical view appears in the System Topology portlet.

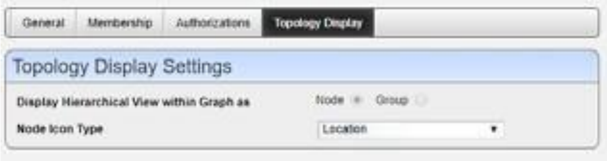

Topology Display Settings

Display Hierarchical View within Graph as   Node ● Group ○

Node Icon Type   Location ▾

Selectedhierarchicalviews’ labelsappearinthe Topology’stitlebar. Configurethefollowingdisplay settings in this panel:

Display Hierarchical View within Graph as — Select either *Node* or *Group*.

Node Icon Type—This appears if you select *Node*. Use the pick list to select among the various icon types as is appropriate for your hierarchical view.

Group Style — Select either *Default (Rectangle Shaded Group)* or *Cloud (Cloud Background Image)*. The group is like the Expand Grouped capability described in Topology and View Configuration on page 239. The Cloud is a cloud icon like the one you can add to views as described in Topology Toolbar on page 235.

Display Hierarchical View Name within Group—This appears if you select *Group*. Check it to display the hierarchical view name as a label within the group.

Hierarchical View

The (non-instanceable) Hierarchical View portlet displays configured hierarchical views authorized for the logged-in user, in the color of the most severe alarm for equipment within that hierarchical view. Because it is non-instanceable, only one can appear on apage.

Filter what appears on a page using the Hierarchical View portlet. Select a hierarchical view, and the rest of the portlets on that page filter their data reporting to reflect that hierarchical view's contents. The only caveat to this type of filtering is that you can add only one Hierarchical View portlet to a page.

Access this page by selecting the Alarms/Events > Hierarchical View from the navigation bar.

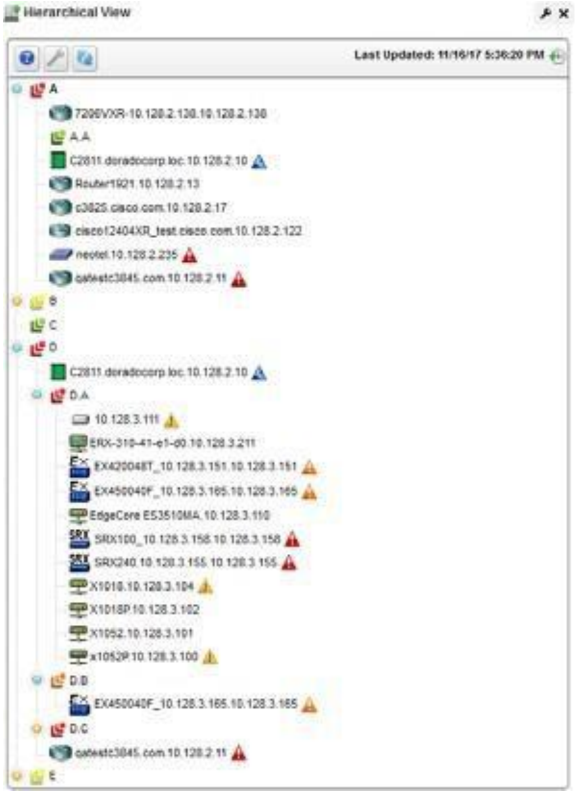
Hierarchical View
Last Updated: 11/16/17 5:36:20 PM
A
7206VXR-10.128.2.130.10.128.2.130
A.A
C2811.doradocorp.loc.10.128.2.10
Router1921.10.128.2.13
c3825.cisco.com.10.128.2.17
cisco12404XR_test.cisco.com.10.128.2.122
neotel.10.128.2.235
qatestc3845.com.10.128.2.11
B
C
D
C2811.doradocorp.loc.10.128.2.10
D.A
10.128.3.111
ERX-310-41-e1-d0.10.128.3.211
EX420048T_10.128.3.151.10.128.3.151
EX450040F_10.128.3.165.10.128.3.165
EdgeCore ES3510MA.10.128.3.110
SRX100_10.128.3.158.10.128.3.158
SRX240.10.128.3.155.10.128.3.155
X1018.10.128.3.104
X1018P.10.128.3.102
X1052.10.128.3.101
x1052P.10.128.3.100
D.B
EX450040F_10.128.3.165.10.128.3.165
D.C
qatestc3845.com.10.128.2.11
E

You can configure the height of this portlet by clicking on the settings button. A popup screen appears that allows you to enter a minimum height and a maximum height, both of which are in pixels. The minimum height is how tall it will be when the portlet first loads. If the maximum height is greater than the minimum height, then this means that the portlet height can grow when container nodes areexpanded. Ifcontainer nodes areexpanded beyond the maximum height, then a vertical scroll bar appears within the portlet. Note that if you change these values and click apply then you will still need to refresh the browser screen for the portlet to re-render with these new settings.

Expand the Hierarchical View tree to viewits content. Hierarchical Viewcontents sort alphabetically, and alarms appear to the right of equipment displayed.

The hierarchical view selected acts as a filter for a screen's other Cruz portlets. If you select "Folsom" as a location in the Hierarchical View portlet, then onlyitems related to Folsom devices appear in the other portlets on the page. If you select a parent hierarchical view, that expands the selection to include all child hierarchical views' selections. It does not, however select everything. Youcan configure hierarchical viewsin Hierarchical View Editor on page 225. Youmay have to wait a few moments to see a hierarchical view's contents accurately.

Portlets that respond to Hierarchical View or Map Context "filtering" include the following: Audit Trail, Event History, Locations, Vendors, Contacts, Managed Resources, Ports, Authentications, Discovery Profiles, Monitors, Services. If you have no hierarchical views configured, the other portlets appear empty when Hierarchical View is on the same page.

See General > Entity Change Settings on page 44 for the way to set the summary portlet refresh interval. The default is 40 seconds. You can also refresh this portlet manually, which is the equivalent of selecting each node within the tree and selecting both Refresh Members and Refresh Alarm State. The Refresh button at the top of the portlet provides an easier way to do this in one step.

NOTE:

If a Hierarchical View portlet displays unexpected results or no members at all, right-click it to refresh its membership or alarm severity/state. Remember also that the visible changes may take a moment to appear.

Right clicking a hierarchical view displays the following menu items:

Refresh Members—Query the database again to populate any dynamic filter that is part of the hierarchical view.

Details — Opens a details panel with a list of the hierarchical view's contents (*Members*) as well as hierarchical view members' *Alarms* (Alarms and Event History) and *History* (Audit trails and Configurations).

Refresh Alarm State—Re-query the database to update the hierarchical view's alarmstate based on its contents.

Edit Resources—Open an editor screen for the hierarchical view that lets you change common attributes within it.

Topology—Displaya hierarchical view in the System Topology portlet where you can drill in to see its contents (see System Topology on page 234).

Tag— Enter map location coordinates for the hierarchical view. See Hierarchical View in Tenant Sites and Tag for more details.

Share with User

Hierarchical View in Tenant Sites

Within a Multitenant environment, only hierarchical views configured to appear in tenant sites appear there. If that hierarchical view contains equipment only visible on the master site, those devices will appear below the hierarchical view node, but will have no impact in filtering other portlets, like Alarms, for example.

Tag

When creating a hierarchical view or customer tag, Cruz automatically selects the latitude and longitude of the address entered for a location. Taga hierarchical view by right clicking it from the Hierarchical View Manager on page 224.

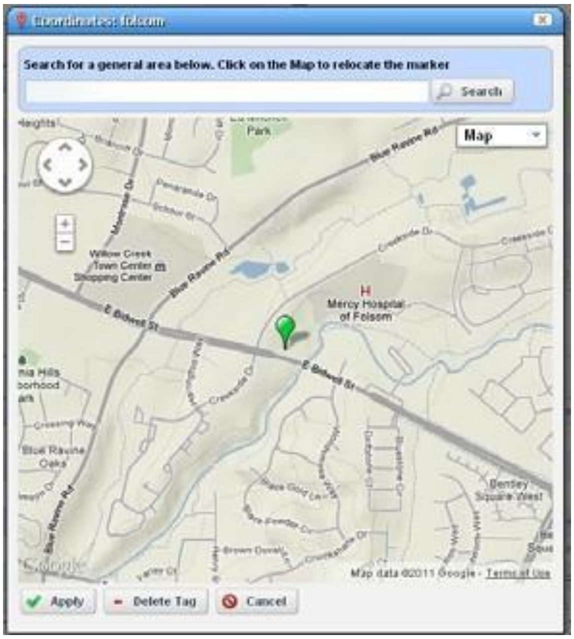

Search for a general area below. Click on the Map to relocate the marker

You can also enter the address in the Search field or click and drag the marker that appears on this screen. Click *Apply* to accept the re-location. A *Delete Tag* button appears when you have created a tag, and lets you remove it. *Cancel* closes the screen.

NOTE:

Youcanzoominoroutonthedisplayedmapwiththe+and-buttonsintheupperleftcornerofthis screen.

Map Context

Inaddition to displayingfiltered-by-Hierarchical Viewportlets, you can viewdiscovered devicesin the *Map Context* portlet, automatically placed by location.

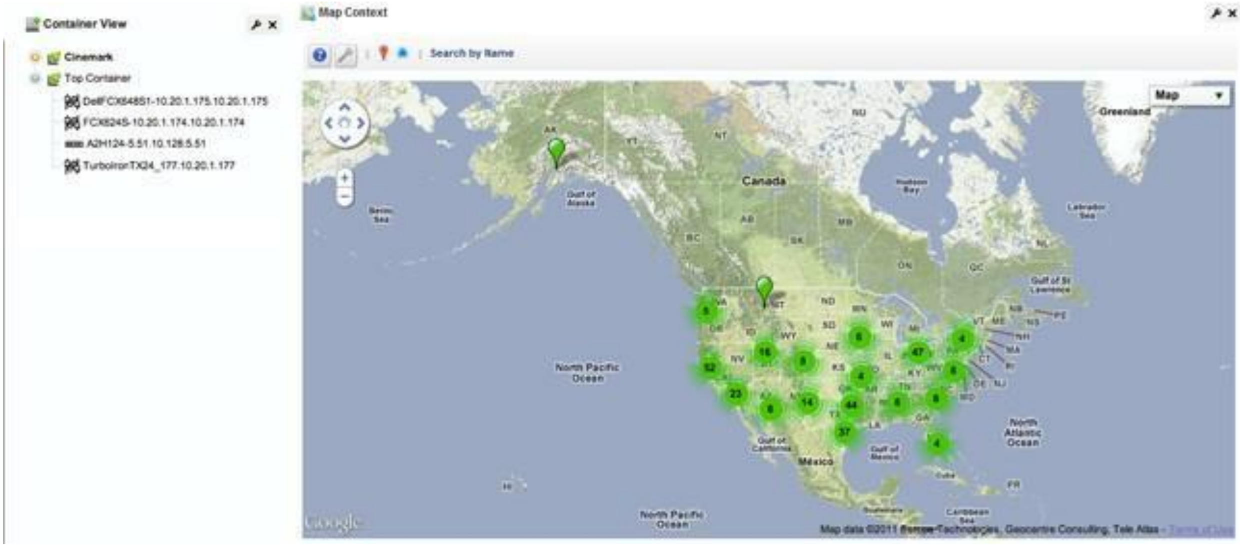

Notice that you can move the center of the map with the arrows in its upper left corner above the zoom in/out (+/-) buttons. The menu in the upper right corner lets you select a *Map* or *Satellite* views, and fine-tune them to include labels, terrain and so on.

In addition to the *Help* and *Settings* icons at the top of this portlet, you can also *Toggle Marker Style* (pushpins or triangles), *Toggle Marker Clustering* (combine markersinto cluster marker when they are near each other), or *Search by Name* for a location. Clustered markers display the number of separate markers combined within them.

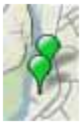

NOTE:

The Search function is case-sensitive. Omit the initial letter if you are uncertain about capitalization for a tagged location.

Clicking the *Settings* icon produces a screen where you can configure the default marker style, whether clustering is enabled, and where you can save the current map boundaries (*Save Current Bounds*), which appear, read-only, below that option.

See General > Entity Change Settings on page 44 for the way to set the summary portlet refresh interval. The default is 40 seconds. This screen also sets the default center of your map.

The page layout controls the width of the map. However, you can control the height of the map with the *Look and Feel* configuration in the *Advanced Styling* tab.

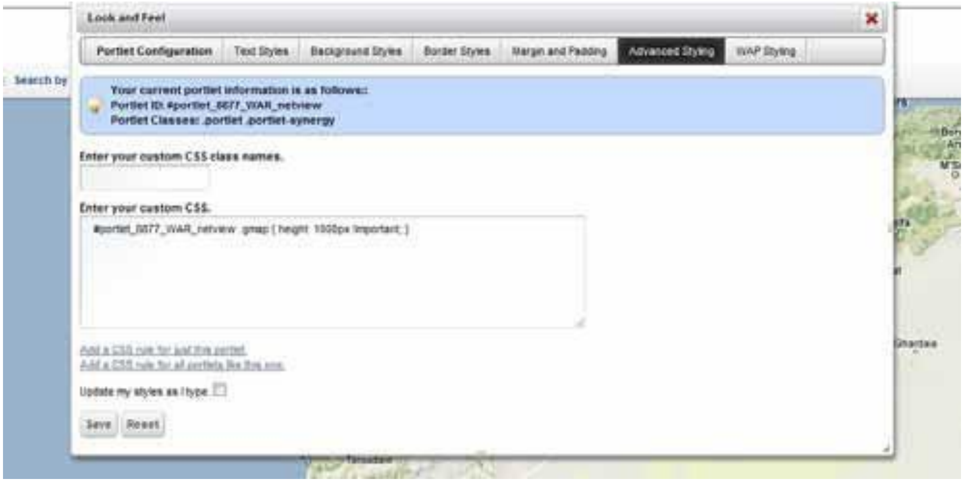

Look and Feel

Portlet Configuration | Text Styles | Background Styles | Border Styles | Margin and Padding | Advanced Styling | WAP Styling

Your current portlet information is as follows::
Portlet ID: #portlet_8877_WAR_netview
Portlet Classes: .portlet .portlet-synergy

Enter your custom CSS class names.

Enter your custom CSS.

```
#portlet_8877_WAR_netview .gmap { height: 1000px !important; }
```

Add a CSS rule for just this portlet.
Add a CSS rule for all portlets like this one.

Update my styles as I type. □

Save | Reset

Add the following line to the custom CSS settings in this tab:

```
    #portlet_8877_WAR_netview .gmap { height: 1000px !important; }
```

This sets the height of the map context portlet to the configured number of pixels, here, 1000.

Access this tab from the drop-down originating with the word *Map* in the top right corner of the portlet.

Configure mapped hierarchical view or customer locations with the *Tag* menu item. See *Tag* on page 230 for an explanation. See Maps and Hierarchical Views Together below for more about their joint capabilities.

Map Context without Hierarchical Views

If a page has no hierarchical view, then the Map Context can act like a Hierarchical View too. It displays all tagged resources within the system (see Tagon page 230). Clicking on a tagged item behaveslikeclickinga Hierarchical Viewportlet, confining displayed resources, alarms, and so on, to those for the selected tag.

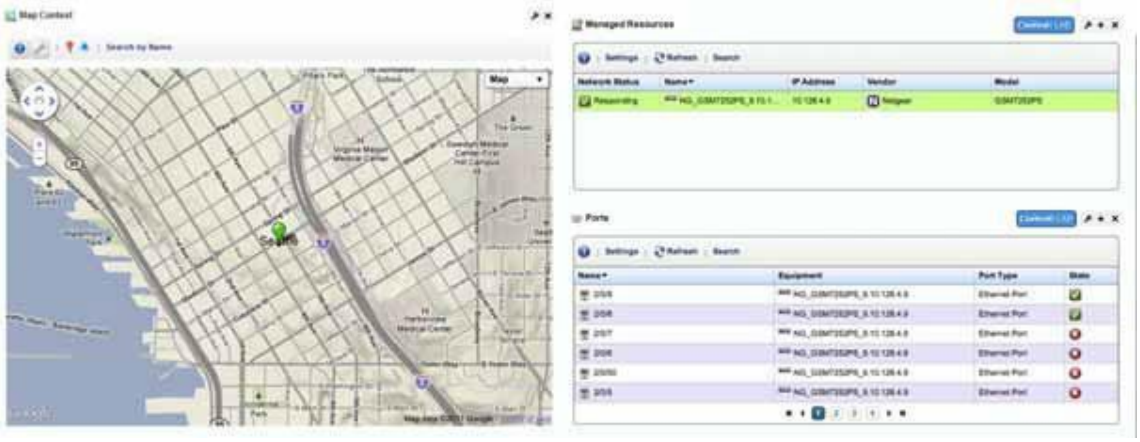

Each tagged coordinate is cross correlated with the Alarm severity table (if there are alarms against it) and its color reflects the current Alarm severity.

Maps and Hierarchical Views Together

A map context portlet is in *Standalone* mode when no Hierarchical View portlet is on the same page.

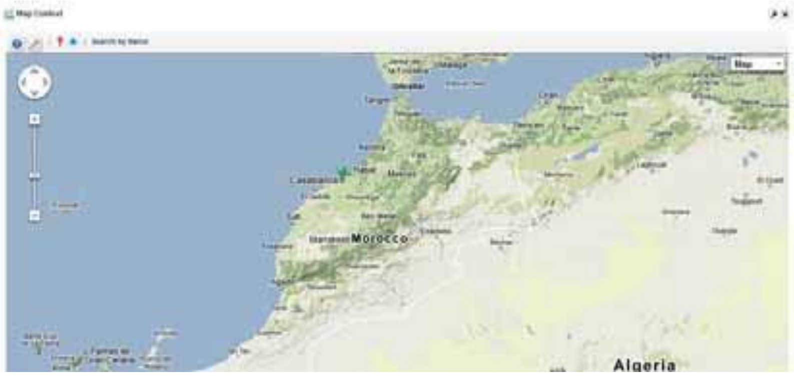

In Standalone mode you can determine exactly what portion of the map appears through the *Settings* option (the wrench icon).

The map context portlet is in *Hierarchical View Context* mode when a Hierarchical View portlet is on the same page. In these Hierarchical View Context configurations, the Hierarchical View portlet determines what appears in the Map portlet, so the Map Context portlet resets its boundaries based upon the geographic position of the selected hierarchical view's members. Forexample, you can select a hierarchical view (Morocco) resulting in two clustered pins for both Casablanca and Tangier.

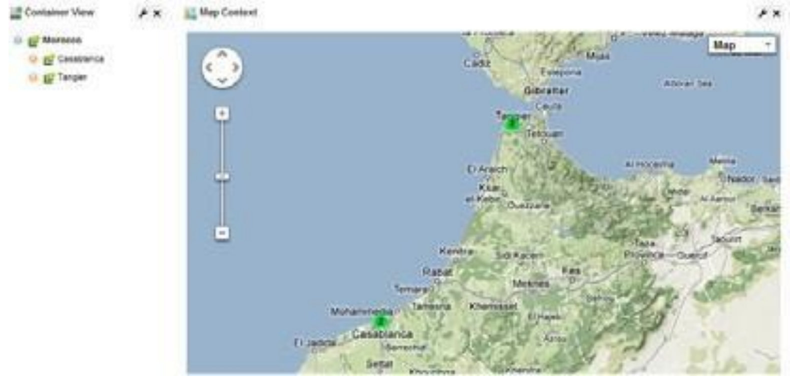

However, if you select Casablanca from the hierarchical view the map automatically changes its presentation and boundaries based upon the members of the new selection. The view zooms to the street level in Casablanca.

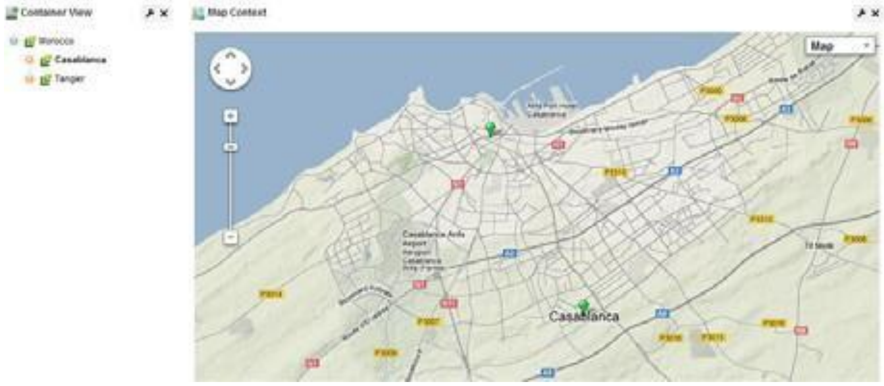

If you select Tangier and map presents the hierarchical view’s sites in a street-level view of Tangier.

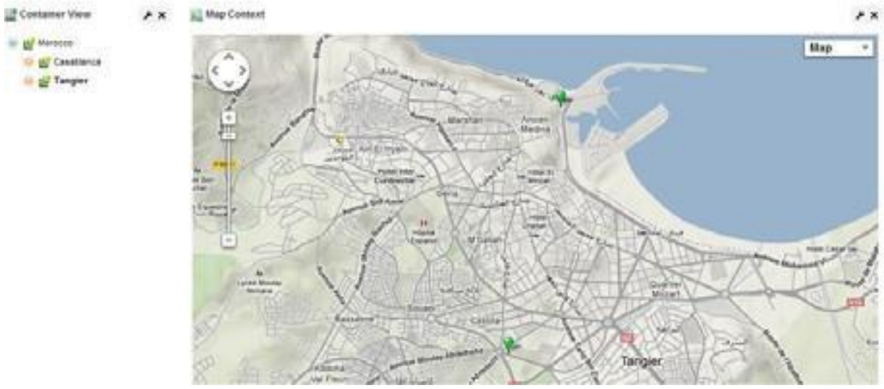

System Topology

The System Topology portlet displays discovered devices, mapping them in relationship to each other. It also lets you store and retrieve views you have arranged, as well as configure the default view (see Topology and View Configuration on page 239 for more about these capabilities). The topology toolbar is common to both this portlet and the Topology Portlet .

See General > Entity Change Settings on page 44 for the way to set the summary portlet refresh interval. The default is 40 seconds.

The color displayed in these topologies indicates the alarm severity of the node or link (“edge”) only. No color or icon indicates a device’s network status or availability, although hovering the cursor over a node displays that information.

NOTE:

You can increase or decrease the size of icons in the Equipment Editor. Right-click a device in the Managed Resources portlet and select *Edit.* In the Extended Details panel, select *Settings* and a *Topology Icon Size* pick list appears as one option to configure.

System Topologyportlet also acts like a filter. Portlets like Alarms and Ports respond to clickinga topology node, by displaying information relevant to only thatnode.

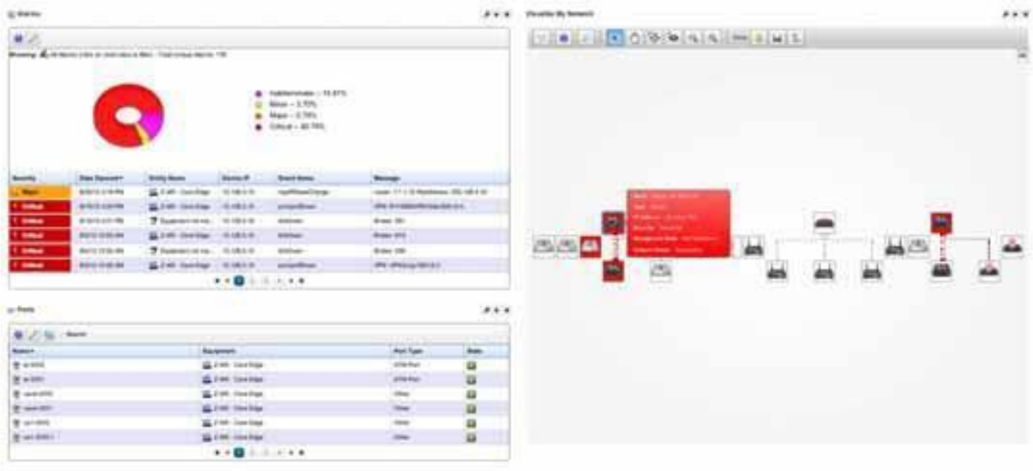

Topology Portlet

Use the Topology portlet to define topologies. This portlet is access from the Managed Resource portlet by right clicking resource s and then selecting Topology. The topology toolbar is common to both this portlet and the System Topology portlet. It also lets you configure the default view (see Topology and View Configuration on page 239 for more about these capabilities).

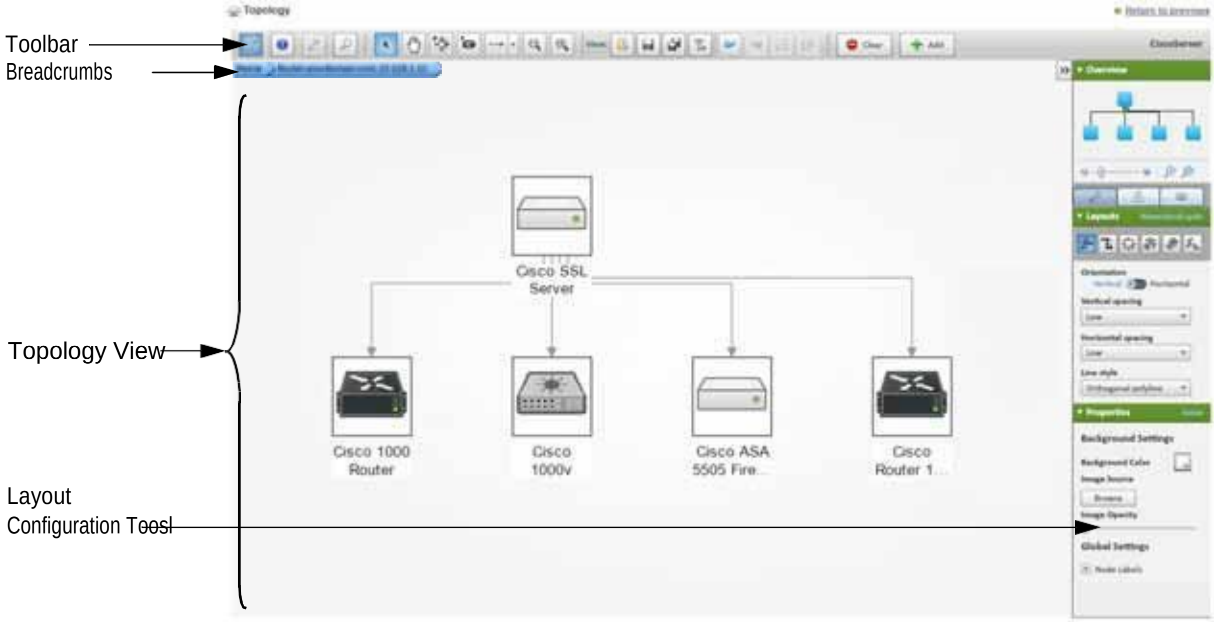

Topology Toolbar

The topology toolbar helps navigate through the topology, load a view, save a view, and so on.

When you click the Toggle Design Mode tool, several additional tools appear that let you manipulate the topology view.

<table>
<tr><th>Tool</th><th>Description</th></tr>
<tr><td></td><td>
<p>Toggle Design Mode (Ctrl +D) shows and hides the design tools, such as the line draw, undo/redo, group/ungroup, and so on.</p>
<p>Click this to turn on the design tools, such as line draw, edge filtering, undo/redo, and so on. You can configure users’ permissions for Design Mode in <i>Control Panel > Permissions manager</i>. To disable Design mode, uncheck Visualizer permission both ADD and DELETE for all assigned roles (typically these include User, Power User, and sometimes Administrator).</p>
<p>To enable Design Mode check Visualizer permission, ADD and DELETE for roles (typically for the Administrator role). To give other-than-administrator users no Design Mode permission, uncheck ADD and DELETE for User and Power User assigned roles.</p>
</td></tr>
<tr><td>?</td><td>Help Click this to turn access the online help for this screen.</td></tr>
<tr><td></td><td>
<p>Default Click the wrench icon to configure the default view. If the Topology portlet is on a page not driven by another Context—for example, Hierarchical—and you have write permission, then this icon appears. Clicking this lets you associate the Topology portlet on the current page to a selected view. To return to the default</p>
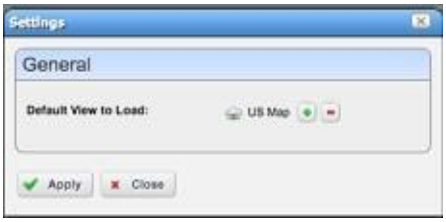
<p>Settings</p>
<p>General</p>
<p>Default View to Load: US Map</p>

<p>network view, click the red minus (-) button in the settings. Any view change requires a page refresh after applying the revised setting.</p>
</td></tr>
<tr><td></td><td>
<p>Search Node Elements within this Graph (Ctrl +S)</p>
<p>Search for a node. This opens a screen displaying the search results, name, type of node and the node’s alarm severity. Click Select / Center Item on Graph to select a listed item.</p>
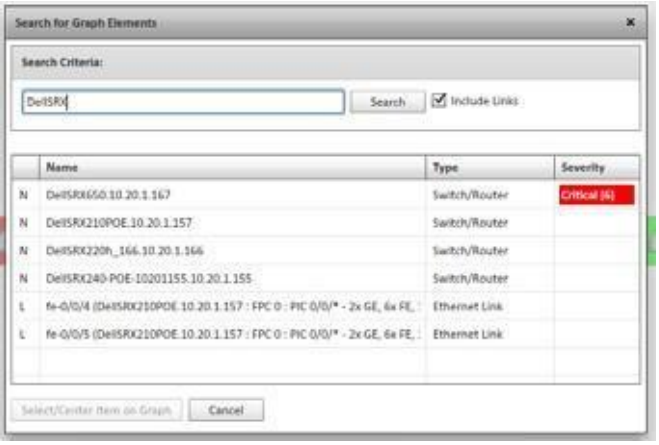
<p>Search for Graph Elements</p>
<p>Search Criteria: DellSRX</p>

<table>
<tr><th></th><th>Name</th><th>Type</th><th>Severity</th></tr>
<tr><td>N</td><td>DellSRX650.10.20.1.167</td><td>Switch/Router</td><td>Critical (6)</td></tr>
<tr><td>N</td><td>DellSRX210POE.10.20.1.157</td><td>Switch/Router</td><td></td></tr>
<tr><td>N</td><td>DellSRX220h_166.10.20.1.166</td><td>Switch/Router</td><td></td></tr>
<tr><td>N</td><td>DellSRX240-POE-10201155.10.20.1.155</td><td>Switch/Router</td><td></td></tr>
<tr><td>L</td><td>fe-0/0/4 (DellSRX210POE.10.20.1.157 : FPC 0 : PIC 0/0/* - 2x GE, 6x FE,</td><td>Ethernet Link</td><td></td></tr>
<tr><td>L</td><td>fe-0/0/5 (DellSRX210POE.10.20.1.157 : FPC 0 : PIC 0/0/* - 2x GE, 6x FE,</td><td>Ethernet Link</td><td></td></tr>
</table>

<p><b>Note:</b> Search also finds links or “edges” between devices and saving a view preserves displayed links’ appearance.</p>
</td></tr>
</table>

Tool	Description
↖	**Selection Tool** The cursor selects nodes. Click and drag around nodes to select several.
✋	**Pan Tool** The hand moves the background.
⚙	**Shortest Path Tool** Click two nodes to highlight the shortest path between them. This simple tool looks for the shortest set of lines and does not take into account factors like route cost, bandwidth, link status, and so on.
👁	**Bifocal Effect** Move the cursor to magnify nodes under it. Handy in a crowded view.
→ ▾	Line Draw Click to select the type of line to draw, then shift + click two icons onscreen to draw the line. hidden with other view tools
🔍	**Zoom Out** These magnifying glass icons change the magnification for the view.
🔍	Zoom In
📂	Load a View opens a saved view.
🗑	Delete and Reset Factory View tool
💾	Save this View tool saves the current. Views include visible nodes and links, but you cannot save the location of these nodes. (See Map Context on page 231 for a possible alternative.)
⏷	**Edge Filtering** tool opens the Edge Filtering window, where you select whether to enable edge filtering and then which links appear in the topology. By default, edge filtering **is not** enabled, and all links are selected.
↩	**Undo** tool reverses the last action.
↪	**Redo** tool puts the last action back.

Tool	Description
	Group tool puts the selected objects, lines, and labels together allowing you to move the in tandem. Ctrl +click to select multiple objects. These two icons group or ungroup selected icons labels and lines together so you can move them in tandem. Ctrl +click to multi-select icons. When you create a group, the Properties panel provides additional configuration parameters. These include the Header panel where you can configure whether the group header is Visible, its Label the Background and Text Color. Click the minus in the header to minimize the group (and plus to expand a minimized group). The Content panel lets you configure whether the group appears as a Panel or Cloud, and its *Background* and *Stroke* colors. 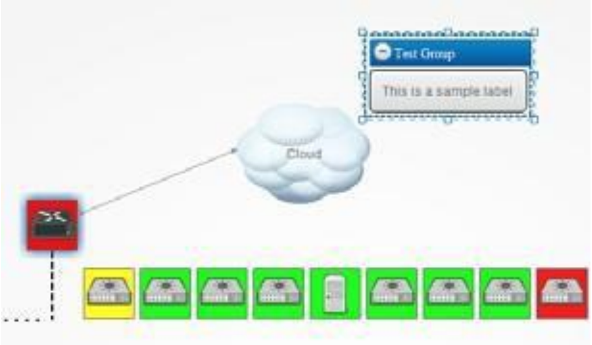
	Ungroup tool separates all objects, lines, and labels and you can no longer move the objects in tandem.
Clear	Clear tool empties the topology view.
Add	Add tool open the Add a Graph Element window where you select the element type (Label, Cloud, or Linked View) and whether it is a static placement. Use the *Properties* panel to configure the font, background color, label contents, and so on, after you select the added element. **Note:** If you configure and save a Drill-in view with the design tools, then that view persists for all drill-ins from that device until you remove it an icon that appears between view Open and Save when it is enabled. Deleting such a drill-in view restores the default settings. 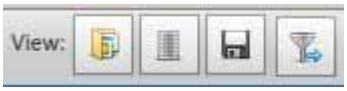 When you add these elements, you can elect to check *Static Placement* and they will not move with graphic elements when they are automatically re-arranged. You can, however, click and drag them.

Topology and View Configuration

Click and drag displayed portions of this view to see other topology parts. To move the display more, click in the Overview Panel. Youcan also expand/collapse the panels on the left of the screen by clicking their title bars. (Figures below display them expanded.)

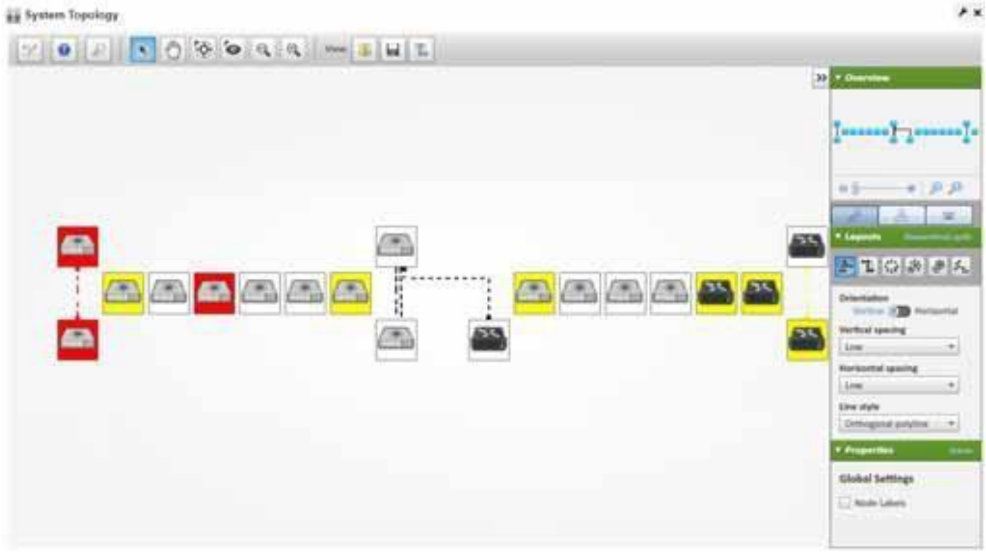

Nodes appear colored according to the alarm severity on the device, and white if no alarm exists for the device. Hover the cursor over an icon or link between icons to see a small screen describing its device (*Name*, *Type, IP address*), network status (*Responding/Not Responding*) and alarm severity. Click an icon to highlight it (or click its name in the Top-Level Nodes Tab list) and its connections to the network. See Alarms in Topologies on page 246 for more about the alarm severities indicated by icons in topology.

Click the Legend Tab to see the lines meaning, links, and alarm colors. Hover the cursor over a link to see its type described.

The topology toolbar helps navigate through the topology, load a view, save a view, and so on.

The right of the topology view displays the following panels:

- Overview Panel
- Layouts Panel
- Properties and Settings > Properties
- Legend Tab
- Top-Level Nodes Tab

Click the triangles to the left of these panels' labels to collapse or expand them.

In addition to the screen components immediately displayed, you can right-click an icon or component, and *Drill in* or *Expand* a device to see its subcomponents. If you expand, then its subcomponents appear with the rest of the topology. If you *Expand Grouped*, then the subcomponents appear in a minimize-able block (hover your cursor to see the block in color and click the circled minus to minimize the group).

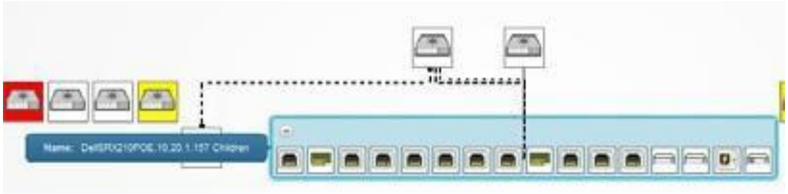

If you drill in, other components do not appear. In addition, you can select the Details menu option to open another browser window with the selected node's Details. The Event History menu item also opens a new browser window with the selected node's event history.

NOTE:

If you want to initiate Actions on a node or its components, do so by right clicking the *Details* screen's Reference Tree.

The Layouts Panel selections determine the arrangement ofsuch expansions or drill-ins.

When you drill in, the path back to the top level appears below the topology.

Click the level where you want to "drill out," or click *Home* to go to the top level.

If you right click the blank area of the screen, you can *Export* it as either a .png image or GML (graphic markup language) or print the displayed topology.

Youcan also right-click to *Remove Node* and delete a device from a view. Youcannot add nodes to a view; you must add them when you create the view. Youcan visualize Managed Resource Groups, however, or simply go to the expanded Managed Resources portlet, select multiple resources (Ctrl +click), and the right-click > Topology.

NOTE:

BecauseTopologyusesAdobeFlash,menuitemsappearforthatsoftwarewhenyouright-clicknodes. Thisincludes *Settings, GlobalSettings* and *AboutFlash* menuitems. Thetextbelowdoesnotdiscuss these since they relate to Adobe products.

Overview Panel

The Overview panel displays a thumbnail of the entire topology view. Click on a location to center on it in the topology view.

Use the slider at to change the magnification of your view. The icons to the right of the slider let you click them to fit visible icons vertically and both vertically and horizontally. Youcan also click and drag the cursor within this overview to change the magnification.

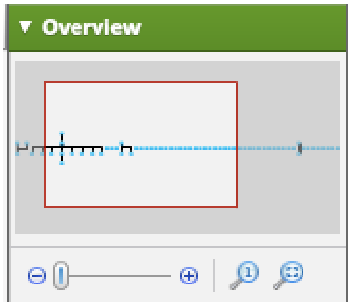

Layouts Panel

The Layouts panel lets you select and configure the automated node layout type that appearsin the topology display.

Access this panel by selecting the Properties and Settings > Layouts tab.

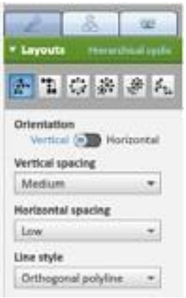

Orientation
Vertical Horizontal

Vertical spacing
Medium

Horizontal spacing
Low

Line style
Orthogonal polyline

Useone of the followinglayout tools to specifyalayout type and its settings. The fields and options available depend on the layout selected.

△ CAUTION:
Layoutsettings becomefixedonceyou saveaview. Forexample, whenyou saveaviewwithalinestyle, thatlinestyleisnotsomethingyoucanalterlater.SavedViewsdonotpreserve linksifthey include non-grouped hierarchical views (single-node representations).

Layout Tool	Description
	Hierarchical-Cyclic layouts arrange connections in a hierarchy. 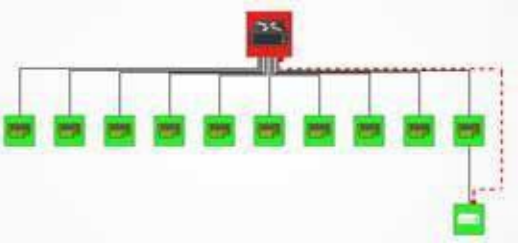 Use the following settings to alter its appearance: • **Orientation** (Vertical or Horizontal) • **Vertical Spacing** (High, Medium, or Low) • **Horizontal Spacing** (High, Medium, or Low • **Line Style** (Orthogonal polyline, Straight, Straight polyline, Curved polyline, or Orthogonal curved)
	Orthogonal layout connections include right-angles. 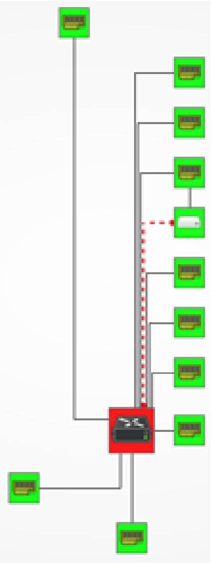 Use the following settings to alter its appearance: • **Grid Spacing** (High, Medium, or Low) • **Use Diagonal Edges** enables edges that have non-right angles

Layout Tool	Description
	Circular layouts arrange all nodes in a circle. 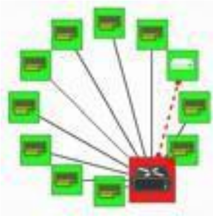 Use the following settings to alter its appearance: • **Layout Angle** (360 or180) • **Nodes spacing** (High, Medium, or Low)
	Balloon layouts display links between managed objects in a balloon tree structure. The root is typically whatever device you expanded or drilled into. 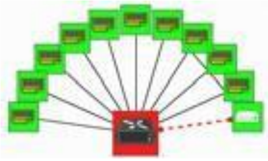 Use the following settings to alter its appearance: • **Root/Child wedge angle sector** uses radio buttons to determine the angle (360, 180). The root sector determines how much of an arc around that root the child nodes fill. The child sector determines the orientation around the child nodes. • **Root selection policy** specifies the item you want at the center of this view (Directed [a pop-up appears with the remaining selections], Most closed/ surrounded/weighted). • **Equal angle distribution** specifies whether to distribute nodes at equal angles.
	Radial layouts arrange nodes in concentric rings. 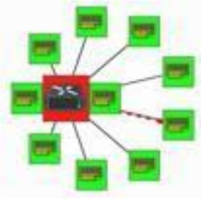 Use the following settings to alter its appearance: • **Layout angle** uses the radio buttons determine the angle (360, 180). • **Root selection policy** specifies the item you want at the center of this view (Directed [a pop-up appears with the remaining selections], Most closed/ surrounded/weighted).
	Organic layout produces a static GEM layout, without any parameters to tune. 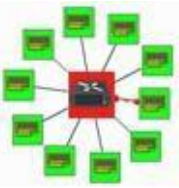

Properties and Settings >Properties

Thispanelconfigures the System Topologyviewproperties. Thispanelhasthefollowingfields(you must click the *Design Mode* icon in the upper left corner to see all of them).

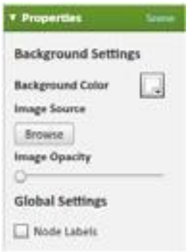

Background Settings

Thebackground settings allowyou to set:

Background Color—Click the icon to see a color selector where you can select the background color for the System Topology view.

Image Source — Click the *Browse* icon to select a graphic for the background.

Image Opacity — Use the slider to set the background opacity.

Global Settings

Node Labels— Check to label nodes in the System Topology view.

Legend Tab

This displays the meaning ofvariouslink types and alarmseverity colors in the System Topology view. It describes only the type of links that appear onscreen in thetopology.

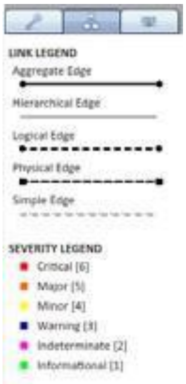

Top-Level Nodes Tab

This displays a legend of icon types followed by a count (in parentheses) of how many of each appear in the topology. The switch at the bottom of this panel centers the display around the selected icon.

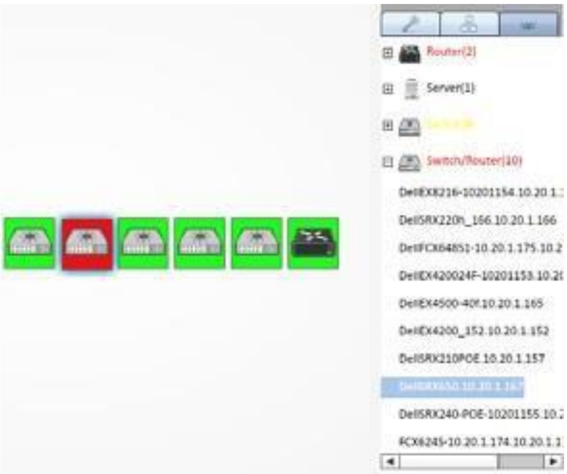
Router(2)
Server(1)
Switch/Router(10)
DellEX8216-10201154.10.20.1.
DellSRX220h_166.10.20.1.166
DellFCX64851-10.20.1.175.10.2
DellEX420024F-10201153.10.2
DellEX4500-40f.10.20.1.165
DellEX4200_152.10.20.1.152
DellSRX210POE.10.20.1.157
DellSRX240-POE-10201155.10.2
FCX6245-10.20.1.174.10.20.1.1

Click the plus (+) to the left of the inventory cate gory icons to display a list of devices in that category in the topology. Click on a list item to highlight that device and its network connection in the topologyview. Acolored glowhighlights it and its network connection(s). The listed inventory changes if you drill in.

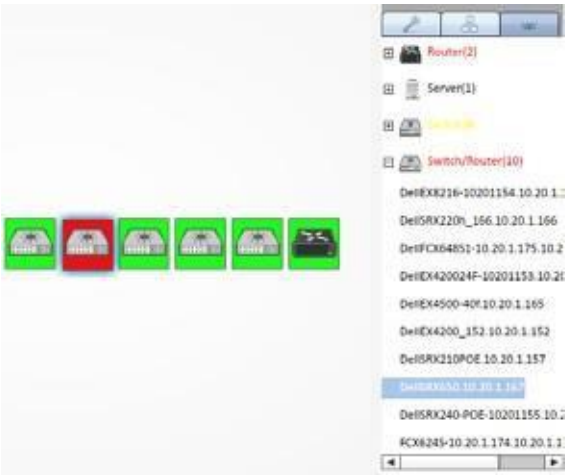
Router(2)
Server(1)
Switch/Router(10)
DellEX8216-10201154.10.20.1.
DellSRX220h_166.10.20.1.166
DellFCX64851-10.20.1.175.10.2
DellEX420024F-10201153.10.2
DellEX4500-40f.10.20.1.165
DellEX4200_152.10.20.1.152
DellSRX210POE.10.20.1.157
DellSRX240-POE-10201155.10.2
FCX6245-10.20.1.174.10.20.1.1

The listed text appears in the alarm color of the device. See Alarms in Topologies on page 246.

Alarms in Topologies

Coloredrectangles appear around topology nodes to indicate the highest alarmon them. Expand or drill in to see alarms on the sub-components. For information about the alarm, hover your cursor over the device or subcomponent, and a tooltip appears describing the alarm’s severity appears. The alarms indicated are like alarms described in the portlet Alarms Portlet on page 276.

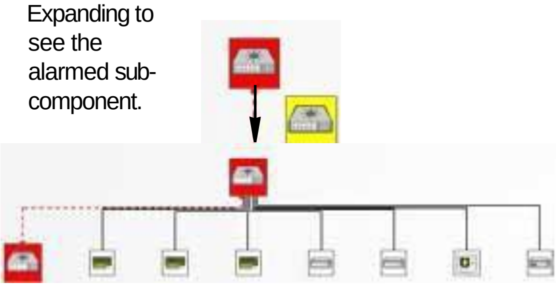

Expanding to see the alarmed sub-component.

Bydefault, un-alarmed nodesappearclear/white. Youcanalter this so theyappear greeninstead. To change this behavior, uncomment the following property:

```
nodes.display.clear.severity.as.green=true
```

Thisproperty is locatedintheserver-overrides. properties.sample filein the\oware\synergy\conf directory, and save the file as serveroverrides.properties in that directory.

Alarm Suppression in TopologyViews

Byadding icons to the devices, the topologyalso displays the alarm suppression and maintenance status of devices in the Topology view.

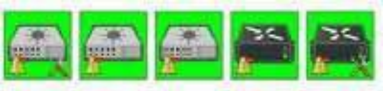

Here are the icons and their significance:

Icon	Device Status
	No icon—The device is unconstrained by the other Administrative States. Changing from Suspended to Normal stops alarm suppression. Standard access, and inclusion in right-to-manage count.
🔔	Alarm Suppression active— Activated from the Managed Resources, Event Management pop-up menu option.
🔧	Decommissioned — While this device is in inventory, it is not active. No device access allowed, no Monitor associations, no event processing, no Management Interfaces, no Authentication, no links, and no services are permitted.
✖	Down— The device is down.
🔧	Maintenance — Neither alarms or polling apply to the device. Does allow resync and Adaptive CLI. Standard device visibility.
📅	Planned — Planned (future) device. No device access allowed, no monitor associations, and no event processing.
🔧	Suspended — Suspends all device-related activities. No device access allowed, Monitoring Suspended, No event processing, Counts against right-to-manage.

You can set these alarm suppression and maintenance statuses in the right-click *Event Management* and *Maintenance* menus in the Managed Resources portlet.

Links in a Topology

When you discover links between devices in your network (see Link Discovery on page 188), they show in a topology view.

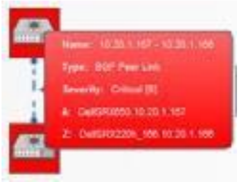

Name: 10.20.1.167 - 10.20.1.166
Type: BGP Peer Link
Severity: Critical [8]
A: DellSRX650.10.20.1.167
Z: DellSRX220h_166.10.20.1.166

Hover the cursor over a link, and a panel the color of the link's alarm severity, appears with the link information (*Name, Type* (for example: Ethernet) , *Severity*, and *A/Z Name*s for the endpoints).

Cruz currently does not support displaying one-ended links.

Topology Views

The Topology Views portlet displays saved views, and when it is on the same page as the System Topology portlet, filters that portlet so it displays the selected, saved topology.

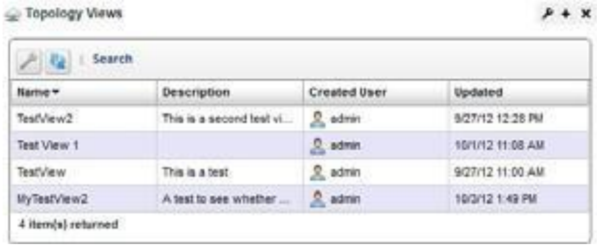

Topology Views

Name	Description	Created User	Updated
TestView2	This is a second test vi...	admin	9/27/12 12:28 PM
Test View 1		admin	10/1/12 11:08 AM
TestView	This is a test	admin	9/27/12 11:00 AM
MyTestView2	A test to see whether ...	admin	10/3/12 1:49 PM

4 item(s) returned

Right clicking selected views lets you *Edit* the title of the view, or its description, or *Delete* the selected view.

Understanding Hierarchical View

Hierarchical views manage what appears in other portlets on the same page, including Topology and Maps portlets. If a page with a Topology port let has no Hierarchical View portlet, clicking a topology node limits other portlets on that page to only that node's information (for example Alarms).

Here are the **context display rules** for how portlets manage each other's displays:

Rule	Description
1	If the Hierarchical View portlet is on a page, its selections drive all portlets that accept context. If you have no hierarchies configured, the other portlets are empty. If the Hierarchical View portlet is not on a page, then the remaining rules may apply.
2	If the System Topology portlet is on a page, it acts like the Hierarchical View portlet and drives all portlets' appearance.
3	If rules 1 and 2 **are not** in effect, the Managed Resources portlet drives Ports and Links portlets' appearance.
4	If rule 1 **is not** in effect and the Topology Views portlet is on the same page as the System Topology portlet, the selected view appears in the Topology portlet.

The context rule displays in a portlet's header if its appearance is being managed byanother portlet as defined by these context display rules.

When a page with a Hierarchical View loads, the hierarchy loads first and then starts polling. If a System Topology portlet is on the same page as a Hierarchical View portlet, the System Topology portlet starts its polling after the page loads, so some lag may occur between the Hierarchical View and Topology Views portlets, depending on your settings. Clicking Context from a portlet or drilling down or expanding nodes in the Topology portlet resets the refresh timer since it may poll different nodes. This can also offset refresh timing for different page elements. You can change refresh timing from the Application Settings window (see Redcell > Application Settings on page 44 for details), but synchronizing such portlets is not likely.

NOTE:

Someportletsmaydisplayaselectedcontextwithoutoperatingasthoughitwasselected.Forexample,if you put ManagedResources.

Some pages **may not** exhibit this default behavior. For example, custom branding on pages may interfere with these rules, in which case, you can see the default behavior by creating a new page with the relevant portlets. You may also try refreshing the page.

Youcan disable the context responses in the Alarms portlet from its preferences menu (click the wrench). When you disable context responses, the Alarmportlet instance's auto refresh displays new, unique alarms, but does not display, for example, a selected hierarchical view's alarms. Otherwise, you must change the context call from the Hierarchical View portlet to refresh alarms.

Using Hierarchies

Use hierarchies as follows.

1 Create the hierarchies you would like for filtering resource views.

For example, create a hierarchy for each customer or location.

△ CAUTION:

By default, hierarchies are configured without any authorizations. Make sure that you configure authorizations so you can see the hierarchy once it is configured. Otherwise, it is not visible.

a. Navigate to the Hierarchical View Manager portlet (Alarms/Events > Hierarchical View).

b. Right-click the portlet and then select New.

The Creating New Root Hierarchical View window is displayed.

c. Enter a name and description.

d. Specify membership, such as customers.

e. Set Authorizations.

f. Choose the topology display object.

g. Click Save.

2 Create a page with the Managed Resources portlet or other hierarchy-filtered portlets (Ports, Alarms, and so on).

3 Add the Hierarchical View portlet to that page.

4 Select the hierarchy by which to filter.

5 Observe the other portlets to see resources assigned to the selected hierarchy, such as, customer or location.

Setting Up GoogleMaps

Google now requires an API key to use their google maps API. Set up Google Maps as follows.

NOTE:

If you do not specify an API key, the default key is used but you may be severely limited in the number of map downloads you can make.

1 Go to https://console.developers.google.com .

2 Create a google account if you do not already have one.

3 Click the APIs & Services Library link.

4 Click Google Maps JavaScript API.

5 Click Create Project.

6 Click Create a Project under API Manager Dashboard.

7 Type in a project name and click yes to agree to terms of service.

8 Click Create.

9 Click Enable.

10 Click Create Credentials.

11 Under “Which API are you using?” select “Google Maps JavaScript API”.

12. Under “Where will you be calling the API from?” select Web browser (Javascript).
13. Click the “What Credentials do I need?” link.
14. Type a name for your API key.
15. Click Create API key.
16. Copy down the API key value.
17. Go to the Cruz Control Panel > Redcell > Application Settings.
18. Click the User Interface tab.
19. Select Google Maps as the map provider.
20. Enter the API key in the Application ID field.
21. Click Save.

Setting Up NokiaMaps

By default, the application uses Google maps. To use the Nokia maps service, you need App ID and App token. Set up Nokia maps service as follows.

1. Get an ID and token from https://developer.here.net .
2. Click Sign In.
3. Click on Register.
4. Create Your Nokia Account.
5. Click Register.
6. Click “Create app” and provide an application name.

   For example: Cruz Operations Center
7. Click Get Started.
8. Click Done.
9. Copy the App ID and App token.
10. Go to the Cruz Control Panel > Redcell > Application Settings.
11. Click the User Interface tab.
12. Select Nokia Maps as the map provider.
13. Enter the API key in the Application ID field.
14. Click Save.

Creating a TopologyView

Creating a topology map of devices or services is as simple as right clicking the items you want to map and selecting Topology. You can also save different topologies after configuration and fine-tune its appearance using the tools provided. Fora detailed description of the Virtualize portlet and its available options, see Topology Portlet on page 235.

⚠ CAUTION:

Ifyou installed a firewall onthe applicationserver, ports 80 and 8080 mustbe open for topology towork.

Create a topology map as follows.

1 Navigate to the Managed Resources portlet.

2 Select the items you want to include in the topology map.

3 Right-clickthe selectionsand thenselect Topology.

The Topology portlet is displayed.

Note that Topology uses Adobe Flash. The pop-up menu includes the Flash Settings, Global Settings, and About Flash menu items.

4 Organize the objects and draw the needed connectors.

5 Save your view.

Your view is now available from the Topology Views portlet.

6 Fine-tune your view using the Overview, Layout, and Propertiesoptions.

7 Save any changes.

If you do not see what you expect, make sure that you refresh your browser so cached images do not interfere with current images.

If you want to initiate Actions on a node or its components, right-click the node, select Details, right-click the reference tree node/component, and then select Actions.

ModifyingTopology Label Length

Labels default to 13 characters, but you can extend them to a wider size by adding the following property to the *installDir*/oware/synergy/conf/server-overrides.properties file:

```
nodes.labels.extended.width=true
```

Otherwise, you can force an extended label with your Extension, so it is automatic. Using your extension, you can do it within the `PortletProvider#getPortalProperties()` call.

For example:

```
public Properties getPortalProperties() {
   Properties props = new Properties();
   props.put("nodes.labels.extended.width", "true");
   return props;
}
```

NOTE:

This allows labels to be much longer but adds the possibility of text bleeding on top of other nodes.

Exporting to Visio

Export a topology view to Visio as follows.

1 Create a topology view if it does not already exist.

See Creating a Topology View on page 251 if you need some detailed steps.

2 Go to the Topology Views portlet.

3 Right-click the view you want to export and then select Export to Visio.

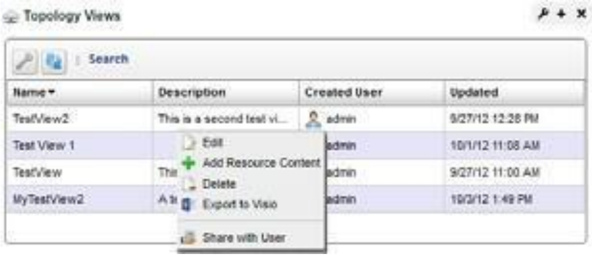

Topology Views

Search

Name	Description	Created User	Updated
TestView2	This is a second test vi...	admin	9/27/12 12:28 PM
Test View 1		admin	10/1/12 11:08 AM
TestView	Thi	admin	9/27/12 11:00 AM
MyTestView2	A t	admin	10/3/12 1:49 PM

A message displaysindicating that theviewwasexported.

4 Click Download Visio File.

This downloads the Visio file to your browser. The equipment and link tooltips are saved in the Visio shape data. To see the shape data in Visio, right-click the shape or link (connector) and then select Data > Shape Data.

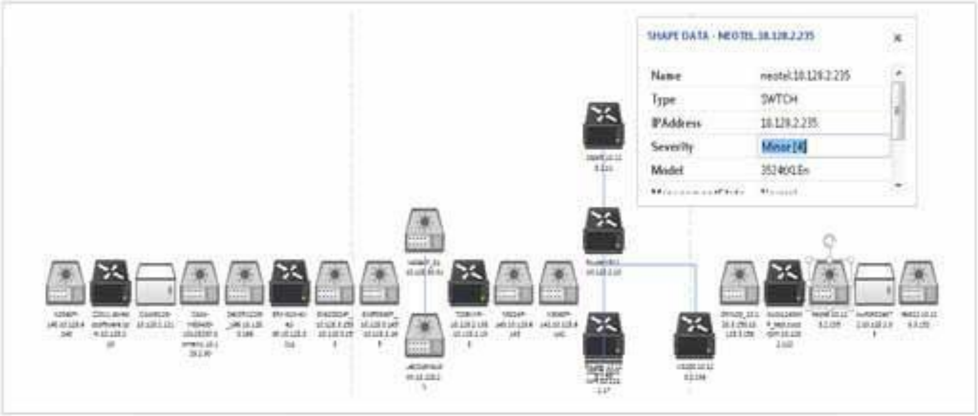

SHAPE DATA - NEOTEL.10.128.2.235

Name	neotel.10.128.2.235
Type	SWTCH
IPAddress	10.128.2.235
Severity	Minor [4]
Model	3524XL-En

5

Generating Reports

This section describes report-related portlets and editors, and tasks related to creating a report template, generating reports, and printing groups of reports. If you already have a good understanding of the portlets and editors, go directly to the tasks you want to perform.

Report Portlets/Editors

This section describes the following report-related portlets and editors:

- Report Templates Portlet
- Report Template Editors
- Reports Portlet
- Expanded Reports Portlet
- Report Editor

Report Templates Portlet

Report Templates are the basis of reports. This portlet displays the *Template Name*, *Description*, *Inventory Entity*, and *Type* in columns.

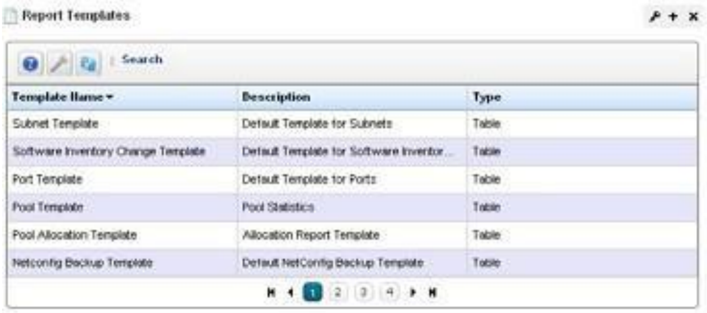

Template Name	Description	Type
Subnet Template	Default Template for Subnets	Table
Software Inventory Change Template	Default Template for Software Inventor...	Table
Port Template	Default Template for Ports	Table
Pool Template	Pool Statistics	Table
Pool Allocation Template	Allocation Report Template	Table
Netconfig Backup Template	Default NetConfig Backup Template	Table

Right clicking in this portlet lets you create a *New* template, *Edit* a selected template (see Report Template Editors for information about subsequent screens), view *Details* or *Delete* a selected template. You can also *Import/Export* report templates to files.

The expanded Report Templates portlet also includes a Reference Treesnap panel displaying a tree for selected templates connecting them to Report Groups and specific reports.

NOTE:

You can create reports related to users and their groups. Create a new report and select *Permission* as the Source of attributes in Report Template editor to begin.

Report Template Editors

Cruz has several report template editors. Creating a *New* template, can make *Comparison*, *Table* and *Trend* templates.

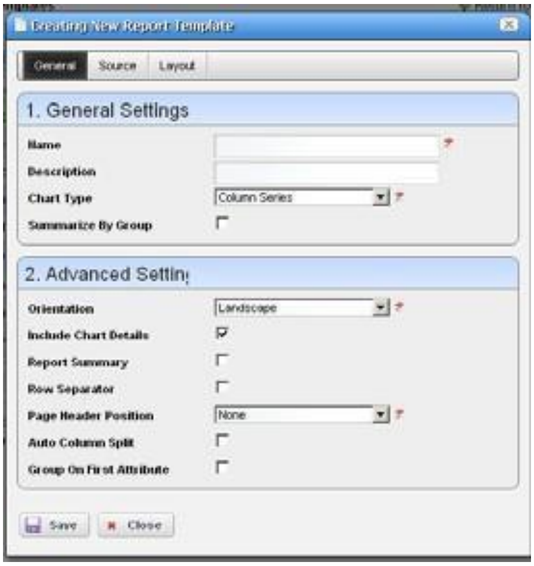

Creating New Report Template

1. General Settings

Name *

Description

Chart Type: Column Series *

Summarize By Group □

2. Advanced Settin

Orientation: Landscape *

Include Chart Details ☑

Report Summary □

Row Separator □

Page Header Position: None *

Auto Column Split □

Group On First Attribute □

Table reports simply report the configured data in tabular form as you have configured the columns. Comparison reportsdisplayselected attributes comparingreporting devices, forexample a summary graph then a list of devices’ ICMP monitor RTT in the following pages.

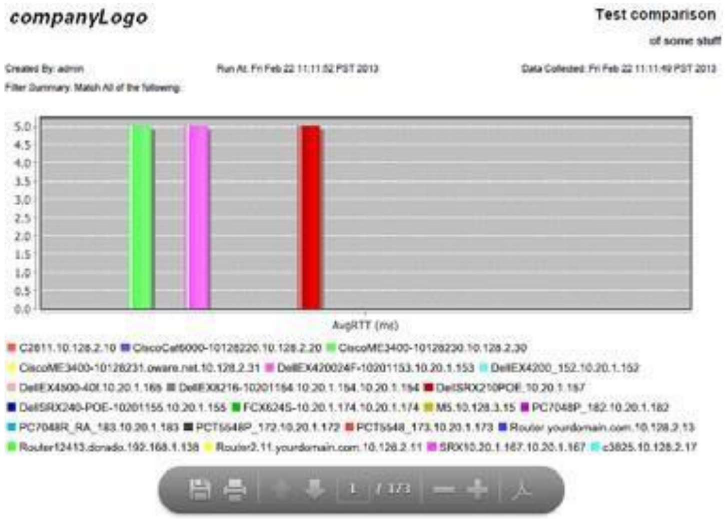

companyLogo

Test comparison

of some stuff

Created By: admin

Run At: Fri Feb 22 11:11:52 PST 2013

Data Collected: Fri Feb 22 11:11:49 PST 2013

Filter Summary: Match All of the following:

5.0
4.5
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
AvgRTT (ms)
C2811.10.128.2.10
CiscoCat6000-10128220.10.128.2.20
CiscoME3400-10128230.10.128.2.30
CiscoME3400-10128231.oware.net.10.128.2.31
DellEX420024F-10201153.10.20.1.153
DellEX4200_152.10.20.1.152
DellEX4500-40I.10.20.1.165
DellEX8216-10201164.10.20.1.164.10.20.1.164
DellSRX210POE.10.20.1.167
DellSRX240-POE-10201155.10.20.1.155
FCX624S-10.20.1.174.10.20.1.174
M5.10.128.3.15
PC7048P_182.10.20.1.182
PC7048R_RA_183.10.20.1.183
PCT5548P_172.10.20.1.172
PCT5548_173.10.20.1.173
Router.yourdomain.com.10.128.2.13
Router12413.dorado.192.168.1.138
Router2.11.yourdomain.com.10.128.2.11
SRX10.20.1.167.10.20.1.167
c3825.10.128.2.17
1
/ 173

A Trend Report displays a data graph with data reported over a polled period.

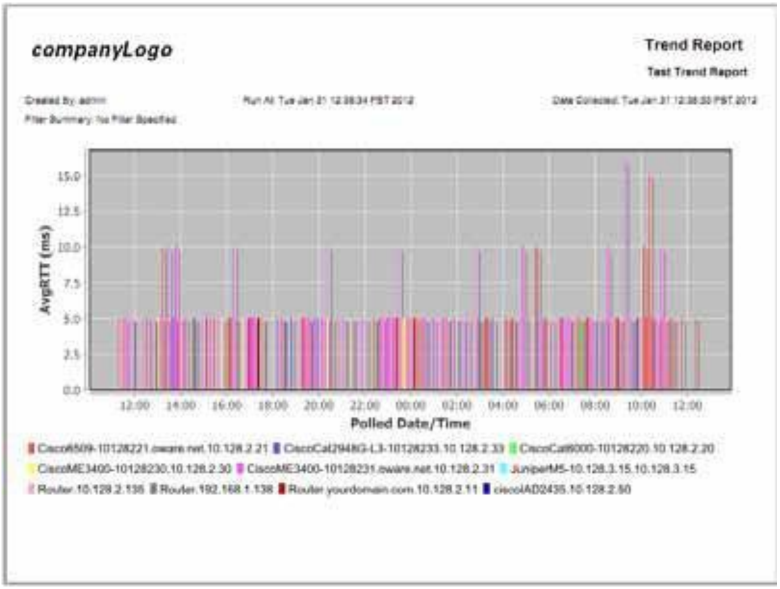

companyLogo

Trend Report

Test Trend Report

Created By: admin

Run At: Tue Jan 31 12:38:34 PST 2012

Data Collected: Tue Jan 31 12:38:30 PST 2012

Filter Summary: No Filter Specified

AvgRTT (ms)
15.0
12.5
10.0
7.5
5.0
2.5
0.0
12:00
14:00
16:00
18:00
20:00
22:00
00:00
02:00
04:00
06:00
08:00
10:00
12:00
Polled Date/Time
Cisco6509-10128221.oware.net.10.128.2.21
CiscoCat2948G-L3-10128233.10.128.2.33
CiscoCat6000-10128220.10.128.2.20
CiscoME3400-10128230.10.128.2.30
CiscoME3400-10128231.oware.net.10.128.2.31
JuniperM5-10.128.3.15.10.128.3.15
Router.10.128.2.136
Router.192.168.1.138
Router.yourdomain.com.10.128.2.11
ciscoIAD2435.10.128.2.50

You can now select more than one attribute for trend reports. Chart generation depends on the number of attributes selected and the number of targets:

- 1 target, n attributes produces 1 chart with all attributes (line series graph only)
- n targets, 1 attribute produces 1 chart with all the targets
- x targets, n attributes produces n charts with x targets on each

This editor has General, Source, and Layout panels.

Youcan edit any but pre-existing templates, whether they have reports attached to them or not. Consider this example:

Template T has three columns; A, B and C. Someone creates a report R against Template T, executes the report, saves the data as a historical report H1. Two weeks later, someone modifies the Template T, removing column C, adding column D.

When executing report R against the revised Template T’, the report now shows columns A, B and D. User saves the report as historical report H2. Here, H1 only has data for columns A, B and C. H2 has data for columns A, B and D.

If you view H1 you see Template T’ is in use and this template creates a report with columns A, B and D. Unfortunately, H1 only has data for columns A, B and C, so the report created has data for columns A and B only. Column D is empty. When viewing H2 you can see Template T’ is in use and can create a report with columns A, B and D. H2 has data for columns A, B and D, so all data appears.

General

The following are fields that appear on these screens. Not all screens have all fields.

General Settings

Name — An identifier for the template.

Description — An optional description of the template.

Chart Type—Select from the available alternatives (*column, line*). This is only available for trend and comparison templates.

Summarize by Group—Group similar results together. This is only available for trend and comparison templates

Advanced Settings

Orientation — Select from *Portrait* and *Landscape*

Include Chart Details—Includes a table with the data after the graph. This is only available for trend and comparison templates

Report Summary—Enables the report summary, which places the total count ofrecordsat the end of the report.

Row Separator — Displays a separator between rows within the report.

Page Header Position —Select *none, top, bottom* or *both.*

Auto Column Split —Enable automatic column splitting. This automatically aligns the columns equally on the report providing the column widths that are most proportional.

Group on First Attribute — Creates a report that groups rows based on the first reported attribute. This creates groups of items in the report whenever the left most column's value changes.

For example, with disabled, a report looks like this:

```
Device Name    Gig/e Port Name      Health Status
M5                         ge/0/0/1                 Up
M5                         ge/0/0/2                 Down
M5                         ge/0/0/3                 Up
M5                         ge/0/0/4                 Unknown
M18                        ge/0/1/1                Up
M18                       ge/0/1/2                 Starting
M18                       ge/0/1/3                 Up
M18                       ge/0/1/4                 Down
```

The same report looks like this with *Group on First Attribute* enabled:

```
Device Name    Gig/e Port Name      Health Status
M5
ge/0/0/1                  Up
ge/0/0/2                  Down
ge/0/0/3                  Up
ge/0/0/4                  Unknown

M18
ge/0/1/1                 Up
ge/0/1/2                  Starting
ge/0/1/3                  Up
ge/0/1/4                  Down
```

Alternativewaysof Grouping Attributes in Reports —The followingareways to turn ongrouping in a report template.

- *Select Group on First Attribute.* This groups output based only on the first attribute as described above.

- *Do not select Group on First Attribute* —In the layout for each column select *group by* for the individual attributes you wish to report together. This method creates separate groups for each attribute, groups within groups appear.

*Select both above.*This method creates a single group using all the columns you have selected in attribute layout and inserts a count for each group.

The Source and Layout tabs are common to all editors.

Summarize Data Only - Generates a report that only shows the raw data, with no column or page headers and no group summaries.

Source

Select the source inventory for a report, and its data types in this screen.

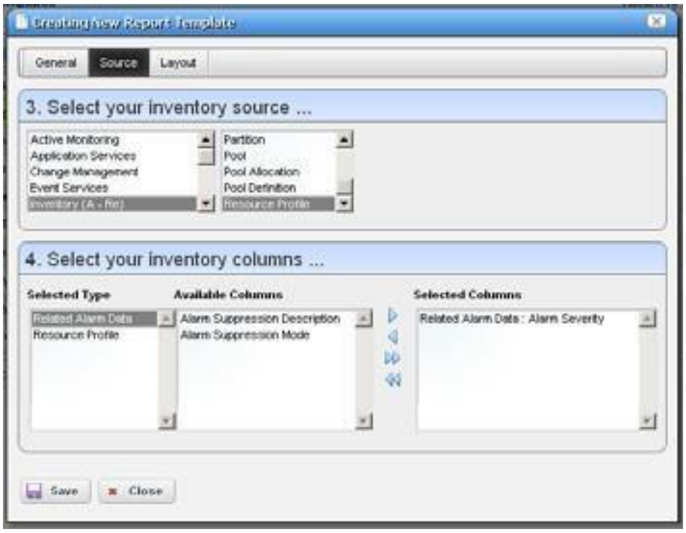

Click the green plus (+) to select the *Inventory Type.* The types of data available for that inventory type appear in the leftmost column in this screen. Click on a *Selected Type* to see its *Available Columns.* Click the arrows to move columns from *Available* to *Selected*. The *Selected Columns* appear in the template’s report.

Layout

This tab outlines the column layout for the template.

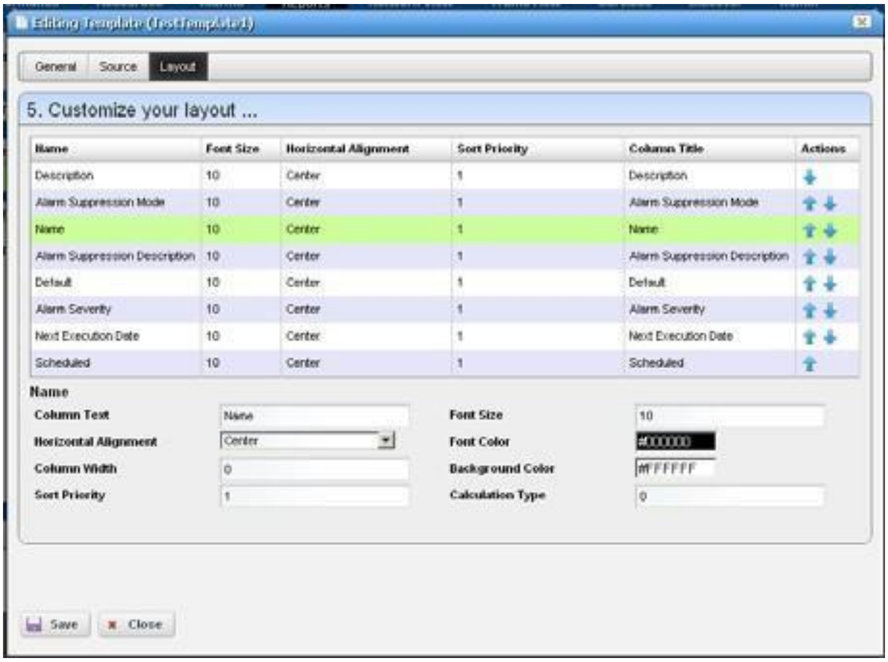

Editing Template (TestTemplate1)

5. Customize your layout ...

Name	Font Size	Horizontal Alignment	Sort Priority	Column Title	Actions
Description	10	Center	1	Description	↓
Alarm Suppression Mode	10	Center	1	Alarm Suppression Mode	↑ ↓
Name	10	Center	1	Name	↑ ↓
Alarm Suppression Description	10	Center	1	Alarm Suppression Description	↑ ↓
Default	10	Center	1	Default	↑ ↓
Alarm Severity	10	Center	1	Alarm Severity	↑ ↓
Next Execution Date	10	Center	1	Next Execution Date	↑ ↓
Scheduled	10	Center	1	Scheduled	↑

Name

Column Text	Name	Font Size	10
Horizontal Alignment	Center	Font Color	#000000
Column Width	0	Background Color	#FFFFFF
Sort Priority	1	Calculation Type	0

Click on the up/down arrows on the right of each row to re-order data columns. Click to select a row, and the editor panel at the bottom of the screen appears. It has the following fields:

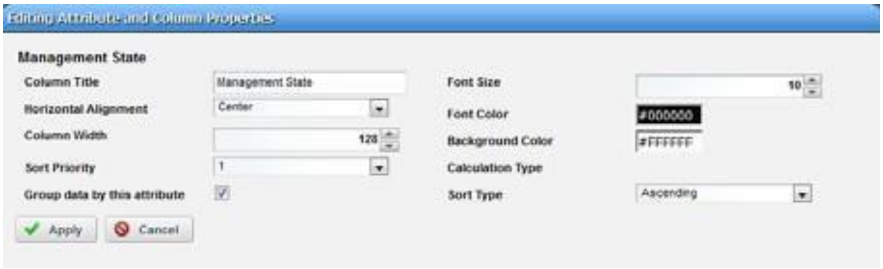

Editing Attribute and Column Properties

Management State

Column Title	Management State	Font Size	10
Horizontal Alignment	Center	Font Color	#000000
Column Width	128	Background Color	#FFFFFF
Sort Priority	1	Calculation Type	
Group data by this attribute	☑	Sort Type	Ascending

Column Text — The column label.

Horizontal Alignment — *Right*, *Left*, *Center* (the default).

Column Width — The column width in characters.

Sort Priority —Configures report sorting. Define the attribute sort order here. Youcan sort within a sort, so you can sort on Name and then by Location and then by IP Address, and so on. The number configures the sort group, so 1 sorts, then 2 within 1, then 3, and so on.

Font Size — The data's font size.

Font/Background Color —The color for the text/background. Click the field to open a color chooser.

Calculation Type—How to calculate for summarizing the numeric data. Select from the available options, which includes *Average*, *Min*, *Max*, *Sum*, and Min/Avg/Max (which shows each of these together). If Min/Avg/Max is selected for at least one attribute areport summary header will be created at the top of the report showing the average max and min values for the attribute calculated over all the data rows in the report. The following screenshot shows an example report summary:

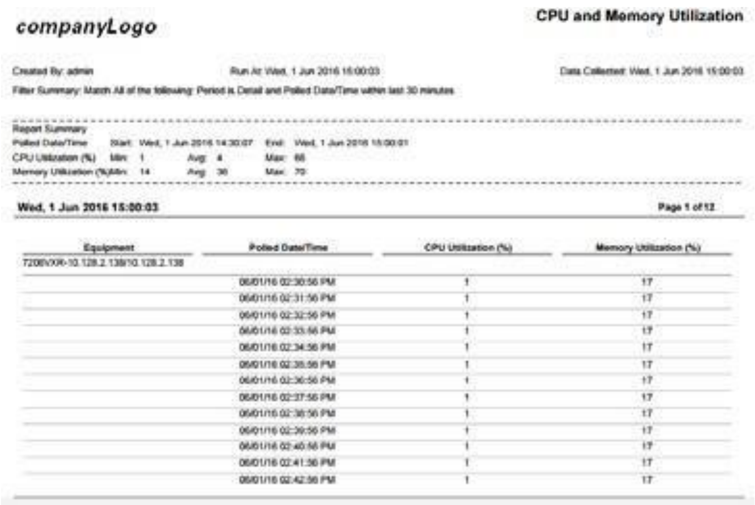

companyLogo

CPU and Memory Utilization

Created By: admin   Run At: Wed, 1 Jun 2016 15:00:03   Data Collected: Wed, 1 Jun 2016 15:00:03

Filter Summary: Match All of the following: Period is Detail and Polled Date/Time within last 30 minutes

Report Summary

Polled Date/Time	Start:	Wed, 1 Jun 2016 14:30:07		End:	Wed, 1 Jun 2016 15:00:01	
CPU Utilization (%)	Min:	1	Avg: 4	Max:	66	
Memory Utilization (%)	Min:	14	Avg: 36	Max:	70	

Wed, 1 Jun 2016 15:00:03   Page 1 of 12

Equipment	Polled Date/Time	CPU Utilization (%)	Memory Utilization (%)
7206VXR-10.128.2.138/10.128.2.138			
	06/01/16 02:30:56 PM	1	17
	06/01/16 02:31:56 PM	1	17
	06/01/16 02:32:56 PM	1	17
	06/01/16 02:33:56 PM	1	17
	06/01/16 02:34:56 PM	1	17
	06/01/16 02:35:56 PM	1	17
	06/01/16 02:36:56 PM	1	17
	06/01/16 02:37:56 PM	1	17
	06/01/16 02:38:56 PM	1	17
	06/01/16 02:39:56 PM	1	17
	06/01/16 02:40:56 PM	1	17
	06/01/16 02:41:56 PM	1	17
	06/01/16 02:42:56 PM	1	17

Groupdata by thisattribute — Ifchecked, the data willbe grouped bythis attribute, which means that every distinct value for this attribute will only appear once as a group header and the related data for the other attributes will show rows below this header.

Sort Type — Select ascending or descending.

Click Save to preserve any template you have configured, or Close to close the editor screens without saving.

Reports Portlet

This portlet's summary screen lists the available reports that you can run with Cruz.

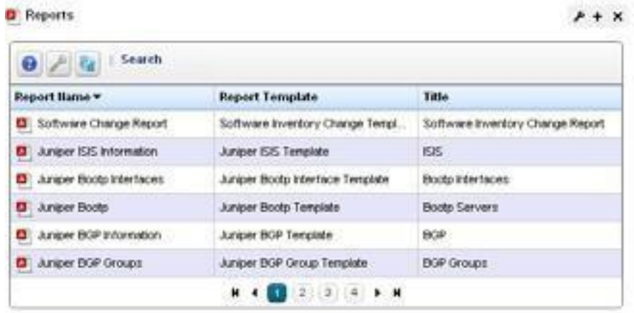

Report Name ▾	Report Template	Title
Software Change Report	Software Inventory Change Templ...	Software Inventory Change Report
Juniper ISIS Information	Juniper ISIS Template	ISIS
Juniper Bootp Interfaces	Juniper Bootp Interface Template	Bootp Interfaces
Juniper Bootp	Juniper Bootp Template	Bootp Servers
Juniper BGP Information	Juniper BGP Template	BGP
Juniper BGP Groups	Juniper BGP Group Template	BGP Groups

The report *Icon*, *Name*, *Template*, and *Subtitle* appear in the columns in this summary screen. Thereportselects the targetequipment, and the template configures the layout and attributes reported. If the Interface details panel is empty, then the Interface reports will have no contents. Some devices have ports, but no interfaces. Use the Ports report for such devices.

Cruz generates reports with only the first 5,000 records by default. Larger reports warn that they have reached the maximum and have only those first 5,000records.

You can change the maximum with the property `com.dorado.redcell.reports.max.report.query.size=5000` in rpt.properties file in /owareapps/reports/lib/.

Larger numbers have an impact on the performance of the report and database.

NOTE:

You must have Adobe’s Acrobat reader installed to view reports.

Right-click a selected report to do the following:

New/Edit/Copy — This opens the Report Editor, described below, to configure a new report, edit or copy an existing, selected report. *Copy* automatically renames the selected report.

New Group — Creates a collection of reports. See Printing Groups of Reports on page 269 for details about how to configure these.

Schedule — Opens a scheduler screen to automate report creation.

Execute Report— When you execute a report, a numbered message notification appears, and a link to the report appears in the *Messages* panel to notify you the report is ready for viewing. Click the magnifying glass to the right of the notification to view either the audit trail or the report.

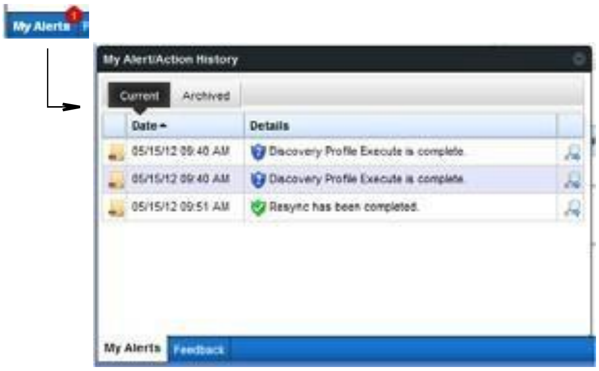

Date	Details
05/15/12 09:40 AM	Discovery Profile Execute is complete.
05/15/12 09:40 AM	Discovery Profile Execute is complete.
05/15/12 09:51 AM	Resync has been completed.

Lengthy Reports may takea sometime to appear onscreen without much indication that they are in process. This is an artifact of the Acrobat plug-in, and outside the scope of Cruz to influence. Acrobat also produces an error if a report has too much data to display meaningfully.

Execute Report (Advanced)—Also lets you schedule configure a fewother things with reports.

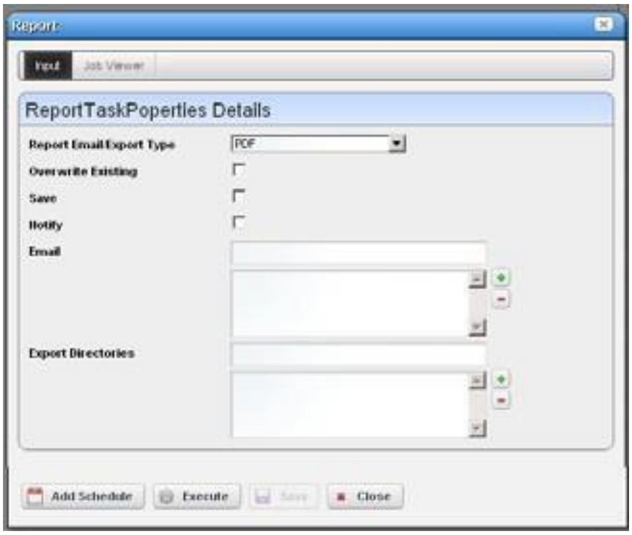

Report Email Export Type: PDF

When you *View or Execute Report (Advance)*, by right clicking either a listed report or a historical instance of that report, a configurat ion screen appears that lets you select several parameters.

These include the following:

Report Email/Export Type —Select the export file type from the pick list. Options can include *CSV*, *HTML*, and *PDF*.

NOTE:

ProgramsotherthanCruzletyoumanipulate mailoutsidethescope ofCruz.Forexample, IFTTT(If This ThenThat) couldsavemail attachments likereports toDropbox accounts. Also: Open CSV outputina spreadsheet for additional formatting options.

Overwrite Existing — Check to activate overwriting any existing report.

Save— Check to activate saving the report to the database.

Notify — Check to activate emitting a notification event.

Email Address — Enter an e-mail destination for the generated report, and click the plus (+) to list it. You can enter several such e-mails.

Export Directory — Enter directory destinations for saved reports as you would e-mail destinations.x

Click *Add Schedule* to schedule the report for future or repeated execution, *execute* to run the report immediately, or *Save* to preserve this report’s configuration. The *Job Viewer* tab displays the report’s progress if you click *Execute*.

CAUTION:

Reports can be large. Typically, the limitations on e-mail within your system are what limit the size of deliverable reports. Bestpractice is to use filters and a limited number of targets to make reports succinct rather thancomprehensive.

Aging Policy—If you automate report generation, you may also want to configure a Database Aging Policy to ensure the volume of reports does not overwhelm your storage capacity. See Implementing DAP on page 80 for more about doing that.

Delete — Removes the selected report from the list display

Delete History — Removes the selected report's history.

Tochange reports' appearance and contents, you must configure their Report Templates Portlet. Also, see Branding Reports on page 268 for instructions about changing the default report logo.

Expanded Reports Portlet

Clicking the plus (+) icon displays the expanded portlet. the expanded portlet adds *Add/Remove Column* to the menu options available in the summary screen.

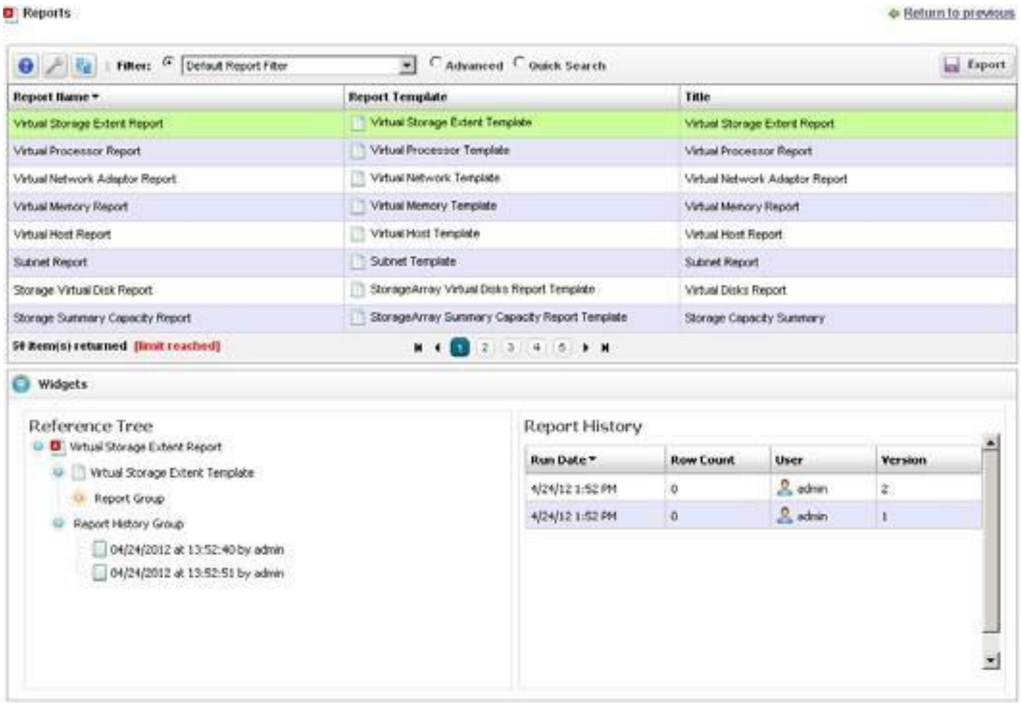

Report Name ▾	Report Template	Title
Virtual Storage Extent Report	Virtual Storage Extent Template	Virtual Storage Extent Report
Virtual Processor Report	Virtual Processor Template	Virtual Processor Report
Virtual Network Adaptor Report	Virtual Network Template	Virtual Network Adaptor Report
Virtual Memory Report	Virtual Memory Template	Virtual Memory Report
Virtual Host Report	Virtual Host Template	Virtual Host Report
Subnet Report	Subnet Template	Subnet Report
Storage Virtual Disk Report	StorageArray Virtual Disks Report Template	Virtual Disks Report
Storage Summary Capacity Report	StorageArray Summary Capacity Report Template	Storage Capacity Summary

59 item(s) returned (limit reached)

Reference Tree

- Virtual Storage Extent Report
  - Virtual Storage Extent Template
    - Report Group
  - Report History Group
    - 04/24/2012 at 13:52:40 by admin
    - 04/24/2012 at 13:52:51 by admin

Report History

Run Date ▾	Row Count	User	Version
4/24/12 1:52 PM	0	admin	2
4/24/12 1:52 PM	0	admin	1

Availablecolumns are like thosedescribedin the summaryscreensection previously. The *Reference Tree* snap panel displays the selected report's connection to devices, historical reports and any report template. Right-click to view the reports in the Historical Reports node.

The Widgetspanels for reportsdisplaya Reference Treeofconnections between the selected report and target equipment, and between the report and any Report Template.

The *Report History* Snap Panel displays the selected report's *Run Date*, *Row Count* and the *User* who ran the report. Right-click a row in this panel, and you can *Delete*, *Print* (the report history) or *Export* (the report history), *View* (the report) or *View (Advanced)*. If you *View* the report, a message with a link to the report appears in the bottom left of the screen.

Report Editor

This editor configures reports, and their targets. It has the following panels:

- General
- Filter

General

This screen configures the *Name, Title* (displayed text in the report), *Subtitle*, and lets you select the *Report Template* for the report (see Report Templates Portlet on page 254 for more about them)

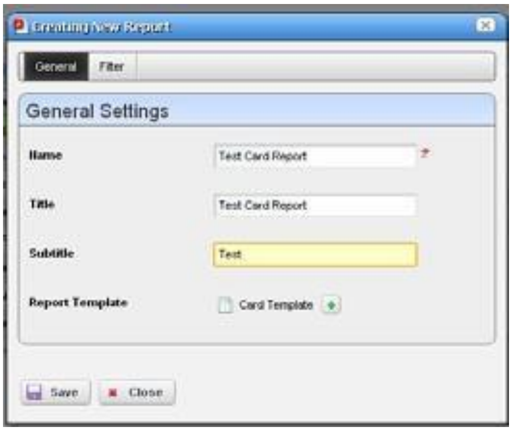

General Settings

Name: Test Card Report

Title: Test Card Report

Subtitle: Test

Report Template: Card Template

Filter

This screen configures a filter to retrieve devices that are the source of the report.

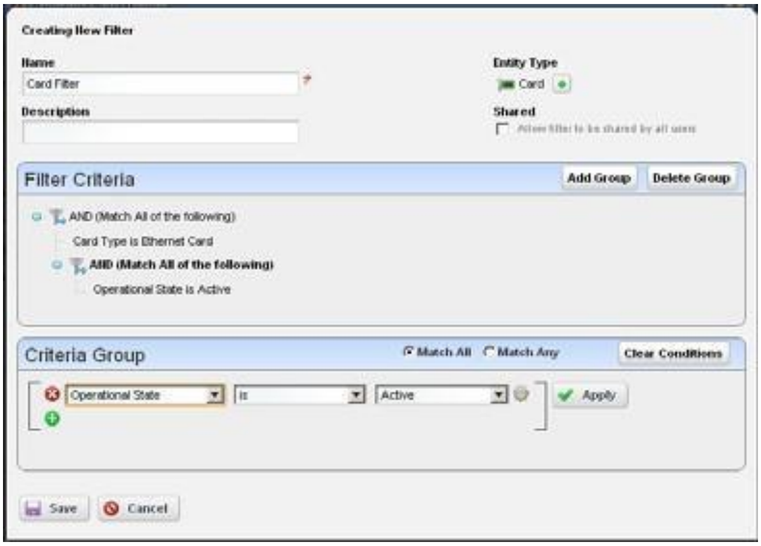

Creating New Filter

Name: Card Filter

Description:

Entity Type: Card

Shared: □ Allow filter to be shared by all users

Filter Criteria

- AND (Match All of the following)
  - Card Type is Ethernet Card
  - AND (Match All of the following)
    - Operational State is Active

Criteria Group ● Match All ○ Match Any

Operational State | is | Active

Click *Add Filter* in the filter panel to select an existing filter, create a new filter, or copy an existing filter. When you create a new filter, you must enter a *Name* and optionally a *Description* for it, select an *Entity Type* with the green plus (+), and elect whether this filter is available to other users *(Shared)*. See Defining Advanced Filters on page 145 for instructions about configuring the filter itself in the lower portion of this screen.

Once you have configured or selected a filter, the *Filter* panel displays its characteristics in tree form. Click *Edit* to re-open the editor, or *Del* to remove the filter. Filters appear only from the entity type of your report template.

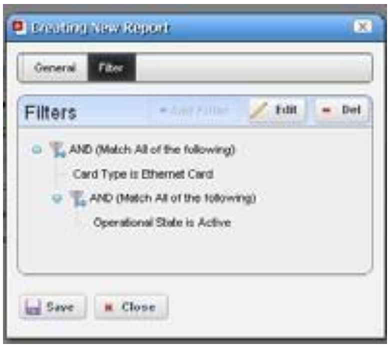

AND (Match All of the following)
Card Type is Ethernet Card
AND (Match All of the following)
Operational State is Active

Creating a Report Template

Formatting counts in making reports useful. Sometimes the output limitations need to inform the formatting you select. For example, PDF output does not handle large numbers of columns well, while CSV (importable into Excel) output has no problem with it. Best practice is to test reports you configure before putting them into production. For a detailed description of the Report Templates portlet and related editors, see Report Templates Portlet on page 254.

Create a report template from the Report Templates portlet as follows.

1 Right-click the portlet and then select New > *template Type.*

The Creating New Report Template window is displayed.

2 Name the template (for example: Test Amigopod Report).

3 Optionally enter a detailed description for the template.

4 Make sure that the advanced settings are correct, such as orientation.

The defaults are Landscape.

5 Click the Source tab.

The Source panel displays a list of inventory sources.

6 Select an inventorysource (for example: Inventoryresources [A - DD] Amigopod).

A list of values related to the selected source is displayed.

7 Select the appropriate value.

The Inventory columns lists are populated.

8 Select a type.

9 Populate the selected columns list from the available columns list (for example: Amigopod: Administrative State, Amigopod: DNS Hostname, Amigopod: Equipment Name, Amigopod:IP Address)

If you populate the selected columns list in the order you want them to appear in the report, they automatically appear in this order on the Layout panel and you can skip to step 12. Otherwise, continue with step 10.

10 Click the *Layout* tab.

The Layout panel lists the column order (top is first, bottom is last) and the default column settings.

11 Click the Edit action tool to modify a column's the font size, color, alignment, and so on and then Apply your changes.

12 Click *Save.*

You have successfully created a template.

Generating a Report

To generate a report, you first configure the report and then generate it from the Reports portlet. For a detailed description of the Reports portlet and related editors, see Reports Portlet on page 260.

NOTE:

If you create a report based on interface monitoring, remember, the interface monitor is disabled as default. Additionally, some reports rely on data collected through Actions and are empty unless the Action supplying data is executed first. For example, to generate a report using VLAN data, you need to execute the “Get VLAN Data” action first. Without data, the report is empty.

For an example of a standard system report, see User Login Report on page 272.

Generate a report from the Reports portlet as follows.

1 Right-click a report and then select New.

The Creating New Report window is displayed.

2 Name the report (for example: Test Juniper Router Report).

3 Enter a title and/or subtitle for the report (such as Juniper Routers).

4 Select a template for the report.

For example, the template configured in Creating a Report Template on page 265.

Note that if you create a template, the first report you create after making that template automatically selects the newly created template.

5 Click the Filters tab.

6 Add a a filter to confine the reports input to certain devices, locations, and so on.

For example, select the existing All Juniper Routersfilter.

7 Click Save.

8 Locate the newly created report in the Reports portlet.

9 Right-click the report and then select Execute Report.

A message of success or failure is displayed.

10 Click My Alerts from the portal status bar.

The My Alert/Action History window displays a notification that the report is ready for viewing.

11 Click the magnifying glass tool for the report message.

The report displays in a new window.

12 Hover your cursor over the lower right corner of the report to show tools that let you expand, zoom out and in, save, or print the report.

Branding Reports

Reports come with a default logo, but you can change that. Put the graphic file (.png, .jpg or .gif) with your desired logo in a directory on the applicationserver.

△ CAUTION:
If you have a distributed installation, make sure this image and property are on all servers.

Brand reports with a different logo as follows.

1 Create an image that is no taller than 50 pixels, and no wider than 50 pixels.

2 Save the image file as .png, .jpg, or .gif.

3 Navigate to the installed.properties file located in the owareapps/installprops/lib/ directory.

4 Alter the image property in the installed.properties file.

```
redcell.report.branding.image=<filename_here>
```

For example:

```
redcell.report.branding.image=C:/installPath/owareapps/redcell/images/
  TestImage.png
```

Notice that you **must** use the forward slashes, (not backslashesas is typicalof Windows) when you specify the path.

5 Verify that the logo changes by generating a report .

Adding Custom Report Images

Youcan upload images that appear, branding reports into the portal’s Documents and Media portlet. Put them in the pre-seeded folder named *Report Images*. By default, reports use the site logo in reports if you have set up no override.

Site Override—To override the logo for every report that does not have an individual custom logo setup, upload the image into the Documents and Media/Reports Images folder and name it report logo. This image becomes the site's default reportlogo.

Individual Report Override —Individual reports. To override the logo for individual reports, upload an image into the Documents and Media/ Reports Images folder. It then appears in the report editor as a selectable branding image. Users can view and select these images if they have the correct permission with write privileges (Permission-RP:ReportTitleImage).

Printing Groups of Reports

You can print a collection of several reports and gene rate a table of contents if the collection is large enough to warrant it. Print groups of reports from the Reports portlet as follows.

1 Right-click and then select New Group.

The Group Report editor is displayed.

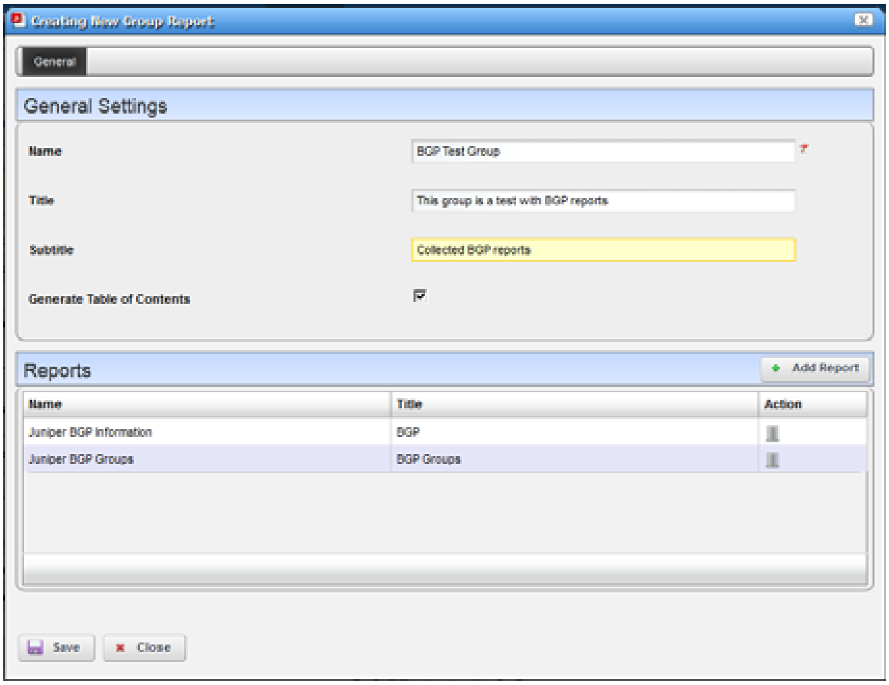

General Settings

Name	BGP Test Group
Title	This group is a test with BGP reports
Subtitle	Collected BGP reports
Generate Table of Contents	☑

Reports

Name	Title	Action
Juniper BGP Information	BGP	🗑
Juniper BGP Groups	BGP Groups	🗑

2 Enter the name and optionally a title and subtitle.

NOTE:

Reports within the report group are sorted by title. If no title specified, the report name is used as the map key.

3 Select whether to generate a table of contents.

The generated table of contents lists reports by group.

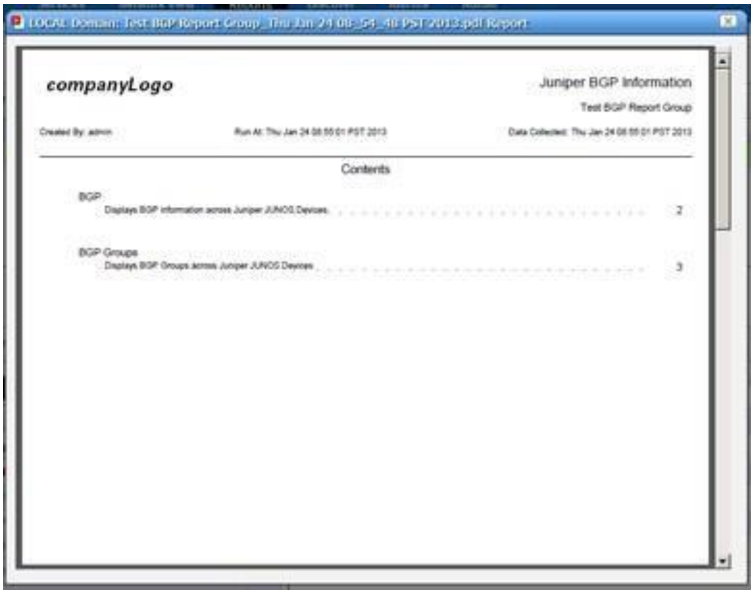

LOCAL Domain: Test BGP Report Group_Thu Jan 24 08_54_48 PST 2013.pdf Report

companyLogo

Juniper BGP Information

Test BGP Report Group

Created By: admin

Run At: Thu Jan 24 08:55:01 PST 2013

Data Collected: Thu Jan 24 08:55:01 PST 2013

Contents

4 Click Add Report(+).

A list of reports from which to select is displayed.

5 Select all that apply.

6 Click Add Selection.

7 Click Done.

8 Click Save.

Example Reports

This section provides the following example reports:

- Cisco Port Groups
- User Login Report
- Network Assessor Reports

Cisco Port Groups

The Cisco Port Group Report collects data on port groups and assembles that to display the current total bandwidth for port groups under a device or card.

This supports two types of port groups, both with 8 ports:

- One type is alternating groups of 8, so on a 48-port card, 1-8 would be one port group, 9-16 would be a second, 17-25 would be a third, and so on.
- The other type is even/odd alternating groups of 8, so on a 48 port card, ports 1, 3, 5, 7, 9, 11, 13, and 15 would be one group, while ports 2, 4, 6, 8, 10, 12, 14 and 16 would be a second group, 17, 19, 21, 23, 25, 27, 29, 31 would be a third group, and so on.

The following two properties in the owareapps/c isco/lib/cisco.properties file configure port collection:

```
com.dorado.cisco.portgroups.blockcards
```

```
com.dorado.cisco.portgroups.evensoddscards
```

Tou must add any card for which you want to gather port group data to one of these two lists, depending on what kind of port groups it wants to collect. Add the card by adding its MODEL NAME(example: WS-X6248-RJ-45). Delimit multiple modelswithcommas. Examplesareinthe cisco.properties file comments.

User Login Report

In addition to reports about inventory, devices, and so on, the Cruz application lets you create a report documenting user logins.

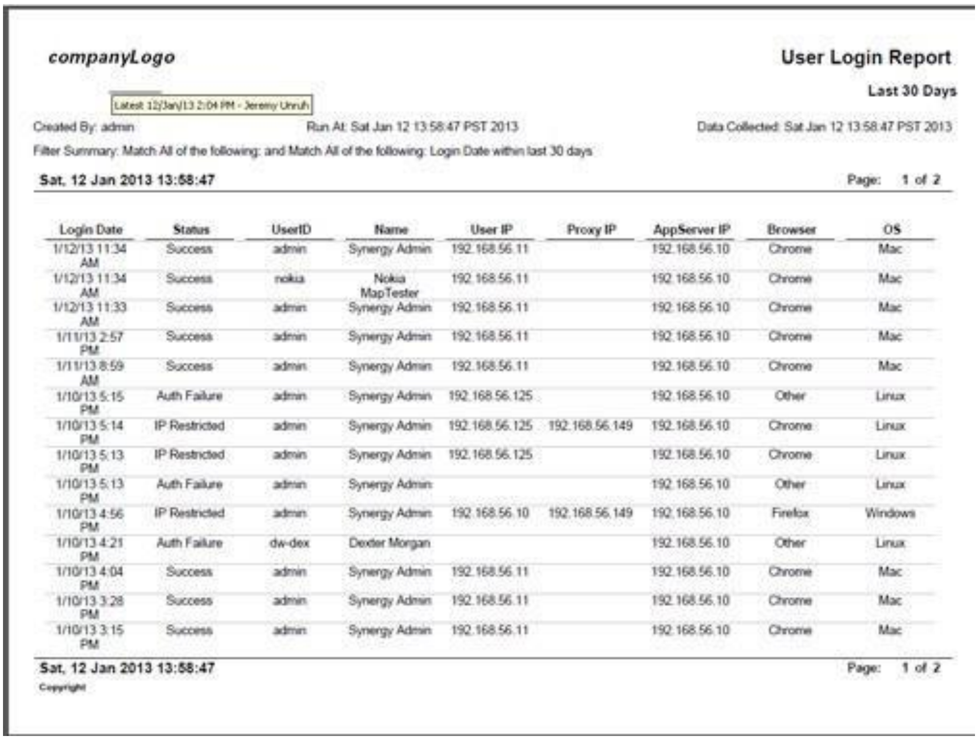

companyLogo

User Login Report

Last 30 Days

Latest 12/Jan/13 2:04 PM - Jeremy Unruh

Created By: admin　Run At: Sat Jan 12 13:58:47 PST 2013　Data Collected: Sat Jan 12 13:58:47 PST 2013

Filter Summary: Match All of the following: and Match All of the following: Login Date within last 30 days

Sat, 12 Jan 2013 13:58:47　Page: 1 of 2

Login Date	Status	UserID	Name	User IP	Proxy IP	AppServer IP	Browser	OS
1/12/13 11:34 AM	Success	admin	Synergy Admin	192.168.56.11		192.168.56.10	Chrome	Mac
1/12/13 11:34 AM	Success	nokia	Nokia MapTester	192.168.56.11		192.168.56.10	Chrome	Mac
1/12/13 11:33 AM	Success	admin	Synergy Admin	192.168.56.11		192.168.56.10	Chrome	Mac
1/11/13 2:57 PM	Success	admin	Synergy Admin	192.168.56.11		192.168.56.10	Chrome	Mac
1/11/13 8:59 AM	Success	admin	Synergy Admin	192.168.56.11		192.168.56.10	Chrome	Mac
1/10/13 5:15 PM	Auth Failure	admin	Synergy Admin	192.168.56.125		192.168.56.10	Other	Linux
1/10/13 5:14 PM	IP Restricted	admin	Synergy Admin	192.168.56.125	192.168.56.149	192.168.56.10	Chrome	Linux
1/10/13 5:13 PM	IP Restricted	admin	Synergy Admin	192.168.56.125		192.168.56.10	Chrome	Linux
1/10/13 5:13 PM	Auth Failure	admin	Synergy Admin			192.168.56.10	Other	Linux
1/10/13 4:56 PM	IP Restricted	admin	Synergy Admin	192.168.56.10	192.168.56.149	192.168.56.10	Firefox	Windows
1/10/13 4:21 PM	Auth Failure	dw-dex	Dexter Morgan			192.168.56.10	Other	Linux
1/10/13 4:04 PM	Success	admin	Synergy Admin	192.168.56.11		192.168.56.10	Chrome	Mac
1/10/13 3:28 PM	Success	admin	Synergy Admin	192.168.56.11		192.168.56.10	Chrome	Mac
1/10/13 3:15 PM	Success	admin	Synergy Admin	192.168.56.11		192.168.56.10	Chrome	Mac

Sat, 12 Jan 2013 13:58:47　Page: 1 of 2

Copyright

Thisreportcanincludethefollowingattributes: Login Date, Status[SUCCESS, AUTHFAILURE, IP RESTRICTION], UserID, User Name, User IP, Proxy IP (if going through a Load Balancer/ Proxy), AppServer IP, Browser [CHROME, FIREFOX, and so on], Operating System [WINDOWS, LINUX, MAC, IPHONE, IPAD, and so on.]

The following attributes are available, but not in the default seeded report to conserve Column space: Portal IP and Browser Version.

When authentication fails, this report does not record the IP address from which the user made the attempt, unless such users are behind a proxy or loadbalancer.

Browser ID, Version and Client appear only by best effort. Browsers do not always send the user-agent and can change standard messaging with extraplugins.

NOTE:

A Default User Sign-On Log DAP exists, which by default keeps the last 30 days.

Network Assessor Reports

If you have the Network Assessor option (assessor.ocp) installed, it includes reports like the following: *Asset Report ALL Resources, Software Version Report, IP - Hostname Only Report, IP - Hostname Report, Configuration Change Report , Hardware Change Report, Software Change Report, NetConfig Backup Status Report, NetConfig Deploy Stat us Report, NetConfig Restore Status Report, Card Report, Firmware Report, Interfaces Report, Inventory Report, Port Report, Primary Contact Report, Subnet Report.* You can also see an *EOL Report ALL Devices* (EOL means "End of Life").

The *EOL Report ALL Devices* report tells which of your discovered equipment has passed its end of life or end of service. A registration script (RegisterEOL) in the **.../owareapps/assessor/bin** directoryregisters EOLinformationdifferent fromthedefaultsthatship with the Network Assessor option. To update your EOL/EOS ("End of Life"/ "End of Service") dates, create a text file (myEOL.txt) with EOL and EOS definitions as a parameter for RegisterEOL. Construct this parameter file as follows:

EOL=SysobjectID, EOL True False, EOL Date (mm/dd/yyyy), EOS True False, EOS Date (mm/dd/yyyy)

Here is an EOL and EOS example text file:

```
# Example (below) Cisco AS5200 Series Universal Access Servers AS5200,,,,
EOL=1.3.6.1.4.1.9.1.109,True,07/14/1999,True,07/23/2010
```

where:

1.3.6.1.4.1.9.1.109 is the SysobjectID for a Cisco AS5200 Universal Access Server. Remember to preceed the objectID with EOL =

True sets the End of Life indicator to true

07/14/1999 is the End of Life date

True sets the End of Service indicator to true

07/23/2010 is the End of Service date

Once you created your EOL/EOFS text file (myEOL.txt), follow these steps to register the file.

1. Save the myEOL.txt file to *installDir* /owareapps/assessor/bin directory.
2. Navigate to *installDir*/owareapps/assessor/bin directory.
3. Execute the RegisterEOL script while the server is running.
   - In Linux, run these commands:
     ```
     . /etc/.dsienv
     ./RegisterEOL myEOL.txt
     ```
   - In Windows, run these commands:
     ```
     oware
     RegisterEOL myEOL.txt
     ```

   The "New EOL Definition was processed" message is displayed.

6

Alarms, Events, and Automation

This section provides tasks and information related to setting up alarms, events, and automation. It also provides what you need to understand alarm propagation and event/alarm life cycles. If you already have a good understanding of the portlets and editors, go directly to the tasks you want to perform.

Alarms

This section describes the following portlets and editors related to alarms, events, and automation:

- Alarms Portlet
- Alarm Editor
- Setting Audible Alerts

Alarms Portlet

The Alarms portlet on the Home page displays the pie chart and the color legend.

Bydefault, access the portlet byselecting Home, Alarms/Events, or Alarms/Events > Hierarchical View, and Topology > Hierarchical View from the navigationbar.

In its summary form, this portlet displays alarms. See General > Entity Change Settings on page 44 for the way to set the summary portlet refresh interval. The default is 40 seconds. If this portlet is on the same page as the Hierarchical View portlet, or if it is in expanded mode, refresh does not occur automatically, but you can refresh it manually.

A small clock icon appears in the upper right corner of this portlet if auto-refresh is enabled. A small speaker icon appears if audible alerts are enabled. See Using Extended Event Definitions on page 333 for more about this feature.

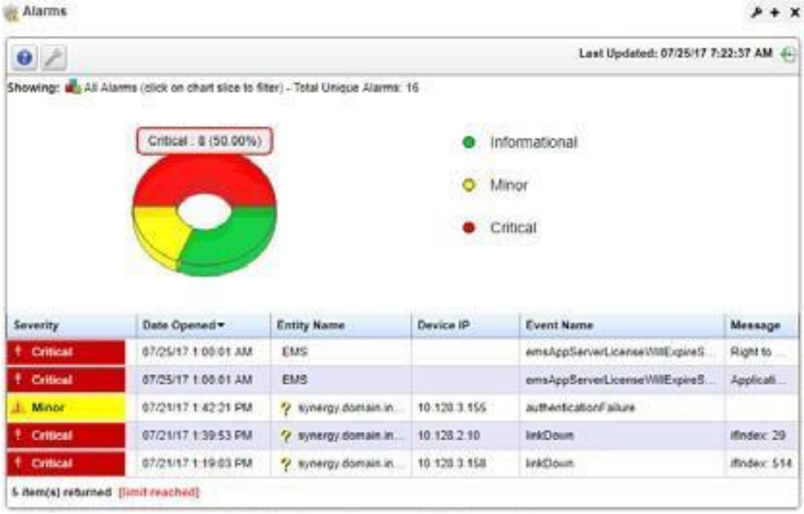

Showing: All Alarms (click on chart slice to filter) - Total Unique Alarms: 16

Severity	Date Opened	Entity Name	Device IP	Event Name	Message
Critical	07/25/17 1:00:01 AM	EMS		emsAppServerLicenseWillExpireS...	Right to ...
Critical	07/25/17 1:00:01 AM	EMS		emsAppServerLicenseWillExpireS...	Applicati...
Minor	07/21/17 1:42:21 PM	synergy.domain.in...	10.128.3.155	authenticationFailure	
Critical	07/21/17 1:39:53 PM	synergy.domain.in...	10.128.2.10	linkDown	ifIndex: 29
Critical	07/21/17 1:19:03 PM	synergy.domain.in...	10.128.3.158	linkDown	ifIndex: 514

5 item(s) returned [limit reached]

The chart acts as a filter, too. For example, clicking the *Critical* alarms slice means only *Critical* alarms are listed. The chart explodes to highlight the selected slice. Hover the cursor over a portion of the chart and a tooltip with information about that slice also appears. Click exploded slices to return the graph and the list to its originalstate.

Different tooltips appear when you hover over the Entity Name and Device IP columns. Such tooltips are available if the question mark appears when you hover over the field.

Forexample, hover over an alarm's entity name and the tooltip displays the alarm's Date Opened, the Entity Name, any alarm Message, Event Name, Alarm and Entity Type, its status as Service Affecting, Notification OID, Equipment, Severity, whether the alarm was Suppressed, or

Acknowledged and the Device IP.

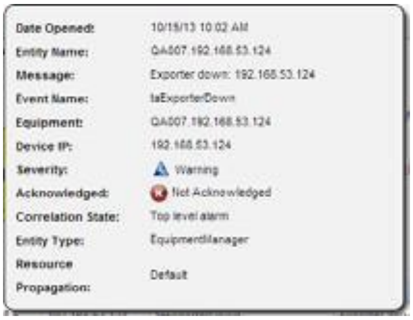

Date Opened:	10/15/13 10:02 AM
Entity Name:	QA007.192.168.53.124
Message:	Exporter down: 192.168.53.124
Event Name:	laExporterDown
Equipment:	QA007.192.168.53.124
Device IP:	192.168.53.124
Severity:	▲ Warning
Acknowledged:	✖ Not Acknowledged
Correlation State:	Top level alarm
Entity Type:	EquipmentManager
Resource Propagation:	Default

Hover over the *Device IP Address* column, and a tooltip appears with information about the alarm source's Model, Vendor, Management State, Discovery Date, and Description, with a title bar that indicates whether the device is running or not. Such tooltips elsewhere also include other device-dependent items. For example, bar graphs to display the % CPU [utilization], % Memory, and Description..

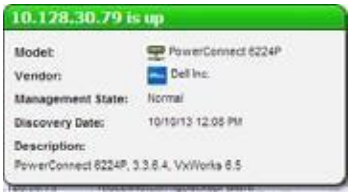

10.128.30.79 is up

Model:	PowerConnect 6224P
Vendor:	Dell Inc.
Management State:	Normal
Discovery Date:	10/10/13 12:05 PM

Description:
PowerConnect 6224P, 3.3.6.4, VxWorks 6.5

By default, the chart appears only when alarms exist. See Alarms Menu on page 280 for details about menu items available when you right-click in the summary and expanded portlets.

The following columns appear in this screen by default:

Severity—The alarm severity indicated by the color of the leftmost icon. The severity only has meaning for Alarms and Security Alarms. Informational Alarms get a severity level of Indeterminate. Closed alarms appear without color.

Date Opened — The date the alarm appeared.

Entity Name — The entity emitting this alarm (often within the Equipment).

DeviceIP — The IP address of the equipment where the alarm appeared.

Event Name — The event associated with the alarm.

Message — The message associated with the alarm.

Open the *Settings* > *Columns* screen to see additional possibilities for columns.

If an alarm is **Service Affecting,** (reflect an impact on a service) it propagates to appear as components of service- and customer-related alarms. The Service Affecting alarm column in this portlet does not appear by default. To see an alarm's propagation, show that column in the Using Extended Event Definitions portlet, where it is concealed by default.

NOTE:

You can select multiple rows in the Alarms portlet and the Managed Resources portlet. Many other columns are available, including those related to suppression, region, any parent alarm, and so on.

See Alarms in Topologies on page 246 for a description of how alarms appear in the topology portlet. The Expanded Alarm Portlet section below describes additional alarm portlet capabilities.

The Settings tool lets you select an Alarm Chart, the expanded Alarms portlet's totals, or no chart. The last two options permit selecting a filter. The Al arm Chart is a filter itself. If no data exists for the chart and the Chart option is on, the portlet returns to "no-chart" mode.

Settings persist if you have Admin rights or the Port let is on your Public/Private pages (like standard behavior).

NOTE:

Changes appear after you click *Apply*. The *Filter* panel disappears when you select the *Show Chart* option.

In addition to the Alarms summary portlet, there is the:

- Expanded Alarm Portlet
- Alarms Menu
- Parent/Child Alarm Correlation: Alarm Details Panel
- Alarm Email

Expanded Alarm Portlet

The expanded Alarm portlet appears when you click the plus (+) in the top right corner of the smaller screen.

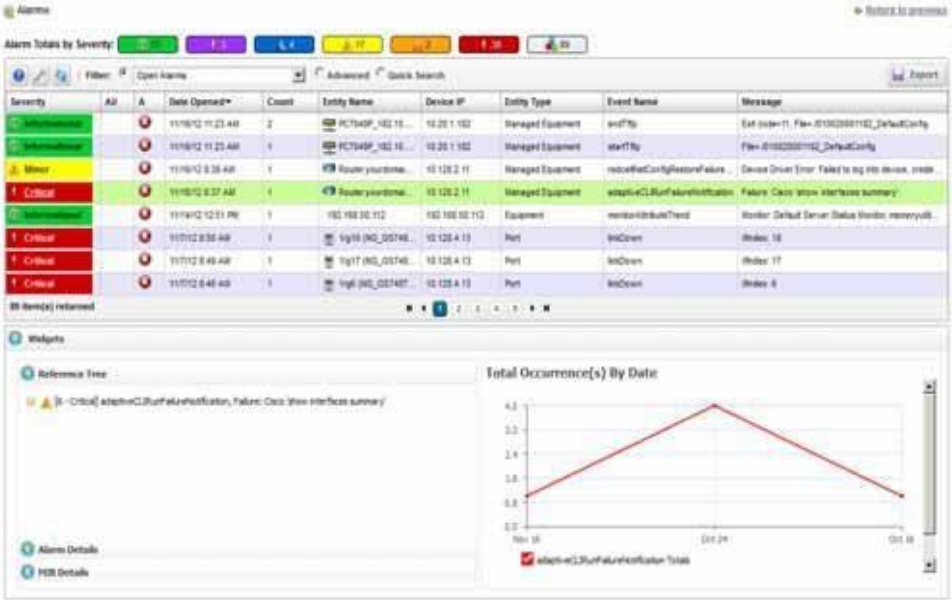
Alarms
Return to previous
Alarm Totals by Severity:
Filter:
Open Alarms
Advanced
Quick Search
Export
Widgets
Reference Tree
[5 - Critical] adaptiveCLIRunFailureNotification, Failure: Cisco 'show interfaces summary'
Alarm Details
MIB Details
Total Occurrence(s) By Date
4.0
3.2
2.4
1.6
0.8
0.0
Nov 16
Oct 24
Oct 18
adaptiveCLIRunFailureNotification Totals

Severity	Ack	A	Date Opened	Count	Entity Name	Device IP	Entity Type	Event Name	Message
Informational			11/16/12 11:23 AM	2	PC7048P_192.16...	10.20.1.182	Managed Equipment	endTftp	[illegible]
Informational			11/16/12 11:23 AM	1	PC7048P_192.16...	10.20.1.182	Managed Equipment	startTftp	[illegible]
Minor			11/16/12 8:36 AM	1	Router.yourdoma...	10.128.2.11	Managed Equipment	[illegible]	Device Driver Error: Failed to log into device, crede...
Critical			11/15/12 8:37 AM	1	Router.yourdomai...	10.128.2.11	Managed Equipment	adaptiveCLIRunFailureNotification	Failure: Cisco 'show interfaces summary'
Informational			11/14/12 12:51 PM	1	192.168.50.112	192.168.50.112	Equipment	monitorAttributeTrend	Monitor: Default Server Status Monitor, memoryutil...
Critical			11/7/12 8:50 AM	1	Vg18 (NG_G5748...	10.128.4.13	Port	linkDown	ifIndex: 18
Critical			11/7/12 8:49 AM	1	Vg17 (NG_G5748...	10.128.4.13	Port	linkDown	ifIndex: 17
Critical			11/7/12 8:48 AM	1	Vg6 (NG_G5748T...	10.128.4.13	Port	linkDown	ifIndex: 6

[illegible] item(s) returned

This displays listed alarms, totals by severity for alarm types found, and Snap Panel details of a selected alarm. By default, this screen adds the first of the following columns to those visible in the Event History portlet. To add the others listed here, right click, and select *Add Columns* to change the screen appearance.

NOTE:

All severity totals appear in expanded view. This display updates automatically when alarms clear.

Instead of the filtering graph on the summary portlet, the expanded portlet gives you the option to select a user-defined filter from the list, create custom filtering by selecting Advanced Filters, or do a quick search.

The following are available additional columns, besides the default columns visible in the Alarms Portlet (summary view):

Assigned User—The user currently assigned to this alarm (right-click to do this). The assigned user can then look for alarms by consulting the Assigned User (AU) column in the display (concealed by default), or byfiltering alarms using Advanced Filters. One can even create an alarm portlet that filters for a single user's assigned alarms.

Acknowledged— Indicates whether thealarm hasbeen acknowledged. Thisiscontrolled by the Acknowledge and Unacknowledge actions that are available from the popup menu.

Count—A count of the instances of the events that correlate to this alarm. Multiples of what is essentially the same alarm appear as a single row but increment this count.

Entity Type — The type of monitored entity.

State — The state (open/closed) of the alarm.

Date Cleared — The date and time that the alarm was closed.

Update Time— The time stamp for when this alarm was updated (for an additional count, the time the last duplicate was received).

Notification OID — The identifier of the notification displayed as an alarm. This is also the unique identifier of the Event Definition upon which the alarm was based, which means that the alarm inherits attributes such as Severity, Resource Propagation, etc. from the Event Definition that has the same Notification OID.

Equipment — The name for the top-level entity (Managed Resource) emitting the alarm.

Date Assigned — The date and time that the alarm was assigned.

Ack Time — The time the alarm was acknowledged.

Cleared By — The user who cleared the alarm.

MIB Text — The alarm's MIB Text.

Location — The location of the Managed Resource associated with the alarm.

Correlated Time—The alarm's date/time of correlation to a parent alarm (caused by or blocked by).

Entity Description — The description of the alarmed entity.

Mediation Partition — The mediation partition that received the event (traps, syslog, etc.) that is correlated to the alarm.

Resource Propagation — Indicates how this alarm propagates through the resource hierarchy, so as to possibly impact the Alarm State of the top-level device and/or the subcomponents that are hierarchicallyassociated with the entity that wasalarmed. Thisvalue is inherited from the Event Definition upon which the alarm is based.

Equipment — The equipment emitting the alarm.

Ack By — The user that acknowledged the alarm.

Correlation State—The role this alarm plays in any parent/child correlation (for example: *Top level alarm, caused by parent, Blocked by parent*).

Has Children—Red for no or green for yes: an indication of whether the alarm has children (see Parent/Child Alarm Correlation: Alarm Details Panel on page 282).

Notes — A text field to take notes about the alarm.

Parent Alarm — The name of the parent alarm (see Parent/Child Alarm Correlation: Alarm Details Panel on page 282).

Domain ID— The Multitenant domain ID emitting the alarm.

Service Affecting —Indicateswhetheror not the alarmaffectsservicesthat are provisioned against the alarmed entity. If this is true, then the alarm state of the entity might propagate to the associated services.

Correlated By —The name of the user who correlated this alarm to a parent.

The **Widgets** panel includes the followinginformation:

Alarm Details —The source, *Severity*, *Message*, *Date Opened*, and so on. See also Parent/Child Alarm Correlation: Alarm Details Panel below.

MIB Details—The *Notification OID*, and *MIB Text* for the selected alarm.

Reference Tree — The connection between the alarm and its source in tree form.

Total Occurrences by Date—A graph of the total occurrences of this alarm, by date.

Alarms Menu

Right clicking an alarm lets you select from the following pop-up menu options:

Edit — Access the editors for the Alarm (see Alarm Editor on page 284) or *Event Definition* (see Event Definition Editor on page 326).

Details(Shift+Click)—This menu item expands to displayoptions for the different detail viewing options. Select Alarm details to open a Details screen for the alarm itself; or select Equipment Details to open a Details screen for the entity emitting it. (see Connected Devices on page 189 for an example of this type of screen). The Alarm Details contains information like the MIB text, any Event Processing Rules invoked, and a Reference Tree for the alarm. It also lets you configure alarm correlation. See Parent/Child Alarm Correlation: Alarm Details Panel on page 282.

Topology — Display a topology map that includes the selected alarm(s). See Presentation Capabilities for more about these maps.

Acknowledge/Unacknowledge Alarm—Acknowledgestheselected Alarm(s). Thecurrentdate and timeappearinthe Ack Timefield. Unacknowledges previously acknowledged alarm(s) and clears the entries in the Ack By and Ack Time fields. The red “unacknowledged” icon appears in the expanded portlet and turns to a green check “acknowledged” icon the alarm has been acknowledged.

Assign User—Assign this alarm to one of the users displayed in the sub-menu by selecting that user. An icon also appears in the expanded portlet indicating the alarm has been assigned to someone.

Clear Alarm—Clearing the alarm removes the alarm from the default alarm viewand marks it as a candidate for the database archiving process (DAP). Essentially it is an indication to the system that the alarm has been resolved/addressed. If your system has enabled propagation policies, clearing recalculates dependent alarms.

Clear Group of Alarms —Sometimes you might have lots of open alarms that are unimportant because they are old and/or of lowseverity. Forexample, perhaps you want to clear all alarms that are informational and are more than a week old. Rather than having to clear them all individually, you can clear them as a group. Before selecting this menu item, you will need to create a filter for the group of alarms that you want to clear. When this menu item is selected, a panel will appear that contains all previously saved alarm filters. When you select a filter from the list and push Execute, it will clear all open alarms that meet the criteria of this filter.

△ CAUTION:
The Clear Group of Alarms action is irreversible.

Direct Access—Open an SNMP Mib Browser to the alarmed device, a CLI Terminal (Telnet window) to the alarmed device, or ICMP Ping the device alarmed. Only those available appear in the subsequent menu.

Email Alarm—Opens the Email Alarm window to email the selected alarm. Enter a subject an e-mail address to which you want to mail the alarm's content, and click add (+) to the list of addresses (the minus deletes them). Then click *Send Email.* Click *Cancel* to end this operation without sending e-mail.

SMTP setup is required to e-mail an alarm. See SMTP Configuration on page 63 for instructions about setting up e-mail from Cruz. See Alarm Email on page 283 for an example of what the content looks like.

Show Performance—Iftheequipment is monitored, thisdisplaysaperformance dashboardfor the alarmed equipment. See Dashboard Views on page 416 for more about these.

Edit Custom Attributes—Opens the Custom Attribute Editor where you define field characteristics, such as whether it is enabled, the label name, and the tooltip.

Aging Policy—This lets you select a policy that determines how long this alarm remains in the database. See Implementing DAP on page 80 for information about configuring such policies.

View as PDF— Create an Acrobat PDF document containing this alarm's contents as displayed in the summary portlet basic columns.

Share with User—Opens the Share with User window where you select the colleague you want to share the selected rules with and then type your message.

NOTE:
To resync alarms—that is, query the device for its alarm state—resync the device.

Parent/Child Alarm Correlation: Alarm Details Panel

Like many other items managed by Cruz, the Alarms portlet has a Details panel (see Connected Devices on page 189, for example). In addition to the items mentioned above, the Alarm Details also have a *Correlations* panel.

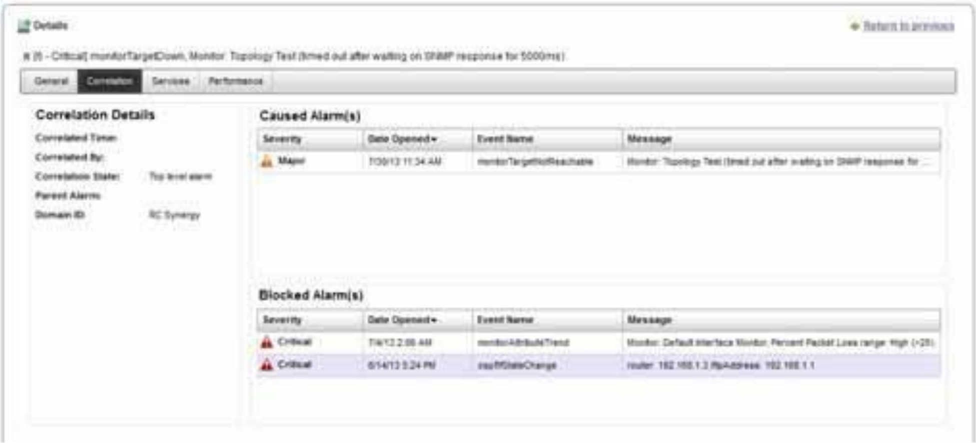

[6 - Critical] monitorTargetDown, Monitor: Topology Test (timed out after waiting on SNMP response for 5000ms)

Correlation Details

Correlated Time:

Correlated By:

Correlation State: Top level alarm

Parent Alarms:

Domain ID: RC Synergy

Caused Alarm(s)

Severity	Date Opened	Event Name	Message
Major	7/30/13 11:34 AM	monitorTargetNotReachable	Monitor: Topology Test (timed out after waiting on SNMP response for ...

Blocked Alarm(s)

Severity	Date Opened	Event Name	Message
Critical	7/4/12 2:08 AM	monitorAttributeTrend	Monitor: Default Interface Monitor: Percent Packet Loss range: High (>25)
Critical	6/14/13 3:24 PM	ospfIfStateChange	router: 192.168.1.3 IfIpAddress: 192.168.1.1

Alarm parent/child correlation lets you correlate one alarm to another so you can conceal the child alarm(s) in standard alarm views. Youcan also correlate one alarm to blockresolution ofanother. In effect, parent alarms conceal (or block) correlated child alarms. This can confine alarms that appear or resolve to those requiring action only.

The *Correlations* panel lists the following:

Correlation Details — The current correlation state of the current alarm. Attributes include: *Correlation Date, Correlated By*, and *Correlation State* along with information about the correlated, parentalarm. Inaddition to manuallyremovingthecorrelation, youcanright-click in this component to navigate to the details of the correlated alarm.

Caused Alarm(s) —The alarms caused by this alarm. Right-click to add alarms to the table or to remove them.

Blocked Alarm(s) — The alarms that are currently blocked from resolution by this alarm. Right-click add alarms to the table (and remove them).

When adding a correlated alarm, you can select any alarm that does not already have a parent alarm, then click *Done* to make it a child alarm.

When you correlate correlating one alarm to another, Cruz understands their correlation state as either *Caused By* or *Blocked By*. By default, the correlation state is *Top Level Alarm*.

Cruz does not support multiple correlation states so one alarm cannot be both *Caused By* and *Blocked By*. However, a parent alarm can have both *Caused By* and *Blocked By* child alarms. Asingle alarm can also be both parent and child. Consider, for example, an alarm that causes several other alarms but is *Blocked By* another alarm. Its parent alarm would appear in the *Correlation Details* panel too.

Alarmscorrelated as *Caused By* clearautomatically when the parent alarm clears. *Blocked by* alarms status as child alarms disappears when their parent alarm clear. This means they become visible again within the alarm views.

Here are a few things you need to know about alarm correlation state:

- Alarm Processing Effects

  An alarm's correlation state does not affect alarm processing behavior. So, if a clearing event enters Cruz, the open alarm clears regardless of its correlation state. Event counts also continue to increase as duplicate events arrive.

- Default Filtering

  By default, all alarm filters exclude child alarms. A filter criteria (*Include child alarms*) can include child alarms, so you can always see all alarms regardless of their correlation state by selecting the *Include child alarms* attribute within the expanded alarms portlet and setting the search criteria to *Is true*.

- Additional Alarm Attributes

  The following attributes reflect the correlation state for an alarm. When another alarm conceals the child alarm (or blocks it), it sets the following too.

  - Correlated By: User who created the correlation
  - Correlation Date: Date the correlation was created
  - Correlation State: CausedBy or BlockedBy
  - Parent Alarm
  - Has Children: Whether the alarm has children

Alarm Email

The e-mail sent by right clicking an alarm has the subject specified when you send it and contains the information within the alarm. For example:

```
Alarm: monitorIntervalSkip
Alarm Attributes:
===============================
Device IP          =
Message            =
Alarm State        = Open
Severity           = 5 - Major
Count              = 1
Date Opened        = Tue Dec 14 22:01:30 PST 2010
Update Date/Time   = Tue Dec 14 22:01:36 PST 2010
Entity Name        =
Entity Type        =
Entity Description =
Equipment          =
Region             = SUPDEMOPartition
Location           =
Assigned By        = OWSystem
Date Assigned      = Thu Dec 16 10:40:24 PST 2010
Assigned User      = qatester
Acknowledged       = false
Ack By             =
Ack Time          =
Cleared By        =
Date Cleared      =
MIB Text          = Monitor session was skipped due to resource
  constraints.  Typically, this implies one or more monitors should run
  less frequently.  This may also be caused by many timeouts which force
  executions to take longer to complete than normal.
Advisory Text     =
```

Alarm Editor

The Alarm Editor window appears when you right-click an alarm and select *Edit Alarm* from the Alarms Portlet.

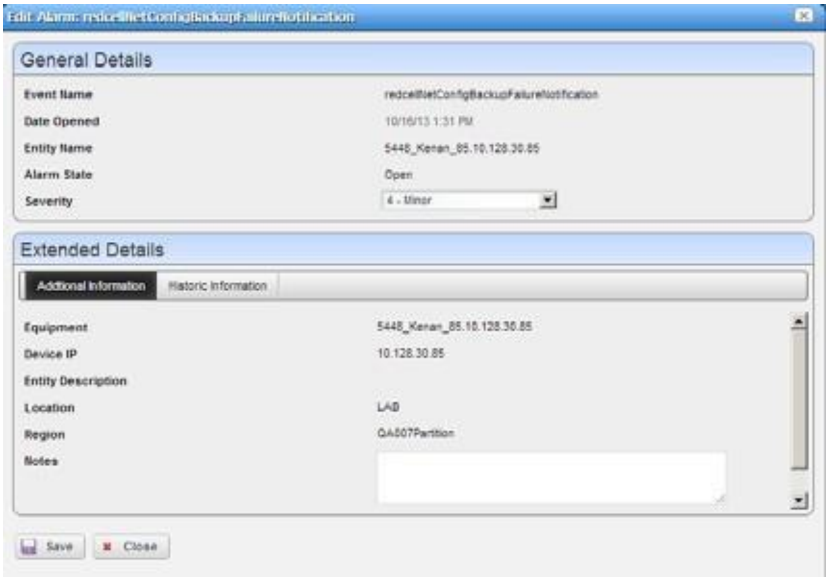

Edit Alarm: redcellNetConfigBackupFailureNotification

General Details

Event Name	redcellNetConfigBackupFailureNotification
Date Opened	10/16/13 1:31 PM
Entity Name	5448_Kenan_85.10.128.30.85
Alarm State	Open
Severity	4 - Minor

Extended Details

Equipment	5448_Kenan_85.10.128.30.85
Device IP	10.128.30.85
Entity Description	
Location	LAB
Region	QA507Partition
Notes	

You also have the option to edit the Event Definition (see Using Extended Event Definitions on page 333) or open the alarmed device’s Details panel (see Connected Devices on page 189). The Alarm Editor window contain the following type of information:

Information Type	Provides
General Details	**Event Name** — The event that triggered the alarm. **Date Opened** — The date the alarm occurred. **Entity Name** — The entity emitting this alarm (often within the Equipment). **Alarm State** — The state of the alarm (Open/Closed). **Severity** — The alarm severity indicated by the color of the leftmost icon. The severity only has meaning for Alarms and Security Alarms. Informational Alarms get a severity level of Indeterminate. Closed alarms appear without color. If you change the severity, you may have to refresh the portlet after you save the changed alarm.

Information Type	Provides
Extended Details: Additional Information	**Equipment** — The equipment (not subcomponent) that triggered the alarm. **DeviceIP** — The IP address of the equipment where the alarm appeared. **Entity Description** —A description of the triggering equipment. **Location** — The location for the alarm. See Locations Portlet on page 217. **Region** — The partition/region for the alarm. **Notes** —A field where you can enter text.
Extended Details: Historic Information	This panel contains primarily read-only fields describing the alarm, including whether it was *Acknowledged, Ack by, Ack Time, Count* and so on.
Extended Details: Custom Fields	If you have created any Custom Fields for Alarms, this panel appears in the editor. See Editing Custom Attributes on page 146 for instructions about these.

Setting Audible Alerts

Audible Alerts produce a sound when a new alarm arrives in the Ala rms summary (not expanded) portlet. The sound occurs when the auto-refresh controller polls for state changes. If you enable Audible Alerts and the Alarmssummary viewreceives **new** alarms, the specified sound occurs.

Cleared alarms do not trigger a n audible alert. **Only new** alarms received triggered an audible alert during au to-refresh. To cut down on audio clutter, only a single Audible Alert sounds no matter how many alarms occur du ring an auto-refresh cycle.

Each browser supports sound differently because of licensing for various sound formats. Audible alarm support exists for most browsers, so if issues occur with a browser the workaround is either to upgrade or use Chrome.

Browsers support MP3 sounds the most, so this is the only format supported for Audible Alerts. The Firefox browser only support OGG format natively and the Internet Explorer browser has issues with most sounds. To support those browsers, the Cruz application plays MP3 sounds through a F lash Object, so browsers need no special plugins.

This section provides instructions for:

- Turning on Audible Alerts
- Adding Custom MP3 Sounds

Turning on Audible Alerts

Turn on Audible Alerts as follows.

1 Navigate to the page containing the Alarms portlet.

NOTE:

For audible alerts to work, the Alarms portlet must be on a page without the Hierarchical View portlet or other context broadcasting that dynamically changes the Alarms portlet's context. Auto refresh does not run when in this environment so as a result the Audible Alerts are not exposed. (See Understanding Hierarchical View on page 248.)

2 Click the Settings (wrench) tool.

The Settingwindowisdisplayed.

3 Enable the audible alerts for new alarms option.

This activates the sound field.

4 Select a sound to plays.

5 Click the play button to preview the selected sound.

By default, the Cruz product ships with four standard alert sounds: Alert, Bell, Chord, and Ding. See Adding Custom MP3 Sounds to add custom sounds.

6 Click Apply.

This Alarms portlet instance now has Audible Alerts enabled.

Adding Custom MP3 Sounds

Add custom MP3 sounds as follows.

1 Go to the Control Panel.

2 Click the Documents and Media section.

3 Click the Add button and then Select Basic Document.

4 Click Choose File and pick an MP3 file to upload from the File section.

Because this interface lets you add any type of media, no file validation occurs. However, Audible Alerts only display audio/MP3 mime-types.

5 Give the new MP3 sound a short title.

For example, if you upload cowsound.mp3, call it Cow Sound.

6 Click Publish.

7 Go back to the Alarms portlet.

8 Click the Settings (wrench) tool.

Verify that there is a new sound to select.

Event History

Not all events appear as alarms. Event History preserves all event information for your system.

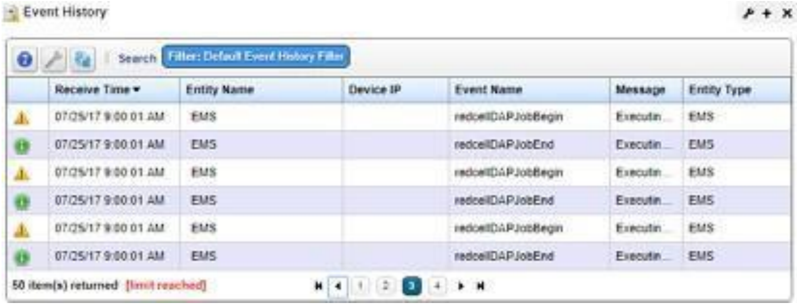

	Receive Time ▾	Entity Name	Device IP	Event Name	Message	Entity Type
⚠	07/25/17 9:00:01 AM	EMS		redcellDAPJobBegin	Executin...	EMS
●	07/25/17 9:00:01 AM	EMS		redcellDAPJobEnd	Executin...	EMS
⚠	07/25/17 9:00:01 AM	EMS		redcellDAPJobBegin	Executin...	EMS
●	07/25/17 9:00:01 AM	EMS		redcellDAPJobEnd	Executin...	EMS
⚠	07/25/17 9:00:01 AM	EMS		redcellDAPJobBegin	Executin...	EMS
●	07/25/17 9:00:01 AM	EMS		redcellDAPJobEnd	Executin...	EMS

50 item(s) returned [limit reached]

The Event Historyportlet (**summary view**) displays an icon whose color reflects any alarm state associated with the event. It also displays the Receive Time, Entity Name, Device IP, and Event Name, Message, and Entity Type. Right-click to edit event definition or entity details, viewevent details, view or modify aging policies, create a PDF, or share and event with users.

NOTE:

Hovering over the *Device IP* column produces a tooltip like the Alarms tooltip.

The default filter for this portlet displays only recent events. If you do not see the desired events, expand the period for which they appear.

Clicking the plus (+) in the upper right corner of the initial portlet view displays the Event History portlet (**expandedview**). As in other expanded portlets, use the filteringcapabilities to further limit the events listed.

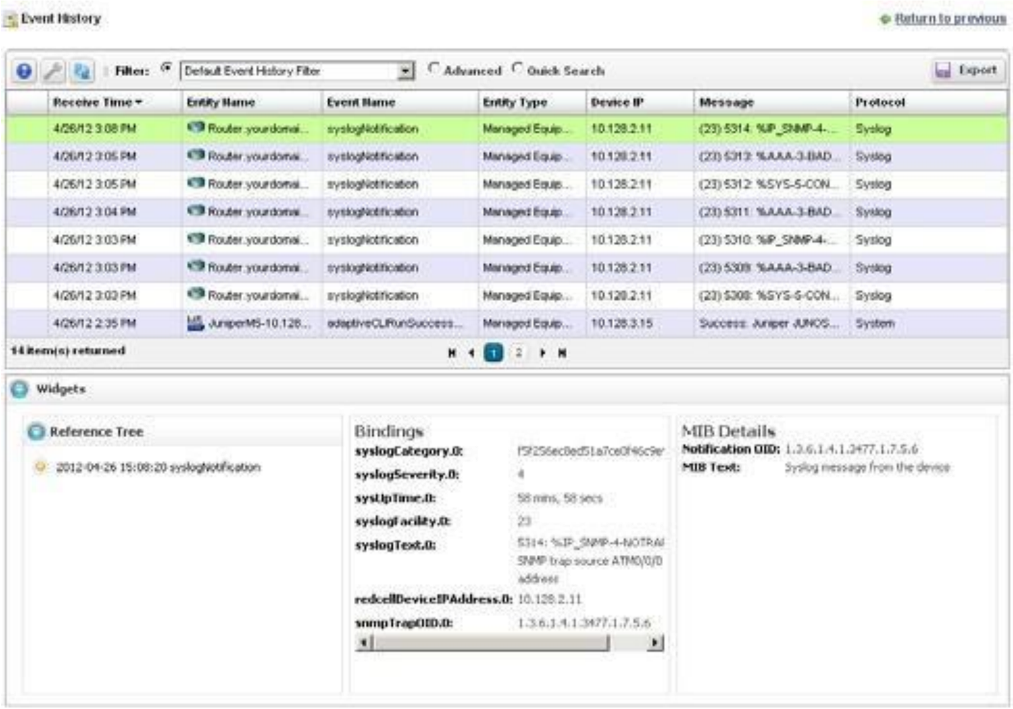

Event History

Return to previous

Filter: ● Default Event History Filter ○ Advanced ○ Quick Search

Receive Time	Entity Name	Event Name	Entity Type	Device IP	Message	Protocol
4/26/12 3:08 PM	Router.yourdomai...	syslogNotification	Managed Equip...	10.128.2.11	(23) 5314: %IP_SNMP-4-...	Syslog
4/26/12 3:05 PM	Router.yourdomai...	syslogNotification	Managed Equip...	10.128.2.11	(23) 5313: %AAA-3-BAD...	Syslog
4/26/12 3:05 PM	Router.yourdomai...	syslogNotification	Managed Equip...	10.128.2.11	(23) 5312: %SYS-5-CON...	Syslog
4/26/12 3:04 PM	Router.yourdomai...	syslogNotification	Managed Equip...	10.128.2.11	(23) 5311: %AAA-3-BAD...	Syslog
4/26/12 3:03 PM	Router.yourdomai...	syslogNotification	Managed Equip...	10.128.2.11	(23) 5310: %IP_SNMP-4-...	Syslog
4/26/12 3:03 PM	Router.yourdomai...	syslogNotification	Managed Equip...	10.128.2.11	(23) 5309: %AAA-3-BAD...	Syslog
4/26/12 3:02 PM	Router.yourdomai...	syslogNotification	Managed Equip...	10.128.2.11	(23) 5308: %SYS-5-CON...	Syslog
4/26/12 2:35 PM	JuniperMS-10.128...	adaptiveCLIRunSuccess...	Managed Equip...	10.128.3.15	Success: Juniper JUNOS...	System

14 item(s) returned

Widgets

Reference Tree

2012-04-26 15:08:20 syslogNotification

Bindings

syslogCategory.0: f5f256ec8ed51a7ce0f46c9e

syslogSeverity.0: 4

sysUpTime.0: 58 mins, 58 secs

syslogFacility.0: 23

syslogText.0: 5314: %IP_SNMP-4-NOTRAI SNMP trap source ATM0/0/0 address

redcellDeviceIPAddress.0: 10.128.2.11

snmpTrapOID.0: 1.3.6.1.4.1.3477.1.7.5.6

MIB Details

Notification OID: 1.3.6.1.4.1.3477.1.7.5.6

MIB Text: Syslog message from the device

The expanded view has columns similar to those described in Alarms Portlet on page 276 or Expanded Alarm Portlet on page 278. Configure these as visible or hidden by clicking *Settings.* The following are some additional columns available.

Receive Time — The date the event was received.

Entity Name — The entity emitting the event.

Event Name — The event identifier.

Entity Type — Typically something like *Managed Equipment.*

Protocol — The protocol that delivered the event. Commonly this is either *SNMPv1, SNMPv2c, or SNMPv3,* indicating that the event originated from some version of an SNMP trap or inform notification. Other possible values that indicate a notification was received from an external system include Syslog and HTTP REST. System; indicating Cruz itself delivered it.

Instance ID — The instance identifier for the event.

Mediation Server IP—The mediation server retrieving the event. If you have a single-server environment, this is blank. It is most useful in a clustered environment.

Location — The location of the entity emitting the event.

Equipment — The equipment emitting the event.

SubType — A classification for the event. For example: *Trap, Inform, etc*.

Notification OID —The object identifier (OID) for the Event Definition upon which this event is based.

Source IP — The source of the event's IP address.

Region— The region emitting the event.

Click a listed alarm to display its details in the **Widgits panel**. The *Reference Tree* displays the event's relationship to any alarms, and to the source device. Click the plus (+) next to an item in the tree to unpack it.

The *Bindings* Snap Paneldisplaystheevent's VariableBinding(varbind) information, includingthe Binding OID, the device's IP address, and other event-specific information.

The *MIB Details* Snap Panel includes MIB information like the Notification OID and MIB Text.

You can right-click the listed events and *Share with User* (see *Sharing a Resource* on page 146), or Modifying Column Settings on page 144.

Automation and Event Processing Rules

Use the Automation and Event Processing Rules portlet to configure and maintain event processingandautomation rules. Youcancreatecustom multitenant domainsandeventprocessing rules (EPRs) within any Multitenant site. Aside fromthe filter criteria, which EPRapplies depends on the EPR's Domain ID compared to the target entity's notification Domain ID. All EPRs with the root site's Domain ID apply to all arriving events. EPRs with another Domain ID only apply to events for entities that have been assigned the same Domain ID.

This portlet is intended for users who are interested in configuring and maintaining event processing and automation rules. For steps to create these rules, see Creating Event Processing Rules on page 292.

Access this portlet by selecting Alarms > Automation from the navigation bar. This portlet has both a summary view and an expanded view. Each view could display different Columns and has the same Pop-Up Menu options available.

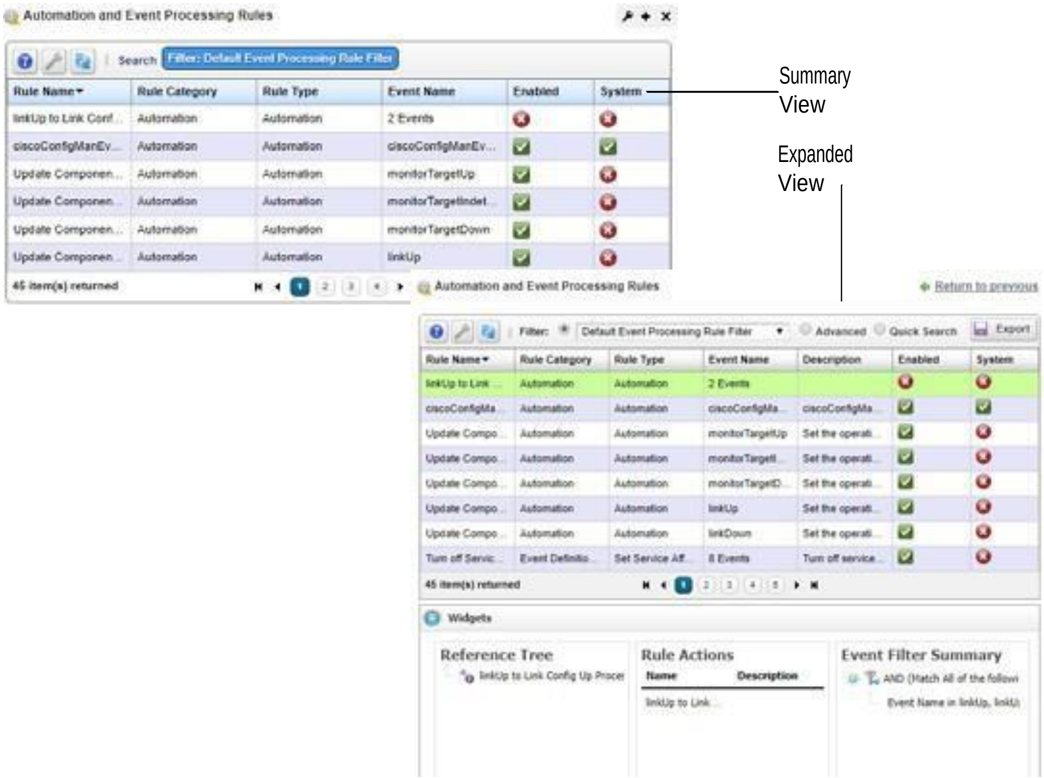

Automation and Event Processing Rules

Rule Name	Rule Category	Rule Type	Event Name	Enabled	System
linkUp to Link Conf...	Automation	Automation	2 Events	☒	☒
ciscoConfigManEv...	Automation	Automation	ciscoConfigManEv...	☑	☑
Update Componen...	Automation	Automation	monitorTargetUp	☑	☒
Update Componen...	Automation	Automation	monitorTargetIndet...	☑	☒
Update Componen...	Automation	Automation	monitorTargetDown	☑	☒
Update Componen...	Automation	Automation	linkUp	☑	☒

45 item(s) returned

Automation and Event Processing Rules

Rule Name	Rule Category	Rule Type	Event Name	Description	Enabled	System
linkUp to Link ...	Automation	Automation	2 Events		☒	☒
ciscoConfigMa...	Automation	Automation	ciscoConfigMa...	ciscoConfigMa...	☑	☑
Update Compo...	Automation	Automation	monitorTargetUp	Set the operati...	☑	☒
Update Compo...	Automation	Automation	monitorTargetI...	Set the operati...	☑	☒
Update Compo...	Automation	Automation	monitorTargetD...	Set the operati...	☑	☒
Update Compo...	Automation	Automation	linkUp	Set the operati...	☑	☒
Update Compo...	Automation	Automation	linkDown	Set the operati...	☑	☒
Turn off Servic...	Event Definitio...	Set Service Af...	8 Events	Turn off service...	☑	☒

45 item(s) returned

Reference Tree

linkUp to Link Config Up Proce

Rule Actions

Name	Description
linkUp to Link ...	

Event Filter Summary

AND (Match All of the followi

Event Name in linkUp, linkU

Columns

Other than the general navigation and configuration options, the Automation and Event Processing Rulesportlets (summary and expand ed views) include the followingcolumns. The columns displayed by default are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Rule Name	A descriptive name given to the rule. This field displays on the summary and expanded views by default.
Rule Category	The stage in the event and/or alarm life cycle that the rule is triggered. For example: • Protocol Translation rules trigger early in the life cycle when non-SNMP messages, such as Syslog are received. • Stream Based Correlation rules also trigger early in the life cycle to minimize the event processing load on the application server and to improve throughput. • Event Definition Override rules trigger after the aforementioned rule categories, but prior to alarm correlation. • Automation rules trigger a little later in the life cycle so that the results from the execution of those rules are taken into account during automation. This field displays on the summary and expanded views by default.

Column	Description
Rule Type	Text that indicates what the rule does. For example, there are several rule types within the Event Definition Override category that change specific event attributes to override default values otherwise inherited from the event definition. To consider a specific example of this, you can make an Event Definition Override Event Processing Rule that sets an event as service-affecting. These rules override the default service affecting field that would otherwise be entirely determined by the event definition (as determined by the notification type). **Note:** The installation provides some seeded event processing rules. You can edit some seeded rules, others you cannot edit. If you edit a system seeded rule and want to restore the default settings, you need to delete the rule from the portlet and then enter `ocpinstall -s` from an oware command line to re-seed the database. This field displays on the summary and expanded views by default.
Event Name	The event associated with the rule. This field displays on the summary and expanded views by default.
Enabled	An indicator that the rule is enabled (checkmark) or not (X). This field displays on the summary and expanded views by default.
System	An indicator that the rule is a system rule (checkmark) or a non-system (user-defined) rule (X). This field displays on the summary and expanded views by default.
Description	Text describing the rule in more detail. This field displays on the expanded view by default.
Widget	Additional information about the selected rule, such as: • **Reference Tree** shows the connection between the rule and its event. • **Rule Actions** list any configured actions associated with the rule. • **Event Filter Summary** shows configured filters for the selected rule. The Widgets field is available **only** from the expandedview.
Component	The component the rule accesses, such as Extreme Networks device driver.
Domain ID	The identifier for the resource domain.
Icon	The icon that represents this resource.
Valid	An indicator that the rule is valid (checkmark) or not (X).

Pop-Up Menu

The Automation and Event Processing Rulespop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
New	Provides access to the following menu options: • **Protocol Translation** displays the Syslog Escalation option. Select this option to open the Creating New Syslog Escalation Rule window and define a syslog escalation rule. • **Stream Based Correlation** displays the Frequency Threshold and State Flutter. options. Select the event processing rule type to open the Creating New *eventType* Event Processing Rule window and define a new event processing rule. • **Event Definition Override** displays the Reject Event, Suppress Alarm, Set Severity, Set Service Affecting, and Device Access options. Select the event processing rule type to open the Creating New *eventType* Event Processing Rule window and define a new event processing rule. • **Automation** opens the Creating New Automation Event Processing Rule window, where you define an automation processing rule.
Edit	Opens the Editing Automation Event Processing Rule window, where you can view the rule details or edit the rule if it is not read-only.
Copy	Creates a new rule derived from the selected rule with some changes to the properties, filtering, members, or actions. **Note:** The Cruz generates a new name, such as CopyOfUpdate Component State. You must change that name before you save the event processing rule.
Audit	Opens the Audit Trail Viewer window, which lists existing audit records and jobs for the selected record.
Delete	Removes the selected rule.
Import/Export	Provides the following actions when available for the selected rule: • Import retrieves a file containing XML compliance descriptions. Some imports can come from a URL. • Export Selection exports the selected description to an XML file. • Export All exports all descriptions to an XML file. Click Download Export File to specify where to save the file. The Import/Export option is useful as a backup or to share descriptors with other projects. You must import data into the correct portlet. For example, you cannot import event data into the Actions portlet. If one type of data depends on another, you must import the other data before importing the data that depends on it.
Share with User	Opens the Share with User window where you select the colleague you want to share the selected rules with and then type your message.

Multitenant Domains and Event Processing Rules (EPRs)

Youcan create custom EPRs within any Multitenant site. Aside from the filter criteria, which EPR applies depends on the EPR's Domain ID compared to the target entity's notification Domain ID. All EPRs with the root site's Domain ID apply to all arriving events. EPRs with another Domain ID only apply to events for entities that have been assigned the same Domain ID.

Creating Event Processing Rules

There are many different types of rules you can create. As an example, the steps provided show how to create a device access rule from the Automation and Event Processing Rules portlet. If you need a detailed description of this portlet, the rule types, or other menu option, see Automation and Event Processing Rules on page 288. Also see Filtering/Settings on page 294, Syslog Escalation Criteria on page 298, and Actions on page 302 for more about the differences available between rule types.

Create a rule from the Automation and Event Processing Rules portlet as follows.

1 Right-click and then select Event Definition Override > Device Access.

The Rule Editor is displayed.

2 Enter a name to identify the rule and an optional description.

3 Select Enabled if you want this rule to begin workingimmediately.

4 Click Next to specify filters.

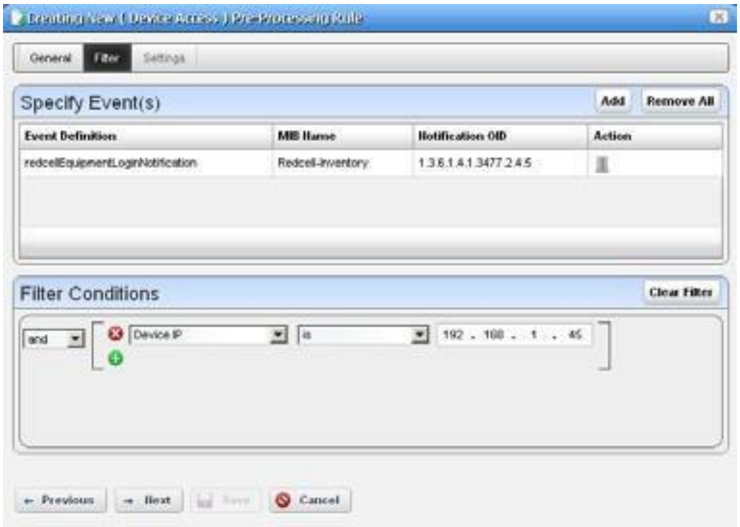

Event Definition	MIB Name	Notification OID	Action
redcellEquipmentLoginNotification	Redcell-Inventory	1.3.6.1.4.1.3477.2.4.5	

and · Device IP · is · 192 . 168 . 1 . 45

5 Add the events to filter.

a. Click Add to show the events list.

b. Select the events you want to add.

c. Click Add Filter to further filter the selected events.

6 Click Next to specify the settings for the selected rule type.

This panel's appearance depends on the type of rule you selected when you clicked New. When you are editing an existing rule, it defaults to that rule's screen. For more about the available alternatives, see Filtering/Settings on page 294.

The Device Access example creates a specific deviceaccess event for user login, logout, login failure, or configuration change.

7 Select the Access Type(Config Change, Login Failure, User Login, User Logout) from the pick list for that field.

8 Enter the Username Variable and/or User Name RegEx match string.

This confines rule response to the selected users.

9 Select Suppress Correlated events if you do not want to see events correlated with this one.

10 Click Save to preserve the event processing rule.

NOTE:

To test these rules, you typically need specialized trap-sending software. However, you can make a rule respond to an internal Cruz event, such as backup failure if you only want to see the outcome. Simply disable your FTP server) and back up a device to get a backup failure event.

Rule Editor

After you select a category and type for a new rule, the Rule Editor is displayed, where you manage the event processing described briefly in Creating Event Processing Rules on page 292. The editor has the following panels:

- General
- Filtering/Settings
- Syslog Escalation Criteria (for Syslog Escalation)
- Actions (for automation rules)

Subcomponent names must cache on the server if you want to refer to them in rules. For example, if you want e-mail whenever a linkDown occurs on a port, then you must cache subcomponents. If you cache subcomponents, it impacts performance, which is why such caching is disabled by default.

To enable caching, set the following property in the installed.properties file and then restart the Application server:

```
com.dorado.redcell.inventory.equipment.subcomponent.cache=true
```

General

The General panel is common to all rule types.

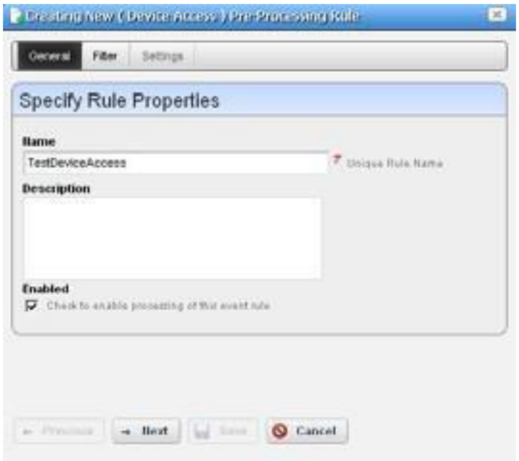

Creating New (Device Access) Pre-Processing Rule

Specify Rule Properties

Name: TestDeviceAccess

Description:

Enabled: ☑ Check to enable processing of this event rule

It contains the following fields:

Name — Enter a text identifier for the rule.

Description — Enter an optional text description of the rule

Alarm Only—Visible only in automation rules. Check this to enable the rule only if an alarm is generated, not suppressed.

Explicit Membership — Visible only in automation rules. This indicates that this rule has explicit membership. If any explicit members are defined for a rule, then the event must be associated with one of these entities for this rule to execute. This box must be checked in order to enable the buttons to add entities on the Explicit Members tab.

Enabled — Select this option to enable the rule.

Filtering/Settings

For all rule types, select the *Event Definition*. Click *Add* to open a screen where you can select events to include in the event you are creating. This includes a filter at the top that you can use to search for specific events. For example, Event Name Contains. Click A dd Selection to include selected items in this filter or Add All to include all displayed events. After you finish event selection, click Done at the bottom of this selection screen.

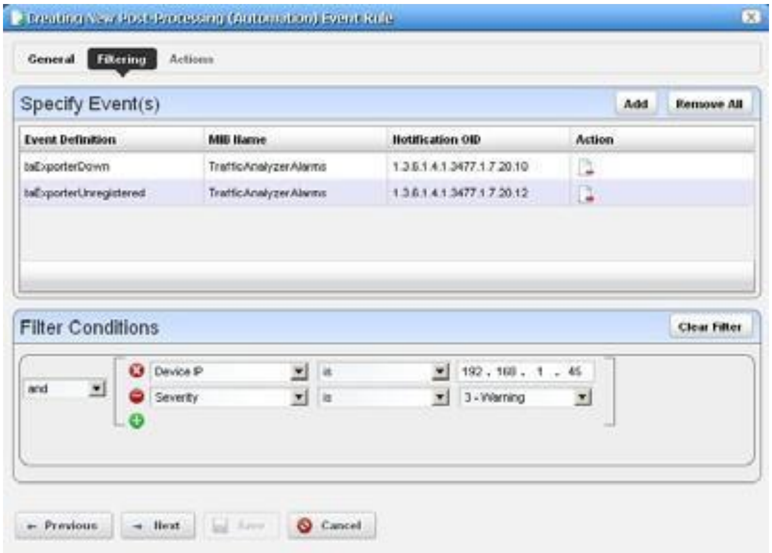

Event Definition	MIB Name	Notification OID	Action
taExporterDown	TrafficAnalyzerAlarms	1.3.6.1.4.1.3477.1.7.20.10	
taExporterUnregistered	TrafficAnalyzerAlarms	1.3.6.1.4.1.3477.1.7.20.12	

and
Device IP is 192 . 168 . 1 . 45
Severity is 3 - Warning

Click *Add Filter* to further filter the selected events. See Defining Advanced Filters on page 145 for more about this feature. After you *Add Filter* the button changes to *Clear Filter* so you can remove any filter from the event rule.

NOTE:

Cruz supports multiple IP addresses per resource. During event processing, filters that include IP address criteria maybehaveincorrectly when Cruz evaluates thefilter. Best practice is usingresource name(s) instead of IP addresses.

The following types are **processing rule types** and a description of their properties.

Rule Type	Properties
Protocol Translation	These rules define how to translate messages that are received from protocols other than SNMP into a format that is more similar to an SNMP trap so that they can become events in Cruz. SNMP traps always have all of the data in their payload contained within variable bindings, and this is also the way events are structured within Cruz so for the system to be able to have events that come from protocols other than SNMP, they first have to undergo this translation. Currently the only rule type within this category is *Syslog Escalation*, and rules of this type translate incoming syslog messages into a payload that is more like an SNMP trap so that they can become events based on the definition syslogNotification.
Stream Based Correlation	These rules allow you to control the stream of events that are processed. You can do this by specifying a correlation pattern that the system should detect among the events within the stream and then also specifying the action that the system should take to modify the stream in some way. Correlation patterns can take into account the frequency of certain types of similar events (as is the casefor *Frequency Threshold* rules) or they can operateon the intermittent fluttering of the raising and clearing of certain types of alarms (as is the case for *State Flutter* rules). You can use rules of types within this category for the purpose of minimizing the number of events submitted to the application server for further processing.
Event Definition Override	By default, events inherit all of their attributes (such as severity, behavior, etc.) directly from the event definition identified by the notification type OID. These rules allow you to override the default attributes such as severity (*Set Severity*), service affecting (*Set Service Affecting*), and/or behavior (with rules of the type Reject Event or *Suppress Alarm*). *Device Access* rules also fall into this category because you can use rules of this type to normalize the device-specific events into standardized events.
Automation	These rules execute specified actions for the rule after the event processing occurs. Note that Automation is both a rule category and a rule type since there is only one type of rule within this category. The following are *Event Definition Override* rule types: • **Reject Event** — This screen presents the Specify Event Filtering portion of the screen without any Settings in the lower screen. Specify events toreject with this selection and filtering. • **Set Severity** — This rule overrides the default alarm severity of anevent selected and filtered in the upper screen. 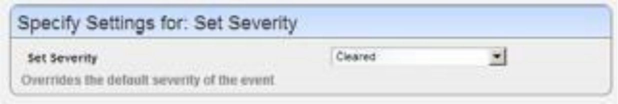 Specify Settings for: Set Severity Set Severity: Cleared Overrides the default severity of the event • **Suppress Alarm** — This screen presents the Specify Event Filtering portion of the screen without any Settings in the lower screen. Specify events/alarms to suppress with this selection and filtering.

Rule Type	Properties
Device Access	The Device Access example creates a specific device access event for user login, logout, login failure or configuration change. Select the *Access Type* (*Config Change, Login Failure, User Login, User Logout)* from the pick list for that field. 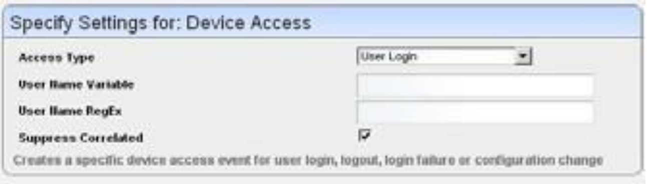 Specify Settings for: Device Access Access Type: User Login User Name Variable User Name RegEx Suppress Correlated ☑ Creates a specific device access event for user login, logout, login failure or configuration change Enter the *Username Variable* and/or *User Name RegEx* match string in those fields. This confines rule response to the selected users. Select *Suppress Correlated* events if you do not want to see events correlated with this one.

Stream Based Correlation has the following rule types:

Frequency Threshold —This rule type changes event behavior based on the frequency of the selected event. For successive events of the same type, associated with the same entity, it suppresses or rejects the first fewreceived, up to the given event threshold, within the pattern expiration time, and publishes the rest.

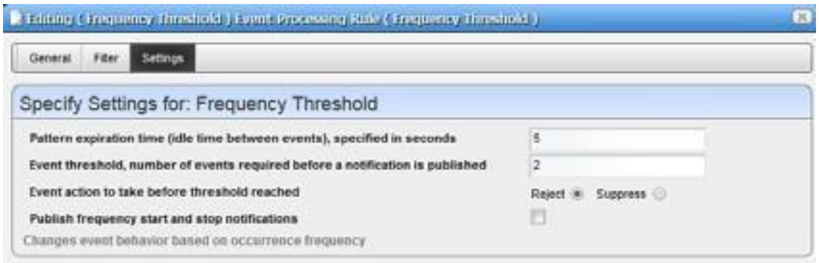

Editing (Frequency Threshold) Event Processing Rule (Frequency Threshold)

Specify Settings for: Frequency Threshold

Pattern expiration time (idle time between events), specified in seconds: 5

Event threshold, number of events required before a notification is published: 2

Event action to take before threshold reached: Reject ● Suppress ○

Publish frequency start and stop notifications ☐

Changes event behavior based on occurrence frequency

Enter the *Pattern expiration time (idle time between events), specified in seconds* and *Event threshold, number of events required before a notification is published*, then select an *Event Action to take before the threshold is reached* (*Reject* or *Suppress* the event). If you *Reject* an event, it does not appear in Event history; if you *Suppress* it, it creates no alarm, but it does appear in the Event history. Check *Publish frequency start and stop notifications* if you want Cruz to keep a record of when this rule starts and stops filtering events.

On receipt of the first event matching the given filter criteria, Cruz enables the selected pattern. It remains active until no matching events are received for at least the number of secondsspecified as the pattern expiration time. The rulealwayswaits this number ofseconds before publishing the event(s), even if the number of matching events crosses the threshold before the pattern expires. Every time the rule reaches its threshold in this time window, it publishes one event and then reset the counter.

For example, consider a pattern configured for 3 events in 10 seconds. If Cruz receives only 2 matching events in a 10 second time window, then it publishes no events. With these same parameters, if Cruz receives at least 3 but less than 6 (3 times 2) events, then Cruz publishes one event. If it receives six events, then it publishes two events (because this amounts to 3 times 2).

State Flutter—This type of rule changes event behavior on transient raising/clearing state changes for events that are correlated to each other. For example, if there is a series of flapping linkUp and linkDown events for the same interface, but you do not want the alarm state of this interface to change rapidly, then you can use a rule of this type to filter out the noise so that the alarm state reflects the most common state of the interface. When rules of this type are configured, the system looks for successive raising and clearing events that correlate to each other and are associated with the same entity and it publishes the final state in after a given number of seconds has elapsed, and it suppresses or rejects the extra events.

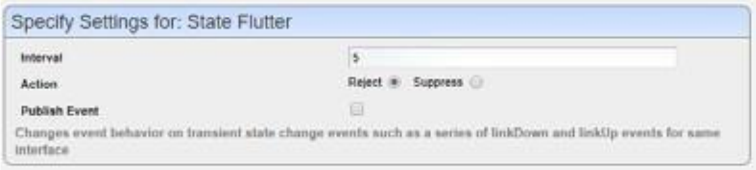

Specify Settings for: State Flutter

Interval: 5

Action: Reject ● Suppress ○

Publish Event: □

Changes event behavior on transient state change events such as a series of linkDown and linkUp events for same interface

To create a rule of this type, you will need to associate it with at least one raising and/or clearing event definition. You can include more than one raising and/or more than one clearing, but for an event to be affected by this type of rule, it must have a correlated pair to another event definition. To add event definitions to the rule, when you are on the Filter tab of the edit screen, click the Add button within the upper panel and select event definitions that correlate to each other through raising/clearin g correlation. If you need to look up this information, you can go to the Event Definitions portlet and bring up the edit screen for any given event definition and then navigate to the Correlations tab. Examples of pairs of event definitions that are related to each other through raising/clearing correlation include linkUp/ linkDown, monitorTargetUp/monitorTargetDown, among manyothers.

After you select the event definitions and also enter any additional filtering as desired, navigate to the Setting tab and enter the Interval (seconds) and the Action (Reject or Suppress the event). If you Reject an event, it does not appear in Event history; if you Suppress it then it creates no alarm, but it does appear in the Event history. Check Publish Event if you want Cruz to keep a record of when this rule starts and stops filtering events.

Cruz always publishes the first raising event matching the given filter criteria. When Cruz receives a correlated event (either the raise or the clear), this activates the State Flutter rule pattern, which will then expire after the given number of seconds has elapsed. Until the pattern expires, it holds all correlated rising and clearing events. This way if the state goes from raise to clear and back to raise in rapid succession, the result will be the final state after the number of seconds has elapsed, but without publishing the extra events and/or creating the alarms.

The following type is the **Protocol Translation rule type**:

Syslog Escalation —This screen presents the *Specify Event Filtering* portion of the screen without any *Settings* in the lower screen. Specify events to select. Then click *Next* to go to the *Escalation* tab.

Automation rules let you modify the *Specify Event Filtering* portion of the screen without any *Settings* in the lower screen. Specify events to select. Then click *Next* to go to the *Actions* tab. See Actions on page 302 for more about that feature.

Syslog Escalation Criteria

This tab of Syslog Event Rules lets you manage events based on matching text and configure messages in response to such matches.

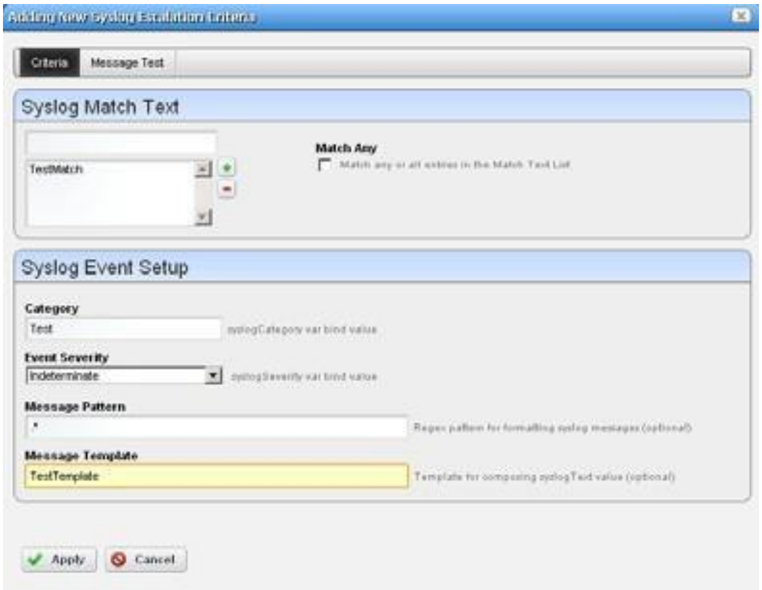

Syslog Match Text

TestMatch

Match Any
□ Match any or all entries in the Match Text List

Syslog Event Setup

Category
Test
syslogCategory var bind value

Event Severity
Indeterminate
syslogSeverity var bind value

Message Pattern
.*
Regex pattern for formatting syslog messages (optional)

Message Template
TestTemplate
Template for composing syslogText value (optional)

- The following sections describe syslog Criteria: Syslog Match Text
- Criteria: Syslog Event Setup
- Template Fields
- Message Test
- Explicit Members

Criteria: Syslog Match Text

In this tab, enter the Syslog Match Text. Click the plus to add matching text to the list below the empty field. There are two options for *How to Apply the Match Text to Syslog Messages*, which is controlled by whether or not this box is checked. Check to match any single entry (only one message match text must be present in the message) or uncheck to match all entries (all message match texts must be present). Forexample, consider that the Message Match Text list contains the following entries: "LOGIN" and "FAILED" and consider a Syslog message that says "USER LOGIN SUCCESS". This message would be a match if the box was checked but it would not be a match if the box was unchecked. If the list contains the same entries but the Syslog message was "USER LOGIN FAILED" then this would be a match regardless of whether the box was checked.

Criteria: Syslog Event Setup

This portion of the Criteria screen sets up the `syslogNotification` event emitted when matching occurs. Here are the fields:

Event Severity — Select the alarm severity of the event emitted when a match occurs. You can choose to assign a specific severity to the syslogNotification events that would be created from this escalation, such as Major, Minor, etc., or you can instead select Indeterminate, which will use the severity within the original Syslog message to determine the severityof the resulting event. For example, if you select Indeterminate and the Syslog message has a severity of Critical then the resulting event will also be Critical. But if you select a specific severity from this drop down, such as Minor, then this the events will be assigned this severity regardless of what severity is found within the original message.

Message Pattern —- An optional regular expression for the text to retrieve and transmit in the created event's variablebindings (varbinds). Syslog escalation uses the retrieved value(s) entered in the template fields to populate the associated varbinds.

Category Template —A directive for how to populate the syslog category varbind value. This is a template field, which means that you can either enter static text (like `Category1`) or a template containingvariables(like `Category%1`).Thispopulatesthe `syslogCategory` varbind with the appropriate text, for example: `Category-LOGIN`. See Template Fields below for more about this type of field.

NOTE:

Whenyoudynamicallypopulate thesyslog category, youcanmore easilybaseextended event definitions(EEDs)onthe `syslogNotification`definition. Forexampleyoucouldchangethebase eventdefinitiontoallowEEDsonthe syslogcategoryvarbind.SeeUsingExtendedEventDefinitions on page 333 for more about EEDs.

Message Template — A directive for how to populate the syslog text varbind value. This is a template field, which means that you can either enter static text (like `syslog message received`) or a template containing variables (like `%1 occurred on %3 for %2`). See the next topic for more information about Template Fields.

Suppress Alarm —Indicates whether or not to suppress the alarm for the resulting event. If you only want events to be created in the Event History but you do not want alarms to also be created, then check this box.

Template Fields

Templatefieldsareassociatedwithspecificvarbinds. Whenasyslogmessagematchesanescalation filter, Cruz creates an event and populates its varbinds using the respective templates and the *Message Pattern*.

Templates have numbered variables— `%1`, `%2`, and so on. Cruz resolves such variables with substrings extracted from the original message text. This means it inserts the first pattern retrieved in place of `%1`, inserts the second pattern retrieved for `%2`, and so on. For example: the *Message Template* field contains `%1 occurred on %3 due to %2`, the *Message Pattern* contains the regular expression `(.*): (.*). IP: (.*)` and a syslog message arrives that matches the *Syslog Match Text* with the contents: `Error: out of memory. IP: 192.168.0.1` then the message text varbind on the resulting event resolves to `Error occurred on 192.168.0.1 due to out of memory`. This works on other template fields too, like Category.

Cruz's syslogNotificationevent also includes avarbind containing theoriginalsyslog message. This can be useful if you want the syslogText varbind to be the product of processing, but you also want to see the original message.

Message Test

This screen lets you test your message against the pattern and/or template. Click the *Test* button to the right of the top field to activate this testing.

Test Message — Enter a message to test.

Test Message Result—The text extracted for the event as it appears in the template after you click the *Test* button.

Click *Apply* to accept these escalation criteria or *Cancel* to abandon them without saving.

NOTE:

The default behavior of the syslogNotification is Reject rather than Suppress. This means that received syslog messages must match an escalation filter to become events. Users who want all received syslog messages to become events must override this default setting and change the behavior to either Suppress (which makes only events but no alarms) or Alarm (which makes alarms). Note that this default behavior only affects messages that do not match any escalation filter. Redcell processes those that match an escalation filter just the same regardless of what is the default behavior of syslogNotification.

Explicit Members

This screen displays the explicit member entities that are associated with this rule. If any explicit members are defined for a rule, then the event must be associated with one of these entities for this rule to be triggered. This is a convenient way of defining entity-specific filters for the rule and you can also use this feature to allow the rule to behave differently based on the entity that is associated with the triggering event. Please note that the Explicit Membership box on the *General* tab must be checked in order to enable the buttons on this tab.

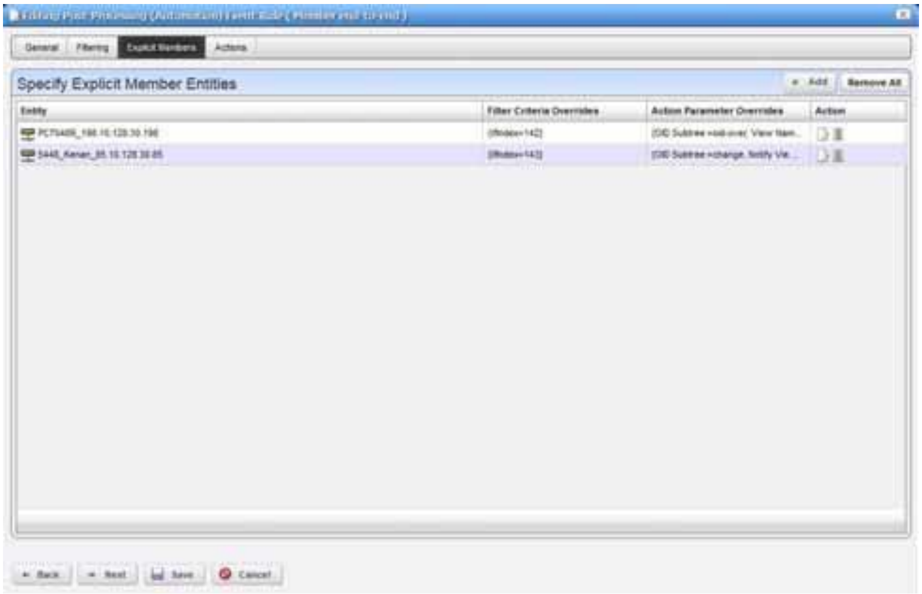

Specify Explicit Member Entities

Entity	Filter Criteria Overrides	Action Parameter Overrides	Action
PCT5405_198.10.128.30.198	{ifIndex=142}	{OID Subtree =oid-over(View Nam...	
5448_Kenan_85.10.128.30.85	{ifIndex=143}	{OID Subtree =change, Notify Vie...	

If explicit membership is enabled, then you can click *Add* and this displays a screen through which you can select an explicit member entity and optionally configure specific overrides that would be associated with this member.

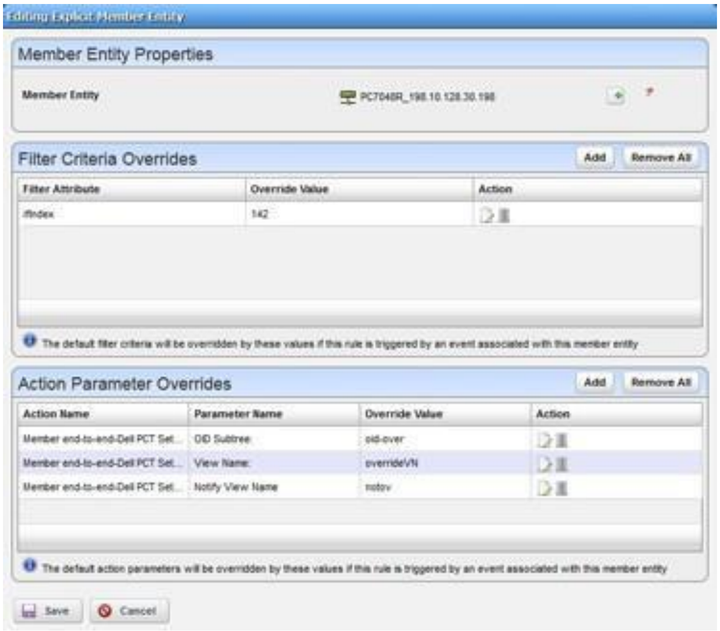

Editing Explicit Member Entity

Member Entity Properties

Member Entity: PC7048R_198.10.128.30.198

Filter Criteria Overrides

Filter Attribute	Override Value	Action
ifIndex	142	

The default filter criteria will be overridden by these values if this rule is triggered by an event associated with this member entity

Action Parameter Overrides

Action Name	Parameter Name	Override Value	Action
Member end-to-end-Dell PCT Set...	OID Subtree:	oid-over	
Member end-to-end-Dell PCT Set...	View Name:	overrideVN	
Member end-to-end-Dell PCT Set...	Notify View Name	notov	

The default action parameters will be overridden by these values if this rule is triggered by an event associated with this member entity

From the *Editing Explicit Member Entity* popup screen, you can select a **Member Entity** from among different possible entity types including Managed Equipment, among others. Youcan also specify **Filter Criteria** Overrides, which are the changes to the filter criteria that should be applied when thisrule is triggered byan event associated with the Member Entity that wasselected. Finally, you can specify **Action Parameter Overrides**, which are changes to the action parameters that should be applied when this rule is triggered by an event associated with the member entity that was selected. If you choose to add an override of either type, a screen will be shown that allows you to specify the attributes that are being overridden along with the override value.

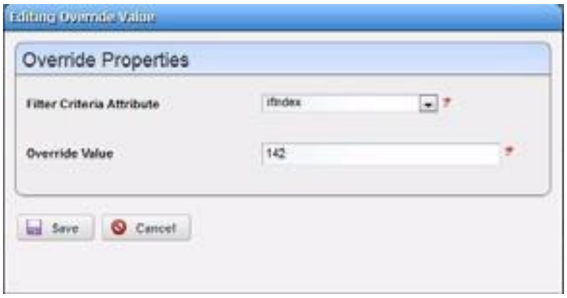

Editing Override Value

Override Properties

Filter Criteria Attribute: ifIndex

Override Value: 142

On the *Editing Override Value* screen, if this is a filter criterion override then a drop-downlist will be shown to select the **Filter Criteria Attribute.** Note that this list will be populated with the filter criteria that this rule is configured with from the *Filtering* tab. If you select an attribute and then enter a value, then this value will override the default configured value on the *Filtering* tab if this rule is triggered by an event associated with the configured member. The filter will then be applied with the override values and this rule will only be triggered if these criteria are met.

If this is an action parameter override, then drop down lists will be shown to select the action and the action parameter. Note that these lists will be populated with the names of the items in the *Actions* tab and with the parameters associated with each respective action. If you select an action and an action parameter and then enter a value, then this value will override the default configured value on the *Actions* tab if this rule is triggered by an event associated with the configured member. The action will then be executed with the override values.

Actions

This screen catalogs the actions configured for the Automation rule you previously configured.

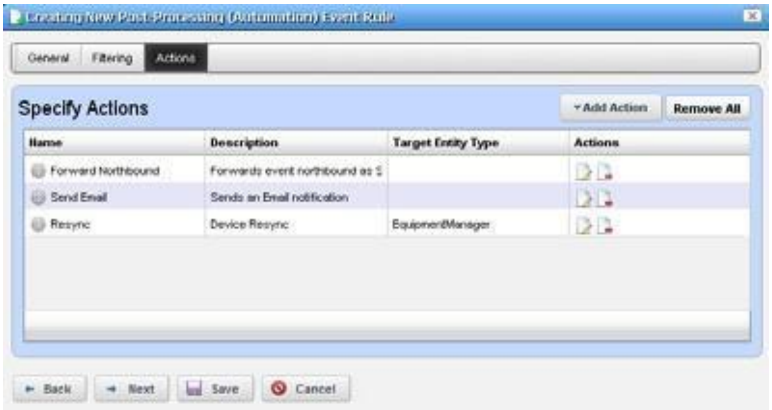

Name	Description	Target Entity Type	Actions
Forward Northbound	Forwards event northbound as S		
Send Email	Sends an Email notification		
Resync	Device Resync	EquipmentManager	

Click *Add Action* to create a new action in the editor. The *Actions* column lets you revise (*Edit this entry*) or *Delete* entries in this table. Click *Save* to preserve the actions configured here or *Cancel* to abandon any edits.

Clicking *Add Action* to configure:

- an action based on available Adaptive CLI actions
- destinations and messages for email and SMS recipients
- automation to forward northbound messages
- config changed

Click *Apply* to accept configured actions or *Cancel* to abandon their editor and return to this screen.

NOTE:

Actions available here are like those for Discovery Profiles on page 163. You can also use actions to Execute Proscan. See.

Action

This screen lets you configure an action based on Adaptive CLI actions available in the system. Select an action. To search for an action, specify search criteria and then click Go. The parameters available vary depending on the selected action. Tooltips provide additional information when you hover over a field.

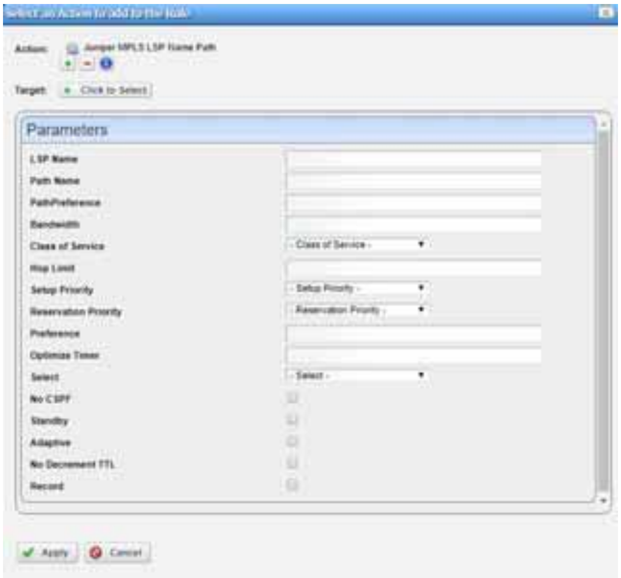

Select an Action to add to the Rule

Action: Juniper MPLS LSP Name Path

Target: Click to Select

Parameters

LSP Name

Path Name

PathPreference

Bandwidth

Class of Service: - Class of Service -

Hop Limit

Setup Priority: - Setup Priority -

Reservation Priority: - Reservation Priority -

Preference

Optimize Timer

Select: - Select -

No CSPF □

Standby □

Adaptive □

No Decrement TTL □

Record □

Optionally select the action device target from the list. If you do not specifyan explicit target, Cruz uses the default entity for the event as the target.

Youcan specify parameter variables, dependent on the event, rule, and selected targets specifics. Do this with either NOTIFICATION or VARBIND.

The following are valid attributes to use in a phrase like `[[NOTIFICATION: <attr name>]]:`

- TypeOID
- AlarmOID
- EntityOID
- EquipMgrOID
- DeviceIP
- SourceIP
- EntityName

NOTE:

ConsulttherelevantportlettofindandverifyanOID.Forexample,theEventDefinitionsportlethasan OID column, and the varbind OIDs appear in the event editor *Message Template* window.

Correctspelling is mandatory, and thesearecasesensitive. NOTIFICATION and VARBINDmust be all caps, and within double brackets. The colon and space after the key word are required.

Cruz converts anything that conforms to these rules and then passes the converted information into the action before execution. Anything outs ide the double square brackets passes verbatim.

For example, if you have the following string:

```
This is the alarm OID [[NOTIFICATION: AlarmOID]] of notification type
 [[NOTIFICATION: TypeOID]] having variable binding [[VARBIND: 1.3.4.5.3]]
```

This string is passed as:

```
This is the alarm OID 1OiE92tUj113G03 of notification type
 1.3.6.1.4.1.3477.1.27.20.7 having variable binding 151.
```

Click *Apply* to accept your edits or *Cancel* to abandon them. For an example, see Extracting Adaptive CLI Attributes from a Syslog Alarm on page 314.

Email

Email actions configure destinations and messages for e-mail and SMS recipients. You can include fields that are part of the event by using the features described in Using Email Action Variables on page 310.

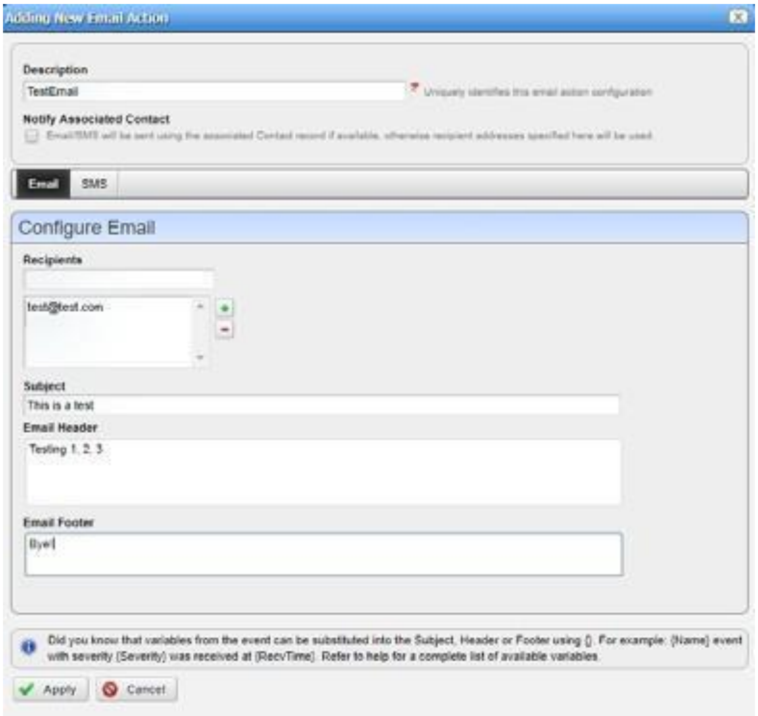

Adding New Email Action

Description: TestEmail

Uniquely identifies this email action configuration

Notify Associated Contact

☐ Email/SMS will be sent using the associated Contact record if available, otherwise recipient addresses specified here will be used.

Configure Email

Recipients: test@test.com

Subject: This is a test

Email Header: Testing 1, 2, 3

Email Footer: Bye!

Did you know that variables from the event can be substituted into the Subject, Header or Footer using {}. For example: {Name} event with severity {Severity} was received at {RecvTime}. Refer to help for a complete list of available variables.

Notice that below the e-mail action Description, you can check to send this mail (and/or SMS) to associated Contacts, if any are available, even if you specify no mail address destination. The SMS tab is like the e-mail tab but limits the number of characters you can enter with a field at its bottom. You must send SMS to the destination phone carrier’s e-mail-to-SMS address. For example, sending text to 916-555-1212 when Verizon is the carrier means the destination address is 9165551212@vtext.com.

When enabled, notification emails go to the Contact associated with the Managed Equipment for the notification event. For the contact's email address, mail goes to the first specified address from either the *Work Email*, *Home Email* or *Other Email* fields in the Contact editor. SMS messages go to the *Pager Email* field for the contact. If a Contact was not found or the required addresses are not specified for the Contact, then Cruz uses the Recipient addresses configured in the Email Action.

NOTE:

Programs other than Cruz let you manipulate mail outside the scope of Cruz. For example, IFTTT (If This Then That) lets you send SMS in countries whose providers do not provide e-mail equivalents to SMS addressing. You can also use such applications to save mail attachments like reports to Dropbox accounts.

This screen has the following fields:

Recipient —Enter an e-mail address in this field and then click add (+) to add a recipient's address. Click delete (-) to remove selected recipients.

Subject — The e-mail subject.

Email Header/Footer — The e-mail's heading and footing.

SMS Body — The e-mail contents to be sent as text.

SMS Max Length—The maximum number of characters to send in the SMS. Typically, this is 140, but the default is 0, so be sure to set to your carrier's maximum before saving.

Here is what an Email looks like when it arrives:

```
Sent: Wednesday, March 02, 2011 2:37 PM
To: techpubs@testsoftware.com
Subject: Web Test
Notification: redcellInventoryAttribChangeNotification
Notification Attributes:
================================
sysUpTime.0                     = 5 hours, 16 mins, 43 secs
snmpTrapOID.0                   = 1.3.6.1.4.1.3477.2.2.1
redcellInventoryAttrName.0      = RedCell.Config.EquipmentManager_Notes
redcellInventoryAttrChangedBy.0 = admin
redcellInventoryAttrNewValue.0  = hello
world
severity
auto
redcellInventoryAttrOldValue.0  = hello
world
severity
```

Forward Northbound

When you want to forward a message, such as an SNMP v2 event (trap) to another host, then configure automation in this screen to do that.

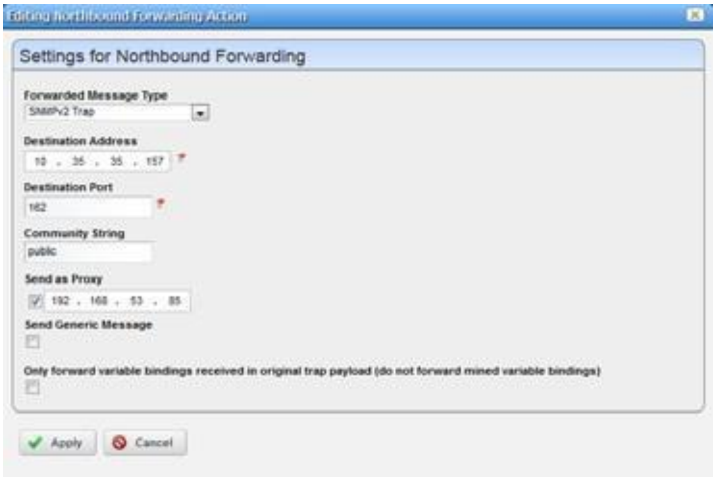

Editing Northbound Forwarding Action

Settings for Northbound Forwarding

Forwarded Message Type: SNMPv2 Trap

Destination Address: 10 . 35 . 35 . 157 *

Destination Port: 162 *

Community String: public

Send as Proxy: ☑ 192 . 168 . 53 . 85

Send Generic Message: □

Only forward variable bindings received in original trap payload (do not forward mined variable bindings): □

Enter the following fields:

Forwarded Message Type — The type of message to forward. This can be an SNMP v2 Trap, a Syslog, or an HTTP REST call. This selection is determined by the type of message the northbound system is expecting. Forexample, you might have a northbound system that is expecting to collect SNMP traps from Cruz. Some northbound systems are configured to collect Syslog messages and other might be configured to collect HTTP REST calls.

Destination Address — The IP address of the northbound destination.

Destination Port — The port on the northbound destination. Note that this auto-populates with the most common port number based on the selected Forwarded Message Type.

Community String — The SNMP community string for the northbound destination. This is only relevant if the Forwarded Message Type is SNMP Trap.

Username — If the Forwarded Message Type is HTTP REST call, the username is required for authentication.

Password — If the Forwarded Message Typeis HTTP REST call, the password is required for authentication.

Send as Proxy—Select this option to send the IP address you specify as the event source. If not selected, the IP address of the mediation server receiving the trap is sent. (See Sending as Proxy on page 307 for more information.)

Send Generic Message —When selected, this sends a generic message, whether this be a trap or a different kind of message. See Forwarding a Generic Message on page 308 for more information.

Only forward variable bindings received in original trap payload (do not forward mined variable bindings) — If selected, only the variable bindings received in the original trap payload is forwarded. If not selected, the mined variable bindings are also forwarded. If the Forwarded Message Type is Syslog, this is irrelevant because Syslog messages do not contain variable bindings.

Config Changed

This action sets a flag on the device notifying that a configuration change has occurred. This also updates the"Last Configuration Change" attribute of the managed resource that is associated with the event that triggered the automation to execute.

Hint: Youcan generate the report named "Configuration Change Report" to see all configuration changes that have occurred for each device.

Forwarding Traps

SNMPv1 and SNMPv3 traps become SNMPv2 Traps

SNMPv1 traps are converted according to the RFC 1908 specification. SNMPv3 trapsare alreadyin the SNMPv2 format and the application simply does not use SNMPv3 security when sending these northbound. The following is the relevant snippet from the RFC 1908 specification:

```
3.1.2. SNMPv1 -> SNMPv2

When converting responses received from a SNMPv1 entity acting in an agent
  role into responses sent to a SNMPv2 entity acting in a manager role:

(1) ...

(2) If a Trap-PDU is received, then it is mapped into a SNMPv2-Trap-PDU.
  This is done by prepending onto the variable-bindings field two new
  bindings: sysUpTime.0 [6], which takes its value from the timestamp
  field of the Trap-PDU; and, snmpTrapOID.0 [6], which is calculated as
  follows: if the value of generic-trap field is enterpriseSpecific,
  then the value used is the concatenation of the enterprise field from the
  Trap-PDU with two additional sub- identifiers, '0', and the value of the
  specific-trap field; otherwise, the value of the corresponding trap
  defined in [6] is used. (For example, if the value of the generic-trap
  field is coldStart, then the application uses the coldStart trap [6])
  Then, one new binding is appended onto the variable-bindings field:
  snmpTrapEnterprise.0 [6], which takes its value from the enterprise
  field of the Trap-PDU. The destinations for the SNMPv2-Trap-PDU are
  determined in an implementation-dependent fashion by the proxy agent.
```

Despite this description, many vendors defined a trap for SNMPv2 and then had to support sending it as S NMPv1 protocol. The assembly of v 2 OID from v1 enterprise and specific is supposed to include an extra zero (0) as in: enterpriseOID.0.specific. However, if a v2 trap is defined that has no '0' in it, it cannot be sent as v1 an d converted back following the specifications.

Sending as Proxy

The application can forward a trap as though it came from a device (sourceIP spoofing) or act as an agent proxy according to the SNMP-COMMUNITY -MIB. If not sending as a proxy, the Cruz application forwards traps from an application server cluster as an SNMPv2 notification as though it is coming directly from the originating agent
(device). This is a common and the desired behavior. Some operating systems prevent packet spoofing as a security measure, so this behavior is optional.

If sending as a proxy, the Cruz application forwards the trap from the IP address given in the field adjacent to the Send as Proxy option when you select it. If you select the Send as Proxy option but leave the adjacent IP address blank, the Cruz application forwards the t rap from the receiving mediation server as sourceIP.

The relevant excerpt from SNM P-COMMUNIT Y-MIB is:

```
--
```

```
 -- The snmpTrapAddress and snmpTrapCommunity objects are included
 -- in notifications that are forwarded by a proxy, which were
 -- originally received as SNMPv1 Trapmessages.
 --


snmpTrapAddress OBJECT-TYPE
        SYNTAX  IpAddress
        MAX-ACCESS accessible-for-notify
        STATUS current
        DESCRIPTION
               "The value of the agent-addr field of a Trap PDU which
               is forwarded by a proxy forwarder application using
               an SNMP version other than SNMPv1.  The value of this
               object SHOULD contain the value of the agent-addr field
               from the original Trap PDU as generated by an SNMPv1
               agent."
  -- 1.3.6.1.6.3.18.1.3 --  ::= { snmpCommunityMIBObjects 3 }

 snmpTrapCommunity OBJECT-TYPE
        SYNTAX OCTET STRING
        MAX-ACCESS accessible-for-notify
        STATUS current
        DESCRIPTION
               "The value of the community string field of an SNMPv1
               message containing a Trap PDU which is forwarded by a
               a proxy forwarder application using an SNMP version
               other than SNMPv1. The value of this object SHOULD
               contain the value of the community string field from
               the original SNMPv1 message containing a Trap PDU as
               generated by an SNMPv1 agent."
  -- 1.3.6.1.6.3.18.1.4 --  ::= { snmpCommunityMIBObjects 4 }
```

The Cruz application alwaysadds `snmpTrapAddress` to every trap forwarded asa proxy,(never adding `snmpTrapCommunity`). It does not keep track of the community string on the traps received.

Forwarding a GenericMessage

The option to forward a generic message forwards a trap or other type of message is found in the following MIB file:

```
../owareapps/eventmgmt/mibs/AssureAlarms-MIB
```

The following is the definitio n, as found in this file:

```
redcellGenericTrap NOTIFICATION-TYPE
```

```
OBJECTS { alarmOID, referencedNotificationTypeOID,
referencedNotificationName, redcellSeverity, redcellEquipmentName,
redcellEquipmentManagerOID, redcellInventoryEntityName,
redcellInventoryEntityType, alarmMessage, redcellNotificationReceivedTime,
redcellEquipmentIPAddress }
    STATUS current
DESCRIPTION
"Generic trap used for forwarding information about another trap while using a
standard trap format and notification type OID."
::= { defaultProcessing 6 }
```

There is an XML version of this definition in the following file:

```
.../owareapps/eventmgmt/server/conf/mibs.xml
```

Here is the XML content, as found in this file:

```
<item>
    <Description>Generic trap used for forwarding information about another
      trap while using a standard trap format and notification type OID.</
      Description>
    <IndexPosition>1</IndexPosition>
    <Name>redcellGenericTrap</Name>
    <OID>1.3.6.1.4.1.3477.1.7.11.6</OID>
    <Status>current</Status>
    <Type>NOTIFICATION-TYPE</Type>
    <Variables>
       <item>
          <Name>alarmOID</Name>
          <OID>1.3.6.1.4.1.3477.1.7.11.1</OID>
       </item>
       <item>
          <Name>referencedNotificationTypeOID</Name>
          <OID>1.3.6.1.4.1.3477.1.7.11.3</OID>
       </item>
       <item>
          <Name>referencedNotificationName</Name>
          <OID>1.3.6.1.4.1.3477.1.7.11.4</OID>
       </item>
       <item>
          <Name>redcellSeverity</Name>
          <OID>1.3.6.1.4.1.3477.1.6.1</OID>
       </item>
       <item>
          <Name>redcellEquipmentName</Name>
          <OID>1.3.6.1.4.1.3477.2.3.1</OID>
       </item>
       <item>
          <Name>redcellEquipmentManagerOID</Name>
          <OID>1.3.6.1.4.1.3477.2.3.14</OID>
       </item>
       <item>
```

```
          <Name>redcellInventoryEntityName</Name>
          <OID>1.3.6.1.4.1.3477.2.1.8</OID>
      </item>
      <item>
          <Name>redcellInventoryEntityType</Name>
          <OID>1.3.6.1.4.1.3477.2.1.5</OID>
      </item>
      <item>
          <Name>alarmMessage</Name>
          <OID>1.3.6.1.4.1.3477.1.7.11.2</OID>
      </item>
      <item>
          <Name>redcellNotificationReceivedTime</Name>
          <OID>1.3.6.1.4.1.3477.1.6.3</OID>
      </item>
      <item>
          <Name>redcellEquipmentIPAddress</Name>
          <OID>1.3.6.1.4.1.3477.2.3.8</OID>
      </item>
    </Variables>
    <ViewType>OBJECT</ViewType>
</item>
```

Using Email ActionVariables

The following are the Email Action variables used to customize the action email content. These variables are classified as follows:

- Basic Variables
- Managed Equipment Variables
- Entity Type: Port
- Entity Type: Interface, Logical interface

Tosuccessfully retrieve custom attributes in an e-mail, you must first create them. See Editing Custom Attributes on page 146.

Youcan also configure more limited variables that are slightly more efficient in performance, if not as detailed as those described in the following section.

Forexample, you can retrieve the followingattributes:

```
{RedCell.Config.EquipmentManager_Custom1}
{RedCell.Config.EquipmentManager_Custom2}
{RedCell.Config.EquipmentManager_LastBackup}
{RedCell.Config.EquipmentManager_LastConfigChange} and
{RedCell.Config.EquipmentManager_HealthStatus}
```

NOTE:

If the entity does not contain/return these values, then the message [No data for <attribute name>] appears in the email instead.

Best practice is to clarify such attributes by combining them with others that spell out their source.

Basic Variables

Here is a description of the basic variables:

Attribute	Description	Email Action Variable
Name	The event/alarm name	{Name}
Message	Description from the event	{Message}
Entity Name	The entity (interface, card...) name	{EntityName}
Equipment Manager Name	The name of the equipment, parent or chassis.	{EquipMgrName}
Device IP address	the IP of the device in alarm	{DeviceIP}
Entity Type	Type of entity (Router, and so on)	{EntityType}
Instance ID	An identifier for the event	{InstanceID}
Protocol Type	Of originating alarm (SNMP, syslog, etc.)	{ProtocolType}
Protocol Sub Type	Inform, Trap, [blank] (for internal events)	{ProtocolSubType}
Receive Time		{RecvTime}
Region	The mediation server partition name.	{Region}
Severity	0 - cleared, through 6 - critical, from Alarm Definition	{Severity}
Source IP address	The IP of the component sending the alarm	{SourceIP}

Managed Equipment Variables

Here is a description of the managed equipment variables. These variables may havea performance impact

Attribute	Description	Email Action Variable
Custom 1	Note that although you can re-name any Custom attribute, you must use the variable’s original name. For example, here, that is {RedCell.Config.EquipmentManager _Custom1}	{RedCell.Config.EquipmentManager _Custom1}
Custom 2		{RedCell.Config.EquipmentManager _Custom2}
Custom 3		{RedCell.Config.EquipmentManager _Custom3}
Custom 4		{RedCell.Config.EquipmentManager _Custom4}
Custom 5		{RedCell.Config.EquipmentManager _Custom5}
Custom 6		{RedCell.Config.EquipmentManager _Custom6}
Custom 7		{RedCell.Config.EquipmentManager _Custom7}

Attribute	Description	Email Action Variable
Custom 8		{RedCell.Config.EquipmentManager_Custom8}
Custom 9		{RedCell.Config.EquipmentManager_Custom9}
Custom 10		{RedCell.Config.EquipmentManager_Custom10}
Custom 11		{RedCell.Config.EquipmentManager_Custom11}
Custom 12		{RedCell.Config.EquipmentManager_Custom12}
Custom 13		{RedCell.Config.EquipmentManager_Custom13}
Description	Description of the equipment	{RedCell.Config.EquipmentManager_DeviceDescription}
DNS Hostname	Hostname of equipment	{RedCell.Config.EquipmentManager_Hostname}
Equipment Type	Equipment Type	{RedCell.Config.EquipmentManager_CommonType}
Firmware Version	Version of the equipment’s firmware	{RedCell.Config.EquipmentManager_FirmwareVersion}
Hardware Version	Version of the equipment’s hardware	RedCell.Config.EquipmentManager_HardwareVersion}
Last Backup	Last Backup	{RedCell.Config.EquipmentManager_LastBackup}
Last Configuration Change	Last Configuration Change	{RedCell.Config.EquipmentManager_LastConfigChange}
Last Modified	Timestamp of Last Modified	{RedCell.Config.EquipmentManager_LastModified}
Model	Model number of the equipment	{RedCell.Config.EquipmentManager_Model}
Name	Component name	{RedCell.Config.EquipmentManager_Name}
Network Status	Network Status	{RedCell.Config.EquipmentManager_HealthStatus}
Notes	Equipment Notes	{RedCell.Config.EquipmentManager_Notes}
OSVersion	OSVersion	{RedCell.Config.EquipmentManager_OSVersion}
Serial Number	Unique identifier for the equipment	{RedCell.Config.EquipmentManager_SerialNumber}
Software Version	Version of the equipment’s software	{RedCell.Config.EquipmentManager_SoftwareVersion}
System Object Id	SNMP based system object identifier	{RedCell.Config.EquipmentManager_SysObjectID}

Entity Type: Port

Here is a description of the Port entity type variables.

Attribute	Description	Email Action Variable
Custom 1	Note that although you can re-name any Custom attribute, you must use the variable's original name. For example, that is {RedCell.Config.EquipmentManager _Custom1}	{RedCell.Config.Port_Custom1}
Custom 2		{RedCell.Config.Port_Custom2}
Custom 3		{RedCell.Config.Port_Custom3}
Custom 4		{RedCell.Config.Port_Custom4}
Encapsulation	Encapsulation	{RedCell.Config.Port_Encapsulation}
Hardware Version	Version of the port's hardware	{RedCell.Config.Port_HardwareVersion}
If Index	SNMP If Index	{Re dCell.Config.Port_IfIndex}
MAC Address	"Typically, a MAC Address, with the octets separated by a space, colon or dash depending upon the device. Note that the separator is relative when used as part of a query."	{RedCell.Config.Port_UniqueAddress}
Model	Model number of the port	{RedCell.Config.Port_Model}
MTU	Maximum Transmission Unit	{RedCell.Config.Port_Mtu}
Name	Port name	{RedCell.Config.Port_Name}
Notes	Port Notes	{RedCell.Config.Port_Notes}
Port Description	Description of the port	{RedCell.Config.Port_DeviceDescrip tion}
Port Number	Port Number	{RedCell.Config.Port_PortNumber}
Slot Number	Slot Number	{RedCell.Config.Port_SlotNumber}
Speed	Speed	{RedCell.Config.Port_Speed}
Subnet Mask	SubMask	{RedCell.Config.Port_SubMask}

Entity Type: Interface, Logical interface

Here is a description of the Interface, Logical interface entity types

Attribute	Description	Redcell Email Action variable
Custom 1	Note that although you can re-name any Custom attribute, you must use the variable's original name. For example, that is {RedCell.Config.EquipmentManager _Custom1}	{RedCell.Config.Interface_Custom1}
Custom 2		{RedCell.Config.Interface_Custom2}
Custom 3		{RedCell.Config.Interface_Custom3}
Custom 4		{RedCell.Config.Interface_Custom4}
Encapsulation	Encapsulation	{RedCell.Config.Interface_Encapsula tion}

Attribute	Description	Redcell Email Action variable
IfIndex	SNMP Interface Index	{RedCell.Config.Interface_IfIndex}
Interface Description	Description of the Interface	{RedCell.Config.Interface_DeviceDescription}
Interface Number	Interface Number	{RedCell.Config.Interface_InterfaceNumber}
Interface Type	Common Interface Type	{RedCell.Config.Interface_CommonType}
MTU	Maximum Transmission Unit	{RedCell.Config.Interface_Mtu}
Name	Interface name	{RedCell.Config.Interface_Name}
Notes	Interface Notes	{RedCell.Config.Interface_Notes}
Port Number	Port Number	{RedCell.Config.Interface_PortNumber}
Slot Number	Slot Number	{RedCell.Config.Interface_SlotNumber}
Subnet Mask	Subnet Mask of the Interface	{RedCell.Config.Interface_SubMask}

Extracting Adaptive CLI Attributes from a Syslog Alarm

This section demonstrates how to pass attribute text from a syslog alarm to an action by detailing the following workflow:

1. Receiving Syslog Events
2. Matching Text from a syslogNotification Event
3. Creating Extended Event Definition for syslogNotification
4. Creating ACLI Action and Pass Varbind Value for Execution
5. Verifying the Results

Receiving Syslog Events

To receive syslog events, you need to configure a device to send syslog traps to the application.

Receive syslog events byconfiguring your device to send syslog traps to the Cruz application as follows.

1. Configure a device to send syslog to the appserver.

   For a Dell S5000 switch (10.128.6.22), add `logging 192.168.0.138`.

   ```
   !
   logging trap debugging
   logging 192.168.54.43
   logging 10.35.35.201
   logging 10.35.35.78
   logging 192.168.50.200
   logging 192.168.0.138
   ```

Verifythat the Event History portlet shows the SyslogNotificationevent name for the defined entity.

The syslogNotification default behavior is suppress. Because we receive many syslogNotifications, you can set the event to Reje ct after confirming the use case. For Cisco, the syslog event may use clogMessageGenerated.

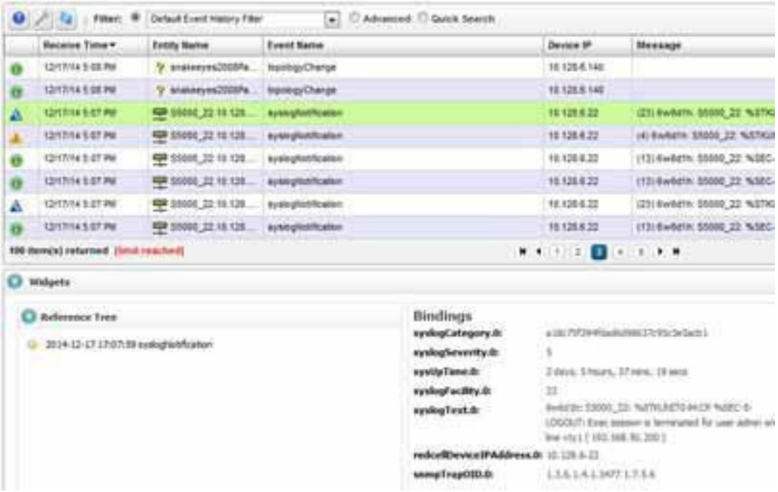

Filter: ● Default Event History Filter ○ Advanced ○ Quick Search

	Receive Time ▾	Entity Name	Event Name	Device IP	Message
	12/17/14 5:08 PM	snakeeyes2008Fa...	topologyChange	10.128.6.140	
	12/17/14 5:08 PM	snakeeyes2008Fa...	topologyChange	10.128.6.140	
	12/17/14 5:07 PM	S5000_22 10.128...	syslogNotification	10.128.6.22	(23) 6w6d1h: S5000_22: %STK
	12/17/14 5:07 PM	S5000_22 10.128...	syslogNotification	10.128.6.22	(4) 6w6d1h: S5000_22: %STKU
	12/17/14 5:07 PM	S5000_22 10.128...	syslogNotification	10.128.6.22	(13) 6w6d1h: S5000_22: %SEC-
	12/17/14 5:07 PM	S5000_22 10.128...	syslogNotification	10.128.6.22	(13) 6w6d1h: S5000_22: %SEC-
	12/17/14 5:07 PM	S5000_22 10.128...	syslogNotification	10.128.6.22	(23) 6w6d1h: S5000_22: %STK
	12/17/14 5:07 PM	S5000_22 10.128...	syslogNotification	10.128.6.22	(13) 6w6d1h: S5000_22: %SEC-

100 item(s) returned (limit reached)

2014-12-17 17:07:59 syslogNotification

Bindings

syslogCategory.0:	[illegible]
syslogSeverity.0:	5
sysUpTime.0:	2 days, 5 hours, 37 mins, 19 secs
syslogFacility.0:	23
syslogText.0:	6w6d1h: S5000_22: %STKUNIT0-M:CP %SEC-5-LOGOUT: Exec session is terminated for user admin on line vty1 (192.168.81.200)
redcellDeviceIPAddress.0:	10.128.6.22
snmpTrapOID.0:	1.3.6.1.4.1.3477.1.7.5.6

MatchingTextfrom asyslogNotification Event

This example matches specific text from an event and associates it to a category by creating a Syslog Escalation rule in the Event Processing Rule portlet.

Match text from a syslogNotification event as follows.

1 Extract the port information from a syslog message.

```
S5000> configure t
```

2 Log in Device and

```
> interface te 0/20
```

3 Shut down an interface:

```
> shutdown
```

4 Bring up the interface

```
> no shutdown
```

5 Locate the `syslogNotification` event in the Event History portlet.

6 Look at the syslogText provided from the `shutdown` and `no shutdown` commands.

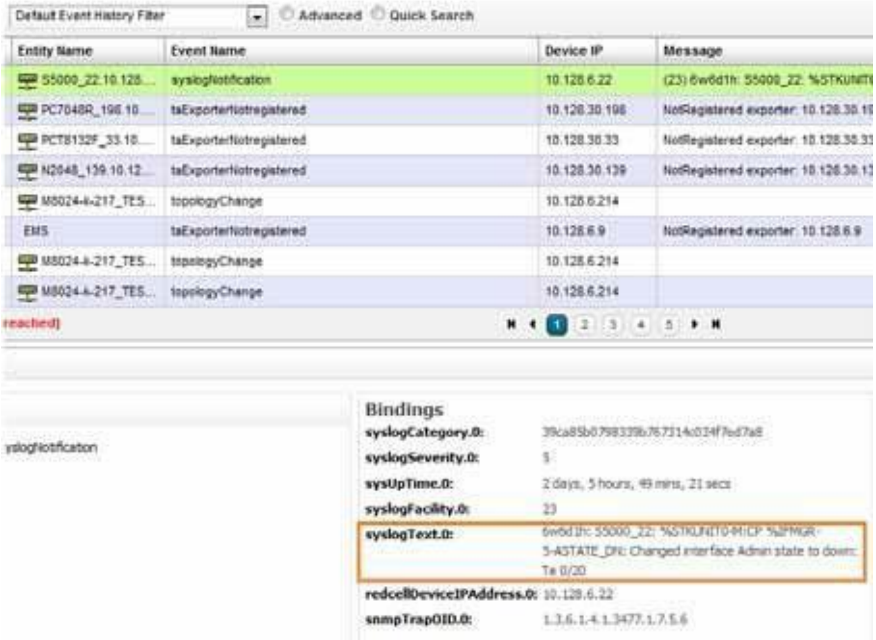

Entity Name	Event Name	Device IP	Message
S5000_22.10.128...	syslogNotification	10.128.6.22	(23) 6w6d1h: S5000_22: %STKUNIT
PC7048R_198.10...	taExporterNotregistered	10.128.30.198	NotRegistered exporter: 10.128.30.19
PCT8132F_33.10...	taExporterNotregistered	10.128.30.33	NotRegistered exporter: 10.128.30.33
N2048_139.10.12...	taExporterNotregistered	10.128.30.139	NotRegistered exporter: 10.128.30.13
M8024-k-217_TES...	topologyChange	10.128.6.214	
EMS	taExporterNotregistered	10.128.6.9	NotRegistered exporter: 10.128.6.9
M8024-k-217_TES...	topologyChange	10.128.6.214	
M8024-k-217_TES...	topologyChange	10.128.6.214	

Bindings

syslogCategory.0:	39ca85b0798339b767314c034f7ed7a8
syslogSeverity.0:	5
sysUpTime.0:	2 days, 5 hours, 49 mins, 21 secs
syslogFacility.0:	23
syslogText.0:	6w6d1h: S5000_22: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to down: Te 0/20
redcellDeviceIPAddress.0:	10.128.6.22
snmpTrapOID.0:	1.3.6.1.4.1.3477.1.7.5.6

7 Extract the relevant port information with these messages.

```
syslogText.0: 6w6d1h: S5000_22: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN:
 Changed interface Admin state to down: Te 0/20
```

```
syslogText.0: 6w6d1h: S5000_22: %STKUNIT0-M:CP %IFMGR-5-ASTATE_UP:
 Changed interface Admin state to up: Te 0/20
```

8 Right-click to create a new Protocol Translation > Syslog rule in the Event Processing Rules portlet to extract information.

9 Make sure that the filtering tab is not blank as you edit.

For this example set `Source IP is not 0.0.0.0.`

10 Add the Escalation criteria extracting the following Syslog information for the `Te 0/20` interface:

```
syslogText.0: 6w6d1h: S5000_22: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN:
 Changed interface Admin state to down: Te 0/20
```

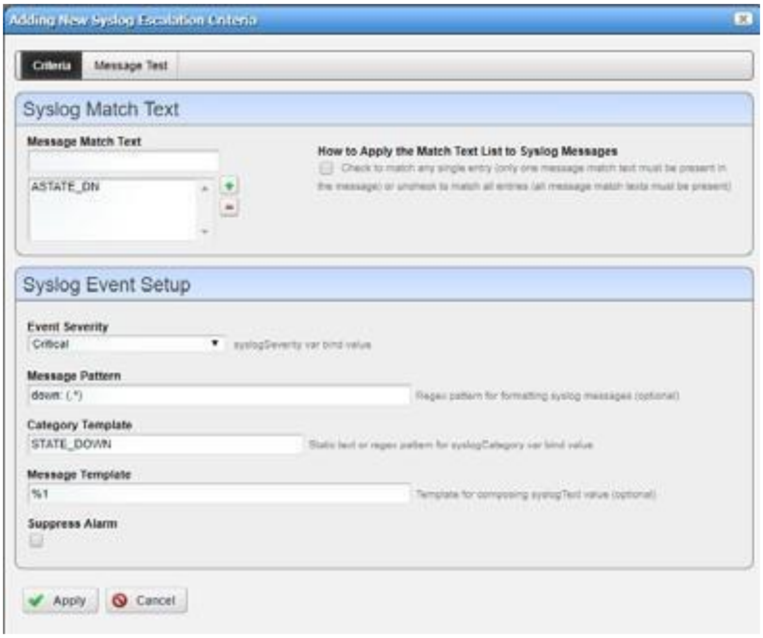

Adding New Syslog Escalation Criteria

Criteria | Message Test

Syslog Match Text

Message Match Text: ASTATE_DN

How to Apply the Match Text List to Syslog Messages

□ Check to match any single entry (only one message match text must be present in the message) or uncheck to match all entries (all message match texts must be present)

Syslog Event Setup

Event Severity: Critical — syslogSeverity var bind value

Message Pattern: down: (.*) — Regex pattern for formatting syslog messages (optional)

Category Template: STATE_DOWN — Static text or regex pattern for syslogCategory var bind value

Message Template: %1 — Template for composing syslogText value (optional)

Suppress Alarm: □

11 Provide the following information:

- **Message Match Text** use `ASTATE_DN` from the `syslogText`.
- **Event Severity** is the alarm severity for `syslogNotification::[Category]`
- **Message Pattern** get from the given `syslogText.0`, parse the port string (`Te 0/20`) with `down: (.*)`
- **Category** Template set up an event definition based on the Category name. The next step creates an Extended Event Definition for this. For example, the next time this type of syslog alarm arrives, it appears as `syslogNotification::[Category]` in this case `syslogNotification::STATE_DOWN`.
- **Message Template** is the regular expression in this rule parses *Message Pattern*, %1 from `(.*)` into `syslogText`

12 Apply the STATE_DN criteria.

13 Parse the `ASTATE_UP` match text similar to:

```
syslogText.0: 6w6d1h: S5000_22: %STKUNIT0-M:CP %IFMGR-5-ASTATE_UP:
Changed interface Admin state to up: Te 0/20
```

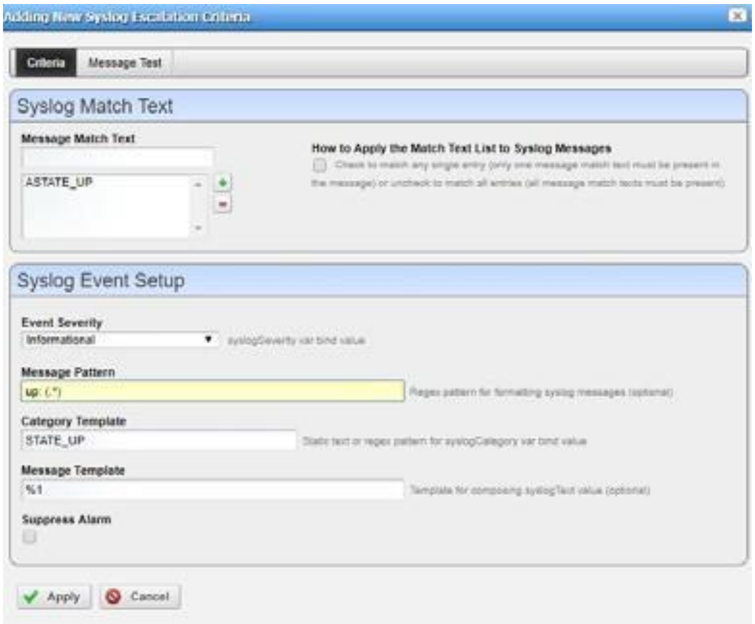

Adding New Syslog Escalation Criteria

Syslog Match Text

Message Match Text: ASTATE_UP

How to Apply the Match Text List to Syslog Messages

□ Check to match any single entry (only one message match text must be present in the message) or uncheck to match all entries (all message match texts must be present)

Syslog Event Setup

Event Severity: Informational — syslogSeverity var bind value

Message Pattern: up: (.*) — Regex pattern for formatting syslog messages (optional)

Category Template: STATE_UP — Static text or regex pattern for syslogCategory var bind value

Message Template: %1 — Template for composing syslogText value (optional)

Suppress Alarm: □

14 Apply and then Save.

Creating Extended Event Definition for syslogNotification

Create extended event definitionevents to associate certain alarms for syslogNotification categories.

1 Go to the Event Definitions port let (Settings > Alarm Definitions).

2 Locate the syslogNotification event.

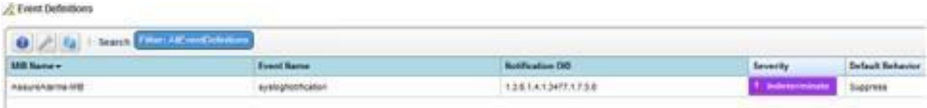

MIB Name	Event Name	Notification OID	Severity	Default Behavior
AssureAlarms-MIB	syslogNotification	1.3.6.1.4.1.3477.1.7.5.0	Indeterminate	Suppress

3 Right-click the syslogNotification event and then select Edit.

The Editing Event Definition window is displayed.

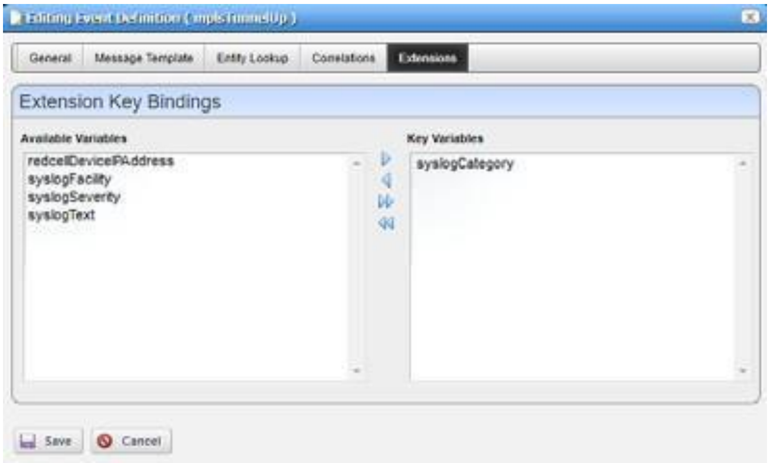

Extension Key Bindings

Available Variables

redcellDeviceIPAddress
syslogFacility
syslogSeverity
syslogText

Key Variables

syslogCategory

4 Click the Extensions tab.

5 Add syslogCategory to the Key Variables list.

6 Click Save.

7 Right-click the syslogNotification event and then select Add Extended Event Definition0.

The Adding Extended Event Definition window isdisplayed.

8 Use the category specified earlier and have `syslogCategory`create the Extended Event Definition.

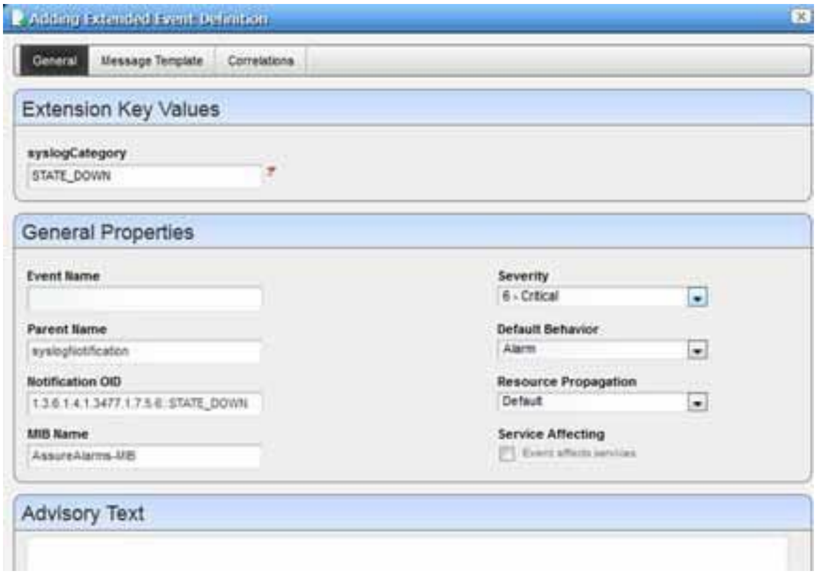

Extension Key Values

syslogCategory: STATE_DOWN

General Properties

Event Name:

Parent Name: syslogNotification

Notification OID: 1.3.6.1.4.1.3477.1.7.5.6::STATE_DOWN

MIB Name: AssureAlarms-MIB

Severity: 6 - Critical

Default Behavior: Alarm

Resource Propagation: Default

Service Affecting: □ Event affects services

Advisory Text

9 Verify that the next alarm is the alarm created during this exercise, after creating the syslog events for syslogNotification::STATE_DOWN and STATE_UP.

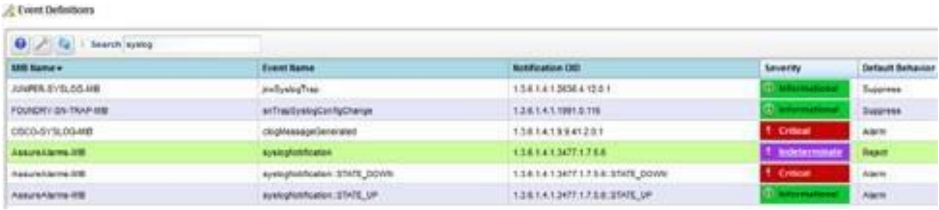

Event Definitions

MIB Name ▾	Event Name	Notification OID	Severity	Default Behavior
JUNIPER-SYSLOG-MIB	jnxSyslogTrap	1.3.6.1.4.1.2636.4.12.0.1	Informational	Suppress
FOUNDRY-SN-TRAP-MIB	snTrapSyslogConfigChange	1.3.6.1.4.1.1991.0.116	Informational	Suppress
CISCO-SYSLOG-MIB	clogMessageGenerated	1.3.6.1.4.1.9.9.41.2.0.1	Critical	Alarm
AssureAlarms-MIB	syslogNotification	1.3.6.1.4.1.3477.1.7.6.6	Indeterminate	Reject
AssureAlarms-MIB	syslogNotification::STATE_DOWN	1.3.6.1.4.1.3477.1.7.5.6::STATE_DOWN	Critical	Alarm
AssureAlarms-MIB	syslogNotification::STATE_UP	1.3.6.1.4.1.3477.1.7.5.6::STATE_UP	Informational	Alarm

10 Log into the device and shutdown an interface and then check the Alarms or Event History:

```
S5000> configure t
> interface te 0/20
> shutdown
```

Then bring up the interface:

```
> no shutdown
```

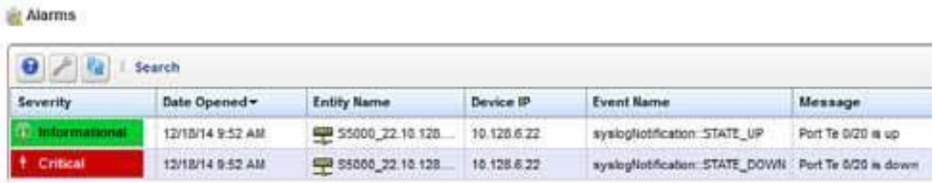

Alarms

Severity	Date Opened ▾	Entity Name	Device IP	Event Name	Message
Informational	12/18/14 9:52 AM	S5000_22.10.128...	10.128.6.22	syslogNotification::STATE_UP	Port Te 0/20 is up
Critical	12/18/14 9:52 AM	S5000_22.10.128...	10.128.6.22	syslogNotification::STATE_DOWN	Port Te 0/20 is down

Creating ACLI Action and Pass Varbind Value for Execution

For this example, create an ACLI action to bring up the interface based on the Varbind extracted from the syslogNotification event. So, if you shut down an interface, what we did previously extracted the portinformation into thesyslogTextattribute, the Adaptive CLI uses the Varbind for the ACLI script we set up here.

Create an ACLI action pass the Varbind value for execution as follows.

1 Navigate to the Actions portlet (Resources > Actions).

1 Create a new action.

a. Right-click the portlet and then select New > Adaptive CLI.

b. Specify the general settings and anyassociated actions.

c. Save the action.

2 Select the Attributes tab.

3 Create a String attribute `Message`.

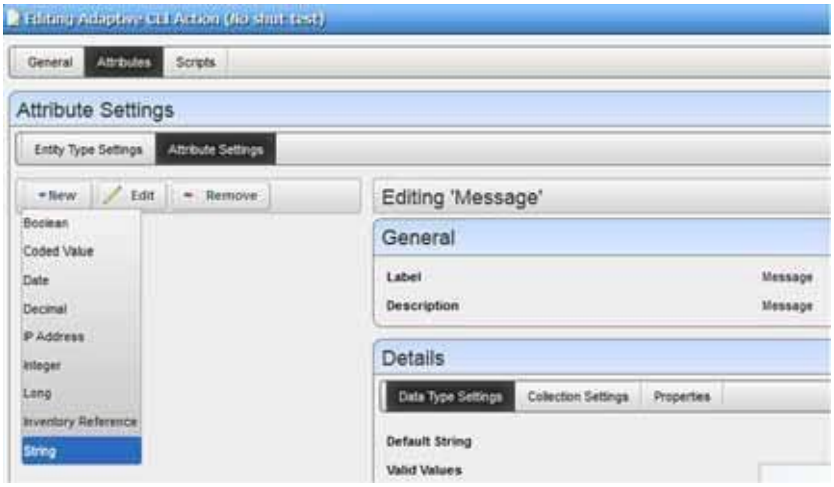

Label: Message

Description: Message

4 Select the Scripts tab.

5 Create an Embedded CLI script.

For example, enter the following commands:

```
configure t
interface [Message]
no shutdown
```

```
exi
```

We will create a post processing rule and pass a Varbind OID so it goes into the Message parameter specified here.

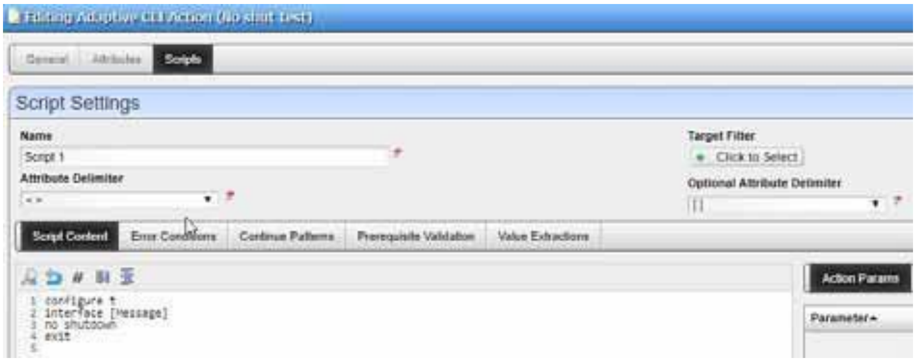

Name: Script 1

Attribute Delimiter: < >

Optional Attribute Delimiter: ||

```
configure t
interface [Message]
no shutdown
exit
```

6 Create an Automation rule.

a. Select Alarms > Automation.

The Automation and Event Processing Rules portlet is displayed.

b. Right-click and then select New > Automation.

c. Add the event with which to trigger an action.

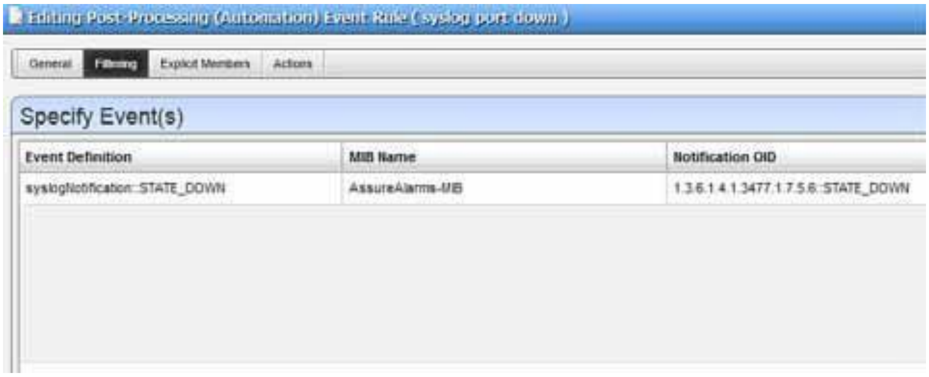

Editing Post-Processing (Automation) Event Rule (syslog port down)

Specify Event(s)

Event Definition	MIB Name	Notification OID
syslogNotification::STATE_DOWN	AssureAlarms-MIB	1.3.6.1.4.1.3477.1.7.5.6::STATE_DOWN

d. Click the Actions tab and then select Add Action > Action.

e. Select the action you made.

f. Pass the Varbind OID to the Adaptive CLI action.

For a Varbind, the syntax is `[[VARBIND: <varbind_oid>]]`.

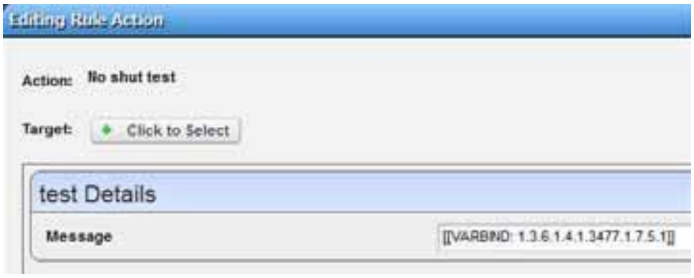

Editing Rule Action

Action: No shut test

Target:

test Details

Message: [[VARBIND: 1.3.6.1.4.1.3477.1.7.5.1]]

The Varbind of interest is the syslogText extracted port string from syslogNotification.

7 Check the Event Details from Event History portlet.

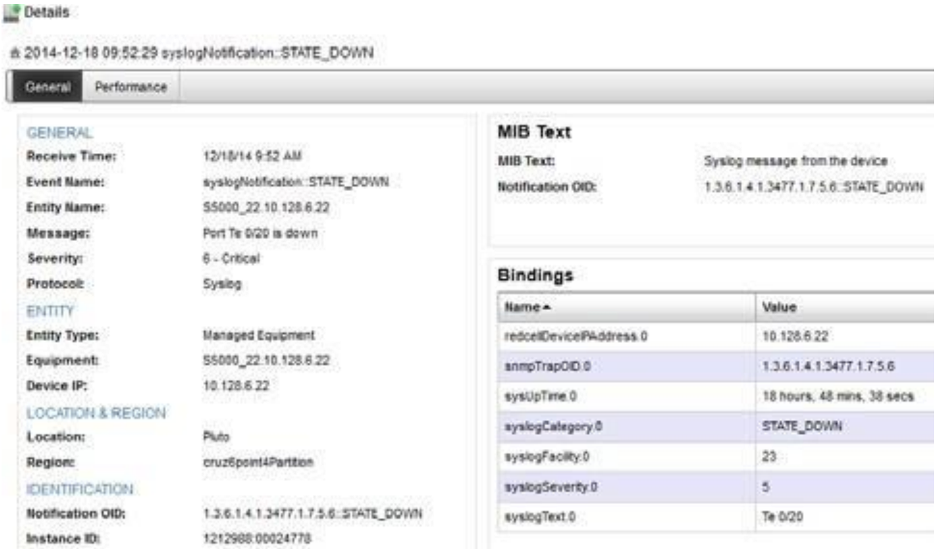

Details

2014-12-18 09:52:29 syslogNotification::STATE_DOWN

GENERAL

Receive Time: 12/18/14 9:52 AM

Event Name: syslogNotification::STATE_DOWN

Entity Name: S5000_22.10.128.6.22

Message: Port Te 0/20 is down

Severity: 6 - Critical

Protocol: Syslog

ENTITY

Entity Type: Managed Equipment

Equipment: S5000_22.10.128.6.22

Device IP: 10.128.6.22

LOCATION & REGION

Location: Pluto

Region: cruz6point4Partition

IDENTIFICATION

Notification OID: 1.3.6.1.4.1.3477.1.7.5.6::STATE_DOWN

Instance ID: 1212988:00024778

MIB Text

MIB Text: Syslog message from the device

Notification OID: 1.3.6.1.4.1.3477.1.7.5.6::STATE_DOWN

Bindings

Name ▲	Value
redcellDeviceIPAddress.0	10.128.6.22
snmpTrapOID.0	1.3.6.1.4.1.3477.1.7.5.6
sysUpTime.0	18 hours, 48 mins, 38 secs
syslogCategory.0	STATE_DOWN
syslogFacility.0	23
syslogSeverity.0	5
syslogText.0	Te 0/20

8 Get the Varbind OID.

a. Go to the Event Definition portlet (Settings > Alarm Definitions).

b. Right-click the event definition and then select Edit.

The Editing Event Definition window is displayed.

c. Click the Message Template tab.

The syslogText OID shows as `1.3.6.1.4.1.3477.1.7.5.1`

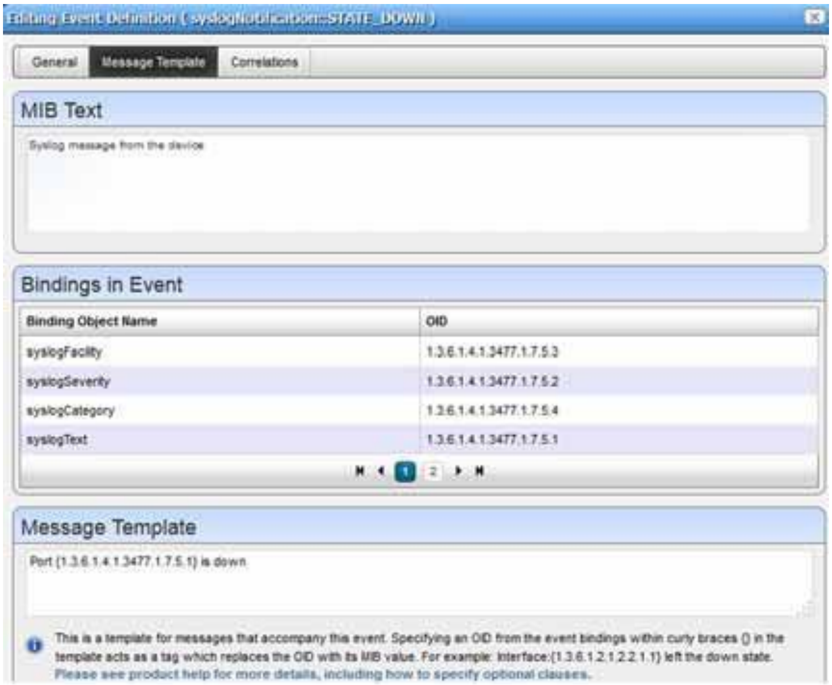

MIB Text

Syslog message from the device

Bindings in Event

Binding Object Name	OID
syslogFacility	1.3.6.1.4.1.3477.1.7.5.3
syslogSeverity	1.3.6.1.4.1.3477.1.7.5.2
syslogCategory	1.3.6.1.4.1.3477.1.7.5.4
syslogText	1.3.6.1.4.1.3477.1.7.5.1

Message Template

Port {1.3.6.1.4.1.3477.1.7.5.1} is down

This is a template for messages that accompany this event. Specifying an OID from the event bindings within curly braces {} in the template acts as a tag which replaces the OID with its MIB value. For example: Interface:{1.3.6.1.2.1.2.2.1.1} left the down state. Please see product help for more details, including how to specify optional clauses.

Verifying the Results

Verify the results as follows.

1 Trigger the *syslogNotification::STATE_DOWN* event byshutting downan interface.

If the Adaptive CLI action successfully brings the interface back up, it triggers the *syslogNotification::STATE_UP event* .

2 Check the audit trail for the Adaptive CLI action.

When successful, it indicates the Adaptive CL I successfully used the syslogTextvarbind message.

Event Definitions

You can define how the Cruz application treats notifications (events) coming into the system. Administrators define event behavior deciding whether it is suppressed, rejected, orgenerates an Alarm. Manage the definitions of events from this portlet.

Bydefault, this portlet is accessed byselecting Settings > Alarm Definitions from the navigation bar.

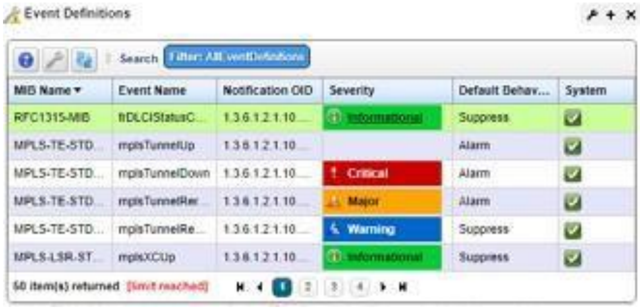

Event Definitions

MIB Name ▾	Event Name	Notification OID	Severity	Default Behav...	System
RFC1315-MIB	frDLCIStatusC...	1.3.6.1.2.1.10....	Informational	Suppress	☑
MPLS-TE-STD...	mplsTunnelUp	1.3.6.1.2.1.10....		Alarm	☑
MPLS-TE-STD...	mplsTunnelDown	1.3.6.1.2.1.10....	Critical	Alarm	☑
MPLS-TE-STD...	mplsTunnelRer...	1.3.6.1.2.1.10....	Major	Alarm	☑
MPLS-TE-STD...	mplsTunnelRe...	1.3.6.1.2.1.10....	Warning	Suppress	☑
MPLS-LSR-ST...	mplsXCUp	1.3.6.1.2.1.10....	Informational	Suppress	☑

50 item(s) returned [limit reached]

From this portlet, you can configure events that, when correlated as described in Automation and Event Processing Rules on page 288, trigger actions.

Columns include the *MIB Name, Event Name, Notification OID, Severity* for associated alarms , and *Default Behavior.* Alter these settings from the Event Definition Editor. Right-click a selected event definition for the following menu items:

Edit — Open the selected event definition in the Event Definition Editor.

Set Behavior — This lets you select from the following options*:*

- *Reject* – Every received message is rejected.
- *Suppress* – The message is tracked in Event History and then ignored.
- *Alarm* – The message is tracked in Event History and then processed through the alarm life cycle, which might create a new alarm, increment an existing alarm, or clear an existing alarm.

Set Severity — Set the alarm severity for the selected event or select the Cleared option.

MIB— This lets you upload a new MIB to your event definitions.

Import/Export—Export the selected configfile to disk or import it fromdisk. Youcanalso import/ export a selected configuration file.

Provides the following actions when available for the selected image:

- Import retrieves a file containing XML compliance descriptions. Some imports can come from a URL.
- Export Selection exports the selected description to an XML file.
- Export All exports all descriptions to an XML file.

Click Download Export File to specify where to save the file.

The Import/Export option is useful as a backup or to share descriptors with other projects.

You must import data into the correct portlet. For example, you cannot import event data into the Actions portlet.

If one type of data depends on another, you must import the other data before importing the data that depends on it.

Share with User—Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Youcan also configure an Aging Policy and View events as PDF in this menu. See Implementing DAP on page 80 for more details.

Tosee an event's propagation policy, you can view the editor panel described below. See also Understanding Alarm Propagation to Services and Customers on page340.

Unknown Traps

Cruz normalizesallincoming trapsnotelsewhere definedas `redcellUnknownTrap`(essentially "none of the above"). When a trap arrives with an OID not in the list of event definitions then the `redcellUnknownTrap` determines its behavior. Youcan configure Cruz to reject all such traps, suppress them so that they become events or allow them all to become full-fledged alarms.

Youcan alsocreate event processingrules to handle (suppress, alarm, send e-mail) anysuch events. See Creating Event Processing Rules on page 292. The `redcellUnknownTrap`'s default behavior is suppress, its default severity is indeterminate. Events and alarms that are unknown still contain the notification OID from the trap. As a formality, `redcellUnknownTrap` has its own notification OID.

Different Forms of Event and Alarm Correlation

Cruz supports different forms of event and alarm correlation:

- **Stream Based Correlation:** When multiple eventsoccur within the mediation server and there are rules that correlate them with each other so that when a certain kind of pattern is detected (frequency threshold, state flutter), this affects which events the mediation server submits to the application server. You can con figure this form of correlation by creating certain types of event processing rules. See Stream Based Correlation on page 295
- **Event to Alarmcorrelation:** Open alarms are uniquely identified by the source entity and the key bindings. This means that if a new alarm comes in that is essentially identical to an existing open alarm, this will increment the count of the open alarm and also in some cases escalate it or de-escalate it (if the incoming alarm has a different severity level or a different message). This is facilitated by the correlation hash field on each alarm record, which encodes all of the fields relevant for correlation. Note that for key bindings, the data within the binding that is used for correlation depends on how the variable binding definition is configured. If it is configured to correlate By Value then the value within the binding is used and conversely if it is configured to correlate By Index then the index of the OID is used. The index is the last segment in the OID string after the defined OID. Forexample, if a variable binding definition has an OID of 1.2.3 and the payload of a trap has a variable binding based on this definition with an OID of 1.2.3.55 then 55 is the index of this binding. There will also be a value associated with this binding in the payload of the trap, which is different from the index. See Key Bindings within Correlations on page 330 for more information.
- **Raising/Clearing Alarm Correlation**: When an open alarmwith a raising severity (Critical, Major, Minor, or Warning) is correlated with another alarm that was recently created whose severity is cleared, this will then clear the raising alarm (it will no longer be opened). See Event Definitions Correlated Events for Raising/Clearing within Correlations on page 330 for more information.
- **Parent/Child Alarm Correlation:** Insome cases, analarmcancauseother alarmsto be created and in other cases an alarm can block other alarms from being addressed. In these cases, the causing or blocking alarm becomes the parent and the alarms that were caused or blocked are the children. This form of correlation can be performed manually or automatically. For information on manual parent/child alarm correlation, see Parent/Child Alarm Correlation: Alarm Details Panel on page 173. For information on automatic parent child alarm correlation, see Event Definitions Correlated for Parent/Child within Correlations on page 330.

Event Definition Editor

This editor lets you modify event definitions from the following panels:

- General
- Message Template
- Entity Lookup
- Correlations
- Extensions

And then you can start adding extended event definitions .

Click *Save* to preserve any modifications you have made or *Cancel* to abandon them.

General

This tab manages basics for Event Definitions.

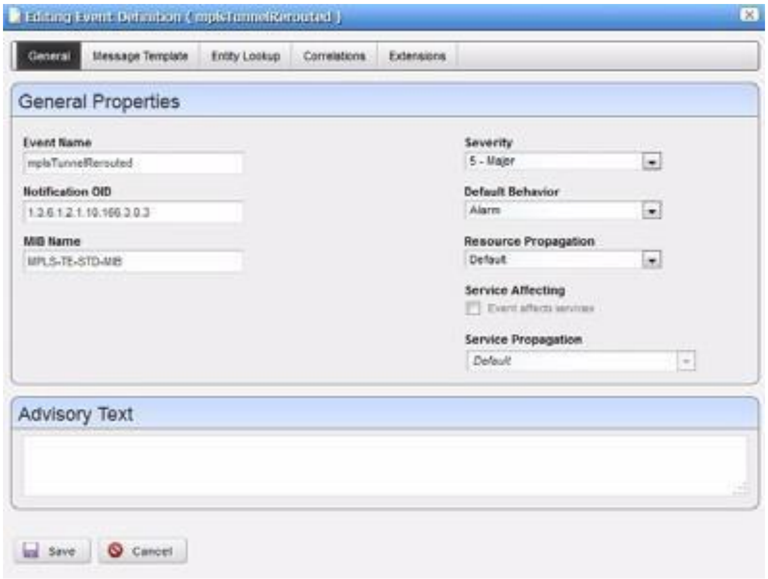

Event Name: mplsTunnelRerouted

Notification OID: 1.3.6.1.2.1.10.166.3.0.3

MIB Name: MPLS-TE-STD-MIB

Severity: 5 - Major

Default Behavior: Alarm

Resource Propagation: Default

Service Affecting: □ Event affects services

Service Propagation: Default

It has the following fields:

Event Name — A text identifier for the event.

Notification OID — The object ID.

MIB Name — The MIB with which this event is associated.

Severity—The severity of any associated alarm. If a new alarm is a clearing severity, then it closes any existing alarm to which it correlates. Otherwise, if a new alarm severity does not match the existing severity then the existing alarm is closed, and a new alarm opened for the new severity.

Default Behavior—The options for behavior (*Undefined, Alarm, Suppress, Reject*). *Alarm* means: Process at the mediation server, generate event history and an alarm. *Suppress* means: Process at the mediation server and generate an event (*not* an alarm). *Reject* means: Reject at the mediation server (do not process)

Resource Propagation — The hierarchical resource propagation behavior for any alarm based on this event definition (either *Default, Impacts subcomponents,* or *Impacts top level*, or *Impacts Top Level and Subcomponents*). (See also Understanding Alarm Propagation to Services and Customers on page 340 for more about how this impacts services and customers.)

An event definition configures "Resource Propagation" (distinct from "Alarm propagation") based on the event type. Do alarms based on this event definition impact the overall device (*Impacts top level*), subcomponents (*Impacts subcomponents),* both (Impacts Top Level and Subcomponents) or just the correlated inventory entity (*Default*)?

Alarm behavior differs for monitorAttributeTrend alarms when an SNMP Interface monitor targets a Port/Interface rather than targeting a device. If the alarm comes from the former, and its correlation state is not *Top Level Alarm*, only the port appears alarmed, not the device. If the monitor target is the device, however, the device appears as alarmed. If the monitor has Port targets, then you must configure propagation to *Top Level Alarm* to see the device alarmed in, for example, the System Topology view.

Service Affecting—Check this if the event has an impact on services. Indicates whether the alarm has an impact on services. If this is checked then alarms based on this event definition propagate calculated alarm states across services and customers that depend on the (directly) alarmed resource.

For example: If a resource has a service affecting alarm, then Cruz propagates the severity of this alarm across all associated services and customers. If the resource alarm is "clear" then all services depending on this resource are "clear" too. If the resource alarm is "critical," then all services depending on that resource are "critical" too.

NOTE:

Alarms imported from previous versions appear as not service affecting, regardless of severity.

For more about propagation, see Understanding Alarm Propagation to Services and Customers on page 340.

Service Propagation— Thisonlyaffectsalarmpropagation if Service Affecting is true for theevent definition. This is like Resource Propagation, but controls only whether alarms affect services associated with entities hierarchically related to the alarmed entity. For example, if a port alarm is created and its event definition specifies that the Service Propagation is Impacts Subcomponents then the alarm propagates only to the services associated with this port's interfaces. This does not affect the services associated with the top-level device.

Advisory Text — The *Advisory Text* appears with the event. Configure it in the text box here.

Message Template

The Message Template panel lets you viewor alter MIB Text, Bindings and the Message Template for the selected event.

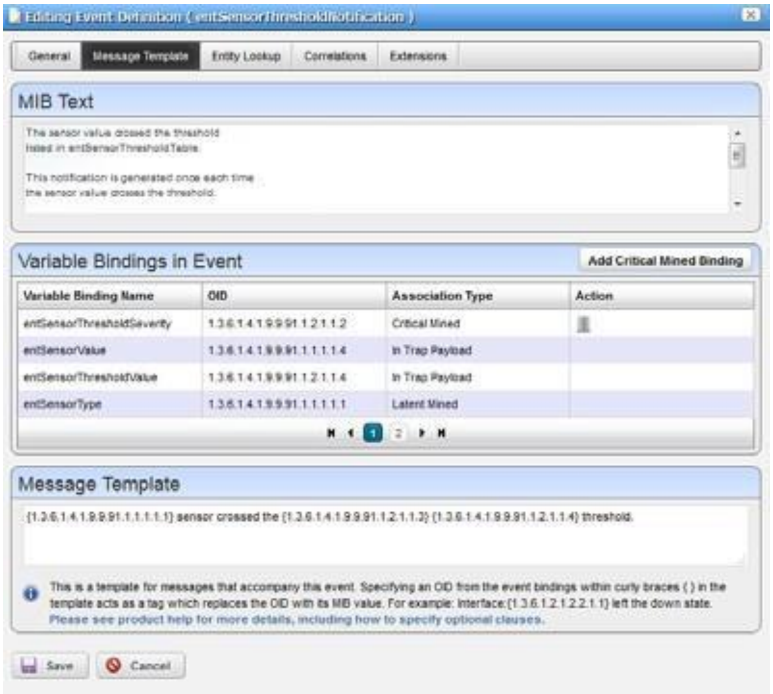
Editing Event Definition (entSensorThresholdNotification)
General
Message Template
Entity Lookup
Correlations
Extensions
Add Critical Mined Binding
Save
Cancel

MIB Text

The sensor value crossed the threshold
listed in entSensorThresholdTable.

This notification is generated once each time
the sensor value crosses the threshold.

Variable Bindings in Event

Variable Binding Name	OID	Association Type	Action
entSensorThresholdSeverity	1.3.6.1.4.1.9.9.91.1.2.1.1.2	Critical Mined	🗑
entSensorValue	1.3.6.1.4.1.9.9.91.1.1.1.1.4	In Trap Payload	
entSensorThresholdValue	1.3.6.1.4.1.9.9.91.1.2.1.1.4	In Trap Payload	
entSensorType	1.3.6.1.4.1.9.9.91.1.1.1.1.1	Latent Mined	

Message Template

{1.3.6.1.4.1.9.9.91.1.1.1.1.1} sensor crossed the {1.3.6.1.4.1.9.9.91.1.2.1.1.3} {1.3.6.1.4.1.9.9.91.1.2.1.1.4} threshold.

This is a template for messages that accompany this event. Specifying an OID from the event bindings within curly braces { } in the template acts as a tag which replaces the OID with its MIB value. For example: Interface:{1.3.6.1.2.1.2.2.1.1} left the down state.
Please see product help for more details, including how to specify optional clauses.

This contains three sections:

MIB Text — A read-only reminder of the MIB contents for this OID.

Variable Bindings in Event — Lists the variable bindings that are expected to be in events based on this definition. This displays the varbind contents of the event, matching the *Variable Binding Name*, *OID* (object identifier), and the association type. There are three ways that bindings are associated to event definitions:

- **In Trap Payload** means that the association is determined by the MIB and thus it being mentioned here is read-only.
- **Critical Mined** means that the binding needs to be mined from the device that emitted the trap so that the event is augmented to have this additional information available for the message template, correlation, entity lookup, extensions, or filtering of event processing rules.
- **Latent Mined** means that the bindings are mined only during alarm correlation. This h is too late for such bindings to be available for certain features, such as entity lookup and extensions.

Click the *Add Critical Mined Binding* button to display a list of variable binding definitions from which to select. Critical mined bindings are configured using this option. Latent mined bindings are configured from the variable binding definition level. See **Variable Binding Definitions** on page 338**.**

Message Template —A template for messages that accompany this event. Specifying an OID within the curly braces {} in the template acts as a tag which replaces the OID with its MIB value. For example: Interface: {1.3.6.1.2.1.2.2.1.1} left the down state.

You can also add optional messages surrounded by double brackets [[ ]]. if the event definition has the message "aindex: {1.2.3}[[, bindex: {1.2.4}]]" and {1.2.3} is defined as say "1" but {1.2.4} is not defined then this resolves to "aindex: 1". If they are both defined (say {1.2.4} is "2") then this resolves to "aindex: 1, bindex: 2"

If a message template exists for an existing, correlated alarm and the generated text does not match the original alarm, then Cruz closes the existing alarm, and generates a new one. Leaving this blank transmits the original message.

NOTE:

Putting an OID in curly brackets amounts to a tag replaced by the MIB text for that OID. Look for OIDs and messages in the MIB browser (as described in MIB Browser Tool on page 53).

Entity Lookup

This screen lets you configure how events based on this definition will be associated to entities. By rule, the Equipment attribute of any event will be determined by the source IP address of the original message (trap, syslog, etc.) By default, the Entity attribute will be the same as the Equipment, but there might be information in the variable bindings of the trap through with the event can be resolved to a subcomponent (Port, Interface, etc.)

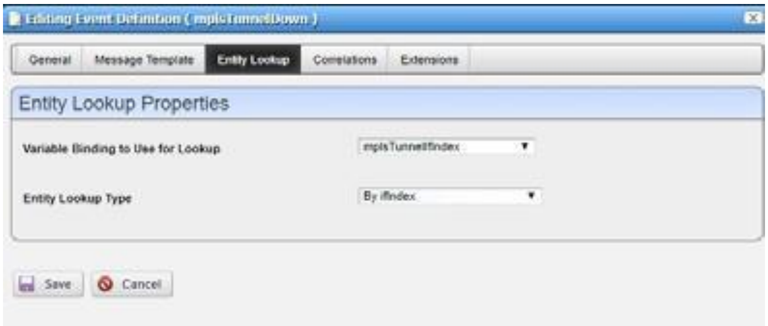

Editing Event Definition (mplsTunnelDown)

Entity Lookup Properties

Variable Binding to Use for Lookup: mplsTunnelIfIndex

Entity Lookup Type: By ifIndex

Variable Binding to Use for Lookup — Indicates the variable binding that should be used for the entity lookup. This drop-down list includes all bindings that are expected to be in the trap payload and also those that are critical mined. Latent mined bindings cannot be used for this purpose.

Entity Lookup Type — IndIndicates the algorithm that should be used to lookup the entity. For example, if the variable binding to use for lookup contains an ifIndex then you should select By ifIndex. It might also be that it instead contains the physical index, in which case there are two options, either By Physical Index (simple) to simply lookup the entity by this given physical index and associate to that entity directly, or By Physical Index (find parent) to instead use the physical index to find the hierarchical parent. This is useful when a vital component of a device, like a power supply or a fan, has sensors attached. Traps coming from this device contain the physical index of the sensor rather than the vital component being monitored. This lookup type lets you associate alarms with the vital components so you can ensure these components are functioning optimally.

Correlations

This screen lets you configure different forms of correlation for the event definition including Raising/Clearing Correlated Events, Parent/Child Correlated Events, and Correlation Key Bindings, the latter of which is essential for event to alarm correlation.

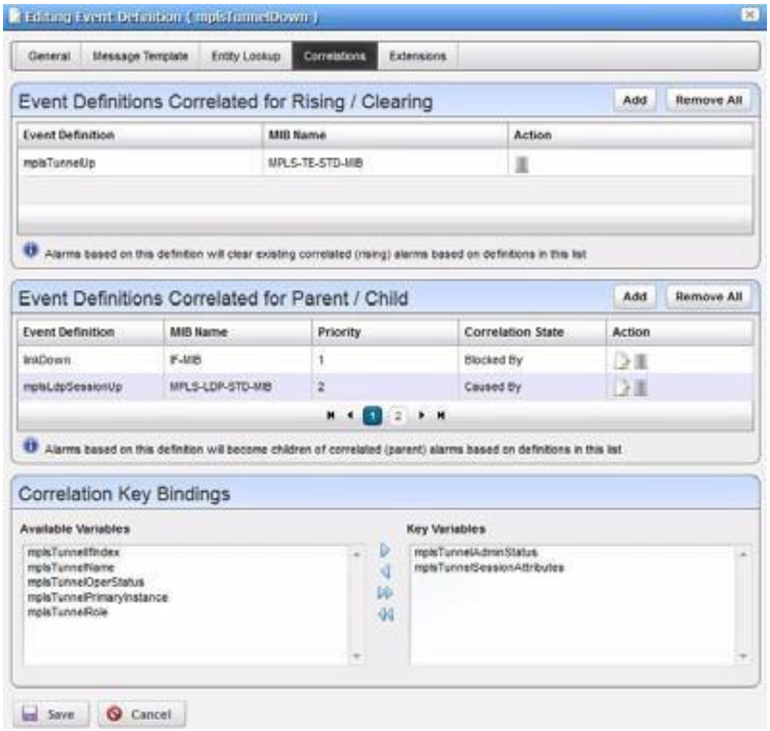

Event Definitions Correlated for Rising / Clearing

Event Definition	MIB Name	Action
mplsTunnelUp	MPLS-TE-STD-MIB	

Alarms based on this definition will clear existing correlated (rising) alarms based on definitions in this list

Event Definitions Correlated for Parent / Child

Event Definition	MIB Name	Priority	Correlation State	Action
linkDown	IF-MIB	1	Blocked By	
mplsLdpSessionUp	MPLS-LDP-STD-MIB	2	Caused By	

Alarms based on this definition will become children of correlated (parent) alarms based on definitions in this list

Correlation Key Bindings

Available Variables

mplsTunnelIfIndex
mplsTunnelName
mplsTunnelOperStatus
mplsTunnelPrimaryInstance
mplsTunnelRole

Key Variables

mplsTunnelAdminStatus
mplsTunnelSessionAttributes

The *Event Definitions Correlated Events for Raising/Clearing* panel lets you control what type of events will automatically clear the correlated alarms that are based on this definition. Within this panel, click *Add* to displaya selector (with filter) to find events definitions to correlate with the one you are editing. What this means is that if there is an existing open alarm based on this definition (we can call it a raised alarm) and then subsequently another alarm is created that is correlated to the raised one, it will cause the raised alarm to be cleared (it will no longer be open). Note that this will only work if the definition you are editing has the severity of cleared and the definition in the list has a severity other than cleared. For example, link down and link up event definitions are correlated as raising/clearing, and thus all events and alarms based on these definitions will inherit these correlations. So, consider that a link down event is received that creates a raised alarm and then subsequently a link up event is created. The link up would then correlate to and clear the link down alarm.

The *Event Definitions Correlated for Parent/Child* panel lets you control what type of alarms will automatically correlate as parent/child relationships that indicate that one alarm caused the other or that one alarm blocks progress on the other. What this means is that an open alarm based on the definition you are editing will form a parent/child correlation to another open alarm based on a definition in this list. Note that alarms base d on the definition you are editing will become correlated children of alarms based on definitions in the list (the later will become correlated parents). For example, when new alarms are created that are based on this definition, Cruz will try to find a correlated match between each new alarm and other existing open alarms, and it does this based on the entries in this list. If it finds a match, it will correlate the parent alarm to the child alarm.

Within this panel, click *Add* to display a popup panel through which you can select an event definition to correlate with the one you are editing.

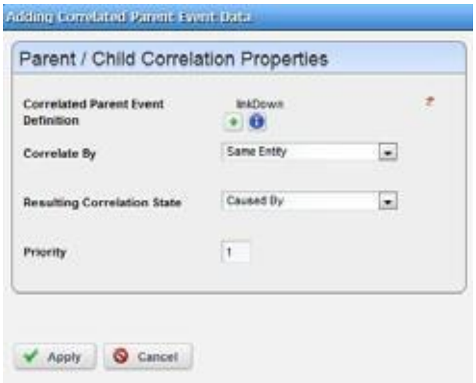

Adding Correlated Parent Event Data

Parent / Child Correlation Properties

Correlated Parent Event Definition: linkDown

Correlate By: Same Entity

Resulting Correlation State: Caused By

Priority: 1

The following fields show on this screen:

Correlated Parent Event Definition — Displaysa selector (with filter) to find the event definition to correlate to the one that you are editing. Alarms based on the definition you select will be correlated parents of alarms based on the definition that you are editing.

Correlated By — This is the algorithm that is used to determine if there is a correlated match between any potential parent and child alarms. Currently the only option is Same Entity, meaning that the two alarms must be associated with the same entity in order to be correlated in this way.

Resulting Correlation State — This indicates the type of parent /child correlation that should be created. Caused By indicates that the parent alarm caused the child alarm to be created and Blocked By indicates that the parent alarm is blocking progress on the child alarm. The most significant difference in terms of howsuch correlations are processed occurs when the parent alarm is cleared. In such scenarios, caused by children will be cleared automatically when the parent is cleared, whereas blocked by children will simply be unblocked when the parent is cleared.

Priority — Indicates the priority that this entry should have in relation to the other entries in the list. If there are multiple entries in this list, then Cruz will first try to find correlated matches based on the entry with priority 1 and then if there are no matches it will try to do so based the entry that is priority 2, and so on.

In Correlation Key Bindings, use the right/left arrows to select *Key Variables* from *Available Variables*. The variables considered keys for event to alarm correlation are the key bindings for the target alarm in the correlation process. This means that if event A is defined to include event B as a correlated event, comparison of the key bindings defined for event B is also considered when comparing a new alarm for event A to an existing alarm for event B.

Extensions

This START panel lets you configure Extension Key Bindings. Extensions allow you to adding extended event definitions.

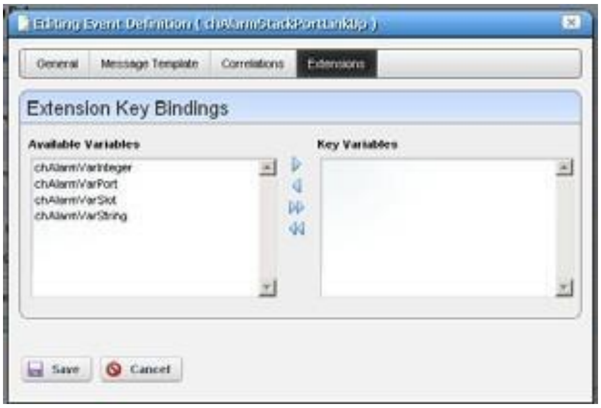

Available Variables
chAlarmVarInteger
chAlarmVarPort
chAlarmVarSlot
chAlarmVarString

Key Variables

In Extension Key Bindings, use the right/left arrows to select *Key Variables* from *Available Variables*. For example, if a device generates one type of trap, but has three key variables (alarmID, state, and severity), you may extend the event definition based on state. One derived definition for the "alarm" state and another for the "clear" state. You probably would extend based on state and severity to define a clearing definition for all events where state is "clear" and then derive a specific "alarm" event for each severity as defined by the vendor.

Adding Extended Event Definitions

The menu item for this screen can appear when yo u right-click an extendable event (one that has one or more Extensions). For an implemented example, see Extending Event Definitions on page 333.

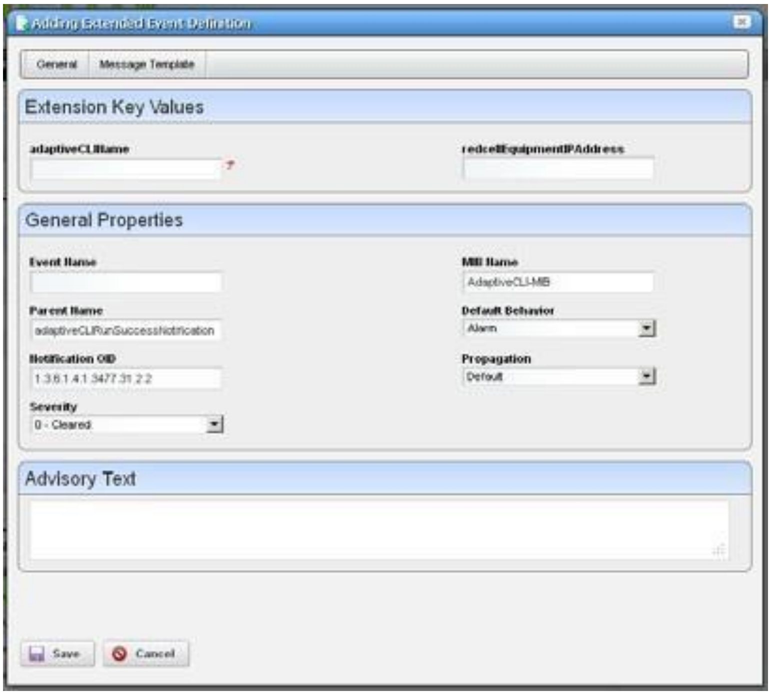

adaptiveCLIName

redcellEquipmentIPAddress

Event Name

MIB Name: AdaptiveCLI-MIB

Parent Name: adaptiveCLIRunSuccessNotification

Default Behavior: Alarm

Notification OID: 1.3.6.1.4.1.3477.31.2.2

Propagation: Default

Severity: 0 - Cleared

The Extended Event Definitions editor has the Extension Key Values fields that are the same as those configured in Extensions, and the following general properties:

Event Name —An identifier for the Extended Event. This field is editable, but if left blank the system creates a name based on the name of the parent event definition and the key bindings entered here.

MIB Name — A MIB identifier for the Extended Event

Parent Name — An identifier for the parent of the Extended Event

Notification OID — As with parent event definitions, this is the object ID. Cruz automatically generates this based on the Notification OID of the parent and key binding values entered. For example, if the parent event definition has a Notification OID of “1.2.3.4” and the key binding values of the extended definition are 5 and 6 (the parent must have previously been configured to have two extension bindings available) then the resulting Notification OID for this new extended event definition will be “1.2.3.4::5:6”.

The remaining fields are as described in the Event Definition Editor General screen. The Message Template tab is as described in Message Template on page 328, only this message template is for the configured extended event definition.

Click *Save* to create or re-configure an extended event or *Cancel* to abandon your edits. See Extending Event Definitions on page 333 for detailed steps.

Using Extended Event Definitions

To handle complex networks, you can configure event responses with a “Base Event Definition” and several“Extended Event Definitions” (EEDs). Youcancustomize events, the reaction to them, and even extend events with EEDs. Modified or extended event definitions override the base set of definitions that comes from the traps and notifications found in loaded MIB files, and their OIDs appear in the Event Definition Editor on page 326.

You can modify Extended Event Definitions two ways:

- Extending Event Definitions
- Extending Event Definition XML

Youcan configure extended event definitions to perform partial matching on variable bindings by including **wildcard characters** in the extension key. The asterisk * matches any number of characters and the question mark?matches a single character. Forexample, an extension key of 1* matches variable bindings that start with a 1 including 11, 12, 134, and so on. An extension keyof 1? matches 11, 12, 13 but does not match 134. An extension key of *3* matches any variable binding thatincludes the number 3 somewhere, including 34, 13, 3, 438.

Extending Event Definitions

Toextend event definitions, you create “child” events that extend a “parent” event. Extend an event definition as follows.

1 Save the parent event for which you configured Extensions.

2 Select it from the Event Definitions manager.

3 Right-click your selection and then select Add Extended Event Definition.

The Adding Extended Event Definitions on page 332 window the Extension Key Values fields configured from the Event Definitions, Extensions panel. The variables also show in the order you selected from the Extensions panel. A suggested name for this event is the parent event name, followed by a colon-separated list of the field values entered on this screen.

Once configured, this extended event overrides the parent.

Extending Event Definition XML

After loading its base event definitions, the application loadsall event definitions in the XMLfiles in the server\conf directory under each module (for example \owareapps\redcell\server\conf \eventdefs.xml) and applies them as overrides. The Cruz application discards any XML event definitions if the notification type is unknown—that is, not in a loaded MIB. If you alter this XML file, you need to import it in Event Definitions Manager before the alterations take effect.

Extended event definitions (EEDs) support system extensions along with system overrides. An extension creates a unique definition for a MIB-based event that only applies to certain instances of the event. The extension appears as a separate event in the event definition manager and you can modify it as you can any other event.

Extension exist primarily to create specialized event definitions for recognizable variations from the basic event.

For details, see the following:

- XML Event Definition
- Extended Event Definitions
- Extended Event Definition Example

XML Event Definition

The format format is of an XML event definition:

```
<bean xsi:type="tns:NotificationDefinitionNP">
    <NotificationOID/>   <!-- required -->
    <ExtensionKey/>     <!-- extensions only - syntax: key1Value[:key2Value]-
      ->
    <Description/>       <!-- message text - supports object OID tag
      replacement: {objectOID} -->
    <Behavior/           <!-- 1=alarm, 2=suppress, 3=reject (default is
      alarm)-->
    <Severity/>          <!-- 0=clearing, 1=info, 2=unkown, 3=warning,
      4=minor, 5=major, 6=critical -->
    <ImpactPropagation/> <!-- 1=up, 2=down, 3=both (default is none) -->
    <KeyBindings>        <!-- OID(s) for bind object value(s) to use as
      correlation key -->
    <item/>
    </KeyBindings>
    <CorrelatedEvents>   <!-- OID(s) for alarms to clear when this event
      occurs. These OIDs can include an extension key as OID::extKey -->
       <item>oid[::keyValue[:keyvalue]]</item>
    </CorrelatedEvents>
</bean>
```

Extended Event Definitions

You can extend any event definition that defines one or more key bindings. The extension key consists of a value for each key binding in the correct order. You must separate each value with a colon (:).

Only extensions should set an extension key. The key format is dictated by the parent (extended) definition’s keybindings. Set keybindings for the extension as needed for proper alarmcorrelation. The following Extended Event Definition Example provide some clarification.

NOTE:

Definition extensions do not inherit any settings from the extended definition.

Extended Event Definition Example

Events extend a basic definition:

```
bigBand 1.3.6.1.4.1.6387
```

Here are the extensions:

```
bigBandAdmin 1.3.6.1.4.1.6387::400
bigbandConfig 1.3.6.1.4.1.6387::5* (with partial matching extension key)
bigbandTraps 1.3.6.1.4.1.6387::400:50
bigbandAlarmTrapPrefix 1.3.6.1.4.1.6387::400:50:0
bigbandSessionAlarm 1.3.6.1.4.1.6387.400.50.0.0.2 (v2 OID)
bigbandEscalation 1.3.6.1.4.1.6387::400:50:0:0:3* (with partial matching
  extension key)
```

Another base event:

```
bigBandCommon 1.3.6.1.4.1.6387.100
```

Other extensions:

```
session 1.3.6.1.4.1.6387.100::50
sessionAlarms 1.3.6.1.4.1.6387.100::50:100
sessionAlarmTable 1.3.6.1.4.1.6387.100::50:100:40
sessionAlarmEntry 1.3.6.1.4.1.6387.100::50:100:40:1
sessionAlarmProgramNumber 1.3.6.1.4.1.6387.100::50:100:40:1:2
sessionAlarmEscalation 1.3.6.1.4.1.6387.100::50:100:40:1*:2? (with
  partial matching extension key)
```

The following example variables are for such extended events:

```
bigbandSessionAlarm TRAP-TYPE
    ENTERPRISE bigbandAlarmTrapPrefix
    VARIABLES {
        sessionAlarmAssertedTime,
        bigbandAlarmOnOrOff,
        sessionAlarmOutputChannelIndex,
        sessionAlarmProgramNumber,
        sessionAlarmPid,
        sessionAlarmAssertedType,
        sessionAlarmAssertedInputChannelIndexOrZero,
        bigbandSessionAlarmSequenceNo,
        sessionAlarmSessionId,
       sessionAlarmAuxiliary1,
       sessionAlarmAuxiliary2
    }
    DESCRIPTION
        "a session alarm is generated for every asserted/removed
         alarm. It contains all parameters as in sessionAlarm table,
         as well as a sequence no. that is incremented by 1 for
         every session-trap generation."
    ::= 2
```

The sample XML for the definitions and extensions:

```
<!-- bigbandSessionAlarm -->
  <bean xsi:type="tns:NotificationDefinitionNP">
    <NotificationOID>1.3.6.1.4.1.6387.400.50.0.0.2</NotificationOID>
    <Description>Program #{1.3.6.1.4.1.6387.100.50.100.40.1.2}</
 Description>
    <Behavior>1</Behavior>
    <Severity>2</Severity>
    <KeyBindings>
      <!-- bigbandAlarmOnOrOff -->
      <item>1.3.6.1.4.1.6387.400.50.1</item>
      <!-- sessionAlarmProgramNumber -->
      <item>1.3.6.1.4.1.6387.100.50.100.40.1.2</item>
    </KeyBindings>
  </bean>

  <!-- bigbandSessionAlarm ext:alarmOn program 12 -->
  <bean xsi:type="tns:NotificationDefinitionNP">
    <NotificationOID>1.3.6.1.4.1.6387.400.50.0.0.2</NotificationOID>
    <ExtensionKey>1:12</ExtensionKey>
    <Description>program #12 alarm ON</Description>
    <Behavior>1</Behavior>
    <Severity>4</Severity>
    <CorrelatedEvents>1.3.6.1.4.1.6387.400.50.0.0.2::0:12</
 CorrelatedEvents>
  </bean>

  <!-- bigbandSessionAlarm ext:alarmOff program 12-->
  <bean xsi:type="tns:NotificationDefinitionNP">
    <NotificationOID>1.3.6.1.4.1.6387.400.50.0.0.2</NotificationOID>
    <ExtensionKey>0:12</ExtensionKey>
    <Description>Program #12 alarm OFF</Description>
    <Behavior>1</Behavior>
    <Severity>0</Severity>
    <CorrelatedEvents>1.3.6.1.4.1.6387.400.50.0.0.2::1:12</
  CorrelatedEvents>
   </bean>
```

Tobuild on the previous BigBand example, the followingexample shows support for an EED with a partial key. It also shows the base event as an indeterminate severity alarm, then extends it for *alarm on* and *alarm off* default behaviors using partial keys. Finally, it is extended twice more, showing example of program-specific event settings.

```
<!-- bigbandSessionAlarm -->
  <bean xsi:type="tns:NotificationDefinitionNP">
    <NotificationOID>1.3.6.1.4.1.6387.400.50.0.0.2</NotificationOID>
    <Description>Program #{1.3.6.1.4.1.6387.100.50.100.40.1.2}</
  Description>
    <Behavior>1</Behavior>
    <Severity>2</Severity>
    <KeyBindings>
      <!-- bigbandAlarmOnOrOff -->
     <item>1.3.6.1.4.1.6387.400.50.1</item>
      <!-- sessionAlarmProgramNumber -->
      <item>1.3.6.1.4.1.6387.100.50.100.40.1.2</item>
    </KeyBindings>
  </bean>

  <!-- bigbandSessionAlarm ext: alarmOff all programs -->
  <bean xsi:type="tns:NotificationDefinitionNP">
    <NotificationOID>1.3.6.1.4.1.6387.400.50.0.0.2</NotificationOID>
    <ExtensionKey>0</ExtensionKey>
    <Description>Program #{1.3.6.1.4.1.6387.100.50.100.40.1.2}</
  Description>
    <Behavior>1</Behavior>
    <Severity>0</Severity>
  </bean>

  <!-- bigbandSessionAlarm ext: alarmOn default-->
  <bean xsi:type="tns:NotificationDefinitionNP">
    <NotificationOID>1.3.6.1.4.1.6387.400.50.0.0.2</NotificationOID>
    <ExtensionKey>1</ExtensionKey>
    <Description>Program #{1.3.6.1.4.1.6387.100.50.100.40.1.2}</
  Description>
    <Behavior>1</Behavior>
    <Severity>6</Severity>
  </bean>

  <!-- bigbandSessionAlarm ext: alarmOn program 12 -->
  <bean xsi:type="tns:NotificationDefinitionNP">
    <NotificationOID>1.3.6.1.4.1.6387.400.50.0.0.2</NotificationOID>
    <ExtensionKey>1:12</ExtensionKey>
    <Description>program #12 alarm ON</Description>
    <Behavior>1</Behavior>
    <Severity>4</Severity>
  </bean>

  <!-- bigbandSessionAlarm ext: alarmOn program 40 -->
  <bean xsi:type="tns:NotificationDefinitionNP">
    <NotificationOID>1.3.6.1.4.1.6387.400.50.0.0.2</NotificationOID>
    <ExtensionKey>1:40</ExtensionKey>
    <Description>#40 special message text</Description>
    <Behavior>1</Behavior>
    <Severity>5</Severity>
  </bean>
```

Variable Binding Definitions

Events contain data elements called Variable Bindings and known as varbinds for short. Just as Events are based on Event Definitions, Variable Bindings are based on Variable Binding Definitions.

Variable Binding Definitions Portlet

Youcan define how the Cruz application treats variable bindings that are found within events. Administrators can fine-tune certain aspects of how variable bindings are processed using this portlet.

By default, this portlet is available from the Settings page byselecting the Alarm Definitions menu option.

You might need to edit a variable binding definition if any of these scenarios occurs:

- Alarm to event correlation is not working properly for events/alarms based on a certain event definition - In this case, you might need to fin d which variable bindings are within the event and change the Correlation Type for one or more of the variable binding definitions. This might then correct the correlation issues for new events and alarms based on the same definitions.
- The message within an alarm is not descriptive enough - In this case, you might need to find which variable bindings are within the event definition upon which the alarm is based and then add latent mined variable bindings so that new alarms based on these definitions will include more information.

- The message within an alarm is unreadable (i.e. it shows up as hexidecimal code) - In this case, you might need to find the which variable bindings are within the event definition upon which the alarm is based and then set it to perform ForcedConversion.

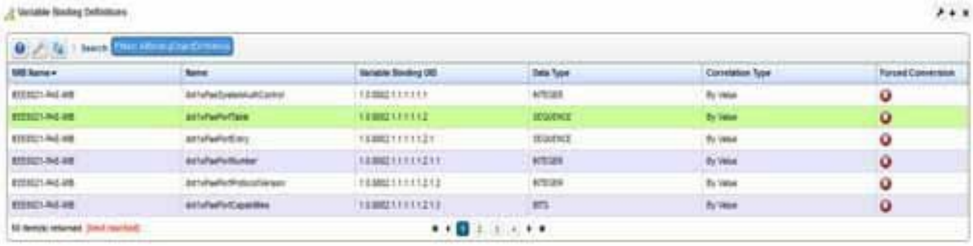

MIB Name ▾	Name	Variable Binding OID	Data Type	Correlation Type	Forced Conversion
IEEE8021-PAE-MIB	dot1xPaeSystemAuthControl	1.0.8802.1.1.1.1.1.1	INTEGER	By Value	●
IEEE8021-PAE-MIB	dot1xPaePortTable	1.0.8802.1.1.1.1.1.2	SEQUENCE	By Value	●
IEEE8021-PAE-MIB	dot1xPaePortEntry	1.0.8802.1.1.1.1.1.2.1	SEQUENCE	By Value	●
IEEE8021-PAE-MIB	dot1xPaePortNumber	1.0.8802.1.1.1.1.1.2.1.1	INTEGER	By Value	●
IEEE8021-PAE-MIB	dot1xPaePortProtocolVersion	1.0.8802.1.1.1.1.1.2.1.2	INTEGER	By Value	●
IEEE8021-PAE-MIB	dot1xPaePortCapabilities	1.0.8802.1.1.1.1.1.2.1.3	BITS	By Value	●

[illegible] item(s) returned (end reached)

Columns include the *MIB Name, Name, Variable Binding OID, Data Type, Correlation Type, and Forced Conversion*. See Variable Binding Definition Editor for how to alter these. Right-click a selected variable binding definition for the following menu items:

Edit — Open the selected variable binding definition in the Variable Binding Definition Editor.

Variable Binding Definition Editor

This editor lets you modify variable binding definitions

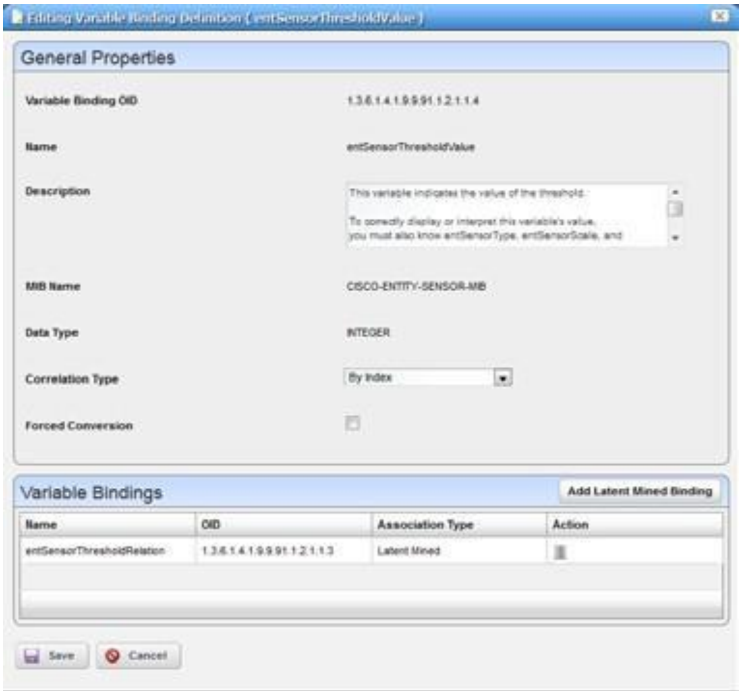

Variable Binding OID	1.3.6.1.4.1.9.9.91.1.2.1.1.4
Name	entSensorThresholdValue
Description	This variable indicates the value of the threshold. To correctly display or interpret this variable's value, you must also know entSensorType, entSensorScale, and
MIB Name	CISCO-ENTITY-SENSOR-MIB
Data Type	INTEGER
Correlation Type	By Index
Forced Conversion	□

Name	OID	Association Type	Action
entSensorThresholdRelation	1.3.6.1.4.1.9.9.91.1.2.1.1.3	Latent Mined	

It has the following fields:

Name — A text identifier for the variable binding.

Variable Binding OID — Unique identifier for the variable binding.

MIB Name — The MIB with which this variable binding is associated.

Data Type — The data type, as defined in the MIB.

Correlation Type— The way that variable bindings based on this definition should be correlated. You might need to change this value if certain types of events and/or alarms do not correlate correctlyfor your system-for example the clearing alarm does not clear the initial raising alarm-then you can edit one or more of the variable binding on the Message Template tab of the event definition. The possible values for this attribute are: *By Value*, where all bindings have both an OID and a value, the value is used to correlate (this is default), and *By Index*, e.g. OID of 1.4.5.3.4389.334 where no binding exists with this exact OID but 1.3.5.3 does exist and is configured this way will use the remainder of the string (4389.334 in this example) to correlate. The correlation value that is extracted from the binding (whether By Index or *By Value*) is only relevant if this particular binding is configured as a key binding within the event definition (see Using Extended Event Definitions on page 333 for an explanation of key bindings) and this is used for event to alarm correlation, not other forms of correlation (see Correlations on page 330 for an explanation of the different forms of correlation).

Forced Conversion — Indicates whether or not forced conversion should be applied to the data within the variable bindings that are based on this definition. This is false by default. If this is set to false, then it is possible that variable bindings base d on this definition will have data that does not conform to the data type it is supposed to have based on what is stated in the MIB. For example, the MIB might say that a certain variable binding is supposed to be an integer, but there could nonetheless be a device that sends data based on this definition as a string. If this is set to true, then this willenable aconversion mechanism so that it willensure that the data will conform to the data type specified in the MIB.

Variable Bindings — This panel contains a list of additional variable bindings to mine after a trap comes in. The variable bindings present in any SNMP trap is supposed to be specified in the MIB, but sometimes more variable bindings are needed to so that the resulting alarm message contains the necessary contextual information. If more variable bindings are needed, but can only be accessed by using data from one that has already come in then this should be specified here where the variable binding you are editing is the one that is expected to come in the payload of the trap and the binding listed within this panel are then mined using the index of the binding that is based on the definition you are editing. You can push the button *Add Latent Mined Binding* and select a binding from the popup and add it to this list. After you save this screen, the bindings you add will show on the Message Template tab of any event definition that contains this binding in its trap payload based on the trap definition as specified in the MIB. Note that any bindings added to this list are latent mined, meaning that they are only mined during alarm correlation. The data from these bindings can be used in the message template and/or key bindings of the respective event definition, but latent mined bindings cannot be used in filters, nor for extensions, and nor for entity lookup. If you would like to use a mined binding for any of those features, you will have to use critical mining rather than latent mining. Critical mined bindings are defined at the event definition level rather than at the variable binding definition level. Please see Using Extended Event Definitions on page 333 for moreinformation.

Understanding Alarm Propagation to Services and Customers

Thefollowingdescribes the use caseswhere Alarm Propagation servicesand customers occurs. This describes the sequence of events/alarms. See also Variable Binding Definition Editor on page 339 for ways to augment propagation possibilities.

Alarm state must propagate to associated entities for each step and might take some time to reach all of them, so matching mentioned may not beinstantaneous, depending on the complexity of the associations. This propagation to services and customers occurs through a background process, running on regular intervals.

A resource can have several levels of services that depend on it, and then customers can depend on them, and so on. Potentially, several levels of dependency and a large database full of services and customers to propagate alarmstates can exist, so propagation processing occursin the background. By default, this process runs every 30 seconds, but you can configure this interval by setting the com.dorado.assure.propagation.AlarmPropagationInterval property. This value is in milliseconds. For example, the to set the interval to 60 seconds, set this property as follows:

```
com.dorado.assure.propagation.AlarmPropagationInterval=60000
```

To prevent any Cruz upgrades from overwriting this setting, it is best to put it into the following file:

```
\owareapps\installprops\lib\installed.properties
```

After changing thisproperty, make sure that you restart the application server for the change to take effect.

NOTE:

Only services associated with the alarmed subcomponents are affected by alarms on the subcomponent, not services connected to the rest of the device. You can also override default service affecting alarm behavior with an Event Processing Rule. See Automation and Event Processing Rules on page 288 for more details about this editor.

The remainder of this section describes what happens when a new alarm arrives, an existing alarm clears, or user actions execute.

A New Alarm Arrives

When a new alarm arrives, here is what happens:

What Happens	Description
Service Affecting Alarm Changes Source Alarm State:	The new alarm changes the alarm state (higher or lower) of the resource that is its source. **Dependencies:** If this resource has services or customers that depend on it, the alarm state matches for all such deployed, dependent services and their associated customers. Without such dependencies, no alarm state changes, besides that of the source.
Parent Resources:	The alarm changes the alarm state of a child of the source and the alarm's Resource Propagation value is Impacts Subcomponents. **Dependencies:** Child equipment matches the top level’s alarm state. All deployed services and their related customers depending on this particular resource component match the resource component’s alarm state.
Child Resources:	The alarm changes the alarm state of parent of the source and the alarm's Resource Propagation value is Impacts Top Level. **Dependencies:** Parent equipment matches the child entities alarm state. All deployed services and associated customers depending on only this resource’s alarmed component have their alarm state match the resource’s component.
No Change to Alarm State:	The new alarm does not change the alarm state of its source, so no services or customers have their alarm state changed

What Happens	Description
Alarm not Service Affecting:	The new alarm is not service affecting. The result is that no change occurs to services' or customers' alarm state.

Existing Alarm Clears

When an existing alarm clears, here is what happens:

What Happens	Description
Clearing Service Affecting Existing Alarm Changes Alarm State:	This changes the alarm state (higher or lower) of a resource. **Dependencies:** All deployed services and associated customers depending on this resource have their alarm state match the resource.
No Dependencies:	No services or customers change their alarm state
Clearing Non-Service Affecting Existing Alarm:	No services or customers have their alarm state changed

User Actions Execute

When user actions execute, here is what happens:

What Happens	Description
Resync the resource's alarm state:	• If the resource's displayed alarm state was incorrect, perhaps because it is a parent or child of a resource whose alarm state has changed, then this corrects it. • If this action changes the alarm state and this resource's most severe alarm is service affecting, then resync makes alarm states propagate to any associated services and customers. If the deployed services have the incorrect alarm state, then resync corrects that inaccuracy.
Viewing alarms associated with a service:	• If the service is deployed, and the target resource has open service affecting alarms, all open service affecting alarms for the target resource appear. • If the service is deployed, but the target resource has only cleared or non-service affecting alarms against it, no alarms appear. • If the service is deployed, and the target resource does not have open service affecting alarms, but at least one descendent entity of this resource has open service affecting alarms against it, those alarms propagate up to the resource. All open service affecting alarms that propagate up (Resource Propagation is *Impacts top level*) for the target resource's descendants appear • If the target resource does not have service affecting alarms, and neither do any service affecting alarms exist for its descendent entities, no alarms appear. • If the service is undeployed, no alarms appear.
Viewing alarms associated with a given customer:	• If at least one service associated with the customer has open, service affecting alarms, all open service affecting alarms for all services associated with this customer appear. • If none of the services associated with this customer have open, service affecting alarms, so alarms appear

What Happens	Description
User views the services impacted by a particular alarm:	• If the alarmed resource has at least one deployed service that depends on it, all deployed services depending on the alarmed resourceappear. • If the alarmed resource does not have any deployed services that depend on it, no services appear.
Deploying a servicewhose target resource has service affecting alarms:	• Before deploying, no alarms appear for the service. After deploying, all open, service affecting alarms for the target resource appear.
Undeploying a service whose target resource has service affecting alarms:	• Before undeploying, all open, service affecting alarms for the target resource should appear. After undeploying, no alarms appear.
Editing a deployed service to change the target from one resource to another:	• If the original resource has service affecting alarms but the new one does not, all open service affecting alarms for the original target resource should appear before the edit. After the edit, no alarms appear. • If the original resource does not have service affecting alarms but the new one does, before editing, no alarms appear. After editing, all open service affecting alarms for the new target resource appear.

Understanding Event Life Cycle

The following diagram shows the Cruz events processing.

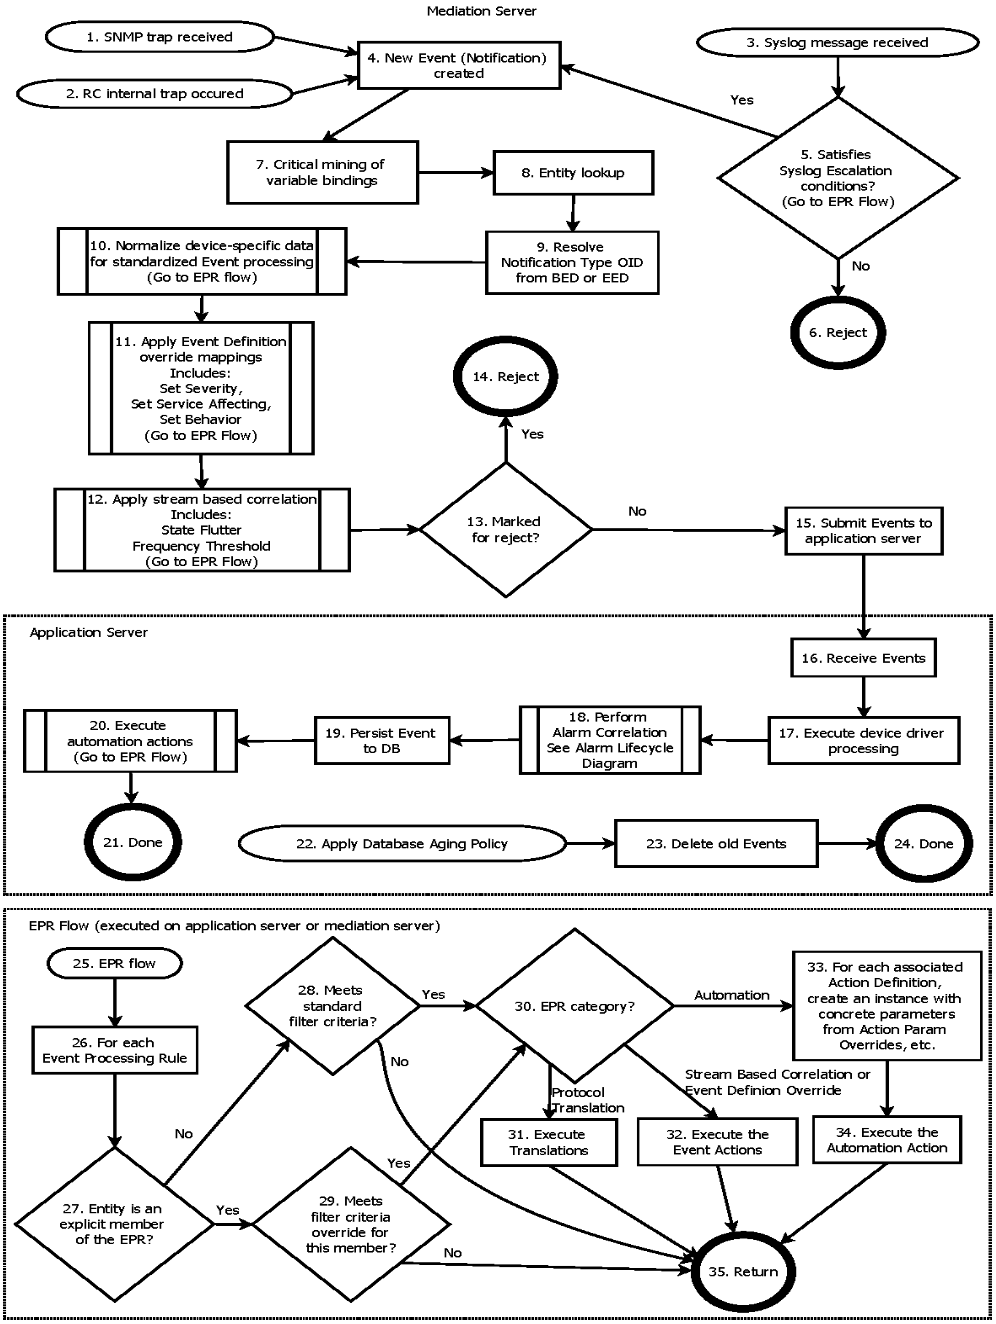
Mediation Server
1. SNMP trap received
2. RC internal trap occured
3. Syslog message received
4. New Event (Notification) created
Yes
5. Satisfies Syslog Escalation conditions? (Go to EPR Flow)
No
6. Reject
7. Critical mining of variable bindings
8. Entity lookup
9. Resolve Notification Type OID from BED or EED
10. Normalize device-specific data for standardized Event processing (Go to EPR flow)
11. Apply Event Definition override mappings Includes: Set Severity, Set Service Affecting, Set Behavior (Go to EPR Flow)
14. Reject
Yes
12. Apply stream based correlation Includes: State Flutter Frequency Threshold (Go to EPR Flow)
13. Marked for reject?
No
15. Submit Events to application server
Application Server
16. Receive Events
17. Execute device driver processing
18. Perform Alarm Correlation See Alarm Lifecycle Diagram
19. Persist Event to DB
20. Execute automation actions (Go to EPR Flow)
21. Done
22. Apply Database Aging Policy
23. Delete old Events
24. Done
EPR Flow (executed on application server or mediation server)
25. EPR flow
26. For each Event Processing Rule
27. Entity is an explicit member of the EPR?
No
Yes
28. Meets standard filter criteria?
Yes
No
29. Meets filter criteria override for this member?
Yes
No
30. EPR category?
Automation
Protocol Translation
Stream Based Correlation or Event Definion Override
31. Execute Translations
32. Execute the Event Actions
33. For each associated Action Definition, create an instance with concrete parameters from Action Param Overrides, etc.
34. Execute the Automation Action
35. Return

The following items correspond to the numbered processes in the diagram.

Process	Description
1 SNMP trap received	The server received an SNMP trap.

Process	Description
2 Cruz internal trap occurred	An internal trap occurred within Cruz. Many situations are considered internal traps that emit an Event. One example is when a monitor polls a certain target and retrieves data for an attribute making the attribute cross a severity threshold. This emits a monitorAttributeTrend Event.
3 Syslog message received	The server received a Syslog message.
4 New Event (Notification) created	Cruz creates a new Event (“Notification”) from the data received. Such Events are specific to Cruz’s internal processing.
5 Satisfies filter conditions?	Syslog Event Processing Rules (EPRs) handle received Syslog messages to determine whether to convert the message into a Cruz Event or discarded. Such EPRs also determine how the Syslog data creates the new Event (what severity, message, and so on, it has).
6 Reject	If the received Syslog message does not satisfy the filter conditions of the Syslog EPRs then it is rejected, and no Event occurs.
7 Critical miningof variable bindings	Sometimes Cruz must return to the device that sent the original trap to “mine” additional variable bindings. This mining can be either critical or latent, depending on whether Cruz needs the additional information early in the Event life cycle or whether it can wait until after an Alarm has been created. If Cruz needs additional data to associate an Event to an entity (which happens in # 8. Entity lookup) then this is critical. On the other hand, if Cruz only needs this additional data for an Alarm message, then this would be better configured as latent mining of variable bindings of the Alarm. Critical mining is configured through the `<MineVarBindObjects>` tag within `eventdefs.xml` files.
8 Entity lookup	By default, Cruz associates Events with the IP address of the device from which the original trap (or other kind of message) came. The entity lookup process associates the Event to a subcomponent or other type of entity if necessary, based on the available variable bindings. Sometimes Cruz must configure Event definitions to associate Events based on them to the appropriate entities. You can configure these through the `<EntityLookup>` tagwithin `eventdefs.xml`files. See Entity Lookup on page 329 for details.
9 Resolve Notification TypeOID from BED or EED	The notification type OID is set on the Event either from the matching base event definition (BED) or if there is an extended event definition (EED) whose extension bindings match those in the Event then this notification type OID is used instead. All default properties (such as severity, behavior, etc.) and EPRs associated with the resolved Event Definition are used, whether this happens to be a BED or EED.
10 Normalize device-specific data for standardized Event processing	If the Event came from a device-specific trap, then an Event Processing Rule of type Device Access may exist that can normalize the payload. This then facilitates standardized Event processing. For example, Cruz can convert an Event based on a Cisco-specific definition reporting on a failed login to a generic Event that still reports on the failed login but is not Cisco-specific. Cruz applies any EPRs whose filter conditions match the incoming Event here.
11 Apply Event Definition override mappings	By default, the Event Definition determines the data attributes of an Event, but some Event Processing Rules can override these defaults. Such EPRs can override attributes like severity, service affecting, and behavior. Cruz applies any such enabled EPRs whose filter conditions match the incoming Event here.

Process	Description
12 Apply stream base correlation	Cruz considers frequent Events a stream that might correlate within the steam base. Users can control stream base correlation by creating Event Processing Rules to detect patterns in the frequency of incoming Events for the purpose of minimizing the number of Events submitted to the application server for further processing. This includes EPRs of type State Flutter and Frequency Threshold. Cruz applies any such enabled EPRs whose filter conditions matches the incoming Event here. If an Event matches a steam base correlator then this step might be the end of its processing.
13 Marked for reject?	Is the behavior of the Event “Reject”? If so, then Cruz does insert it into the database. In most cases rejected Events do not go to the application server for further processing if, but some Events require a trigger for correct system behavior and therefore must be processed by the system even when rejected. Cruz posts such Events to the application server for further processing but does not insert them into the database.
14 Reject	Cruz rejects the Event rather than inserting it into the database. Note that even rejected Events might have some effect, like triggering device driver processing, despite not being persisted to the database. See # 13. *Marked for reject?*
15 Post Events to application server	In distributed environments, much of this processing is going on within the mediation server. Events that make it through to this point are posted to the application server so that they can be inserted into the database and for additional processing to take place that depends on data that needs to be queried from the database, including Alarm Correlation.
16 Receive Events	Application server receives Events from the mediation server.
17 Execute device driver processing	Some managed devices have drivers that feature follow-up Event processing. If the Event originated from a such a device then this step executes that follow-up processing.
18 Perform Alarm Correlation	Go to step one of the Understanding Alarm Life Cycle diagram. This process might create, edit, or clear an Alarm.
19 Persist Event to DB	Saves the Event to the database for future reference.
20 Execute automation actions	If any automation Event Processing Rules satisfy the filter conditions of the Event, then the Cruz application executes the associated actions here. Actions may include sending an email, forwarding the Event northbound as an SNMP trap, signifying that the configuration was changed on the device, and so on.
21 Done	Processing on the application server is done for this Event.
22 Apply Database Aging Policy	This begins the process that applies Database Aging Policies (DAP) to delete old Events.
23 Delete old Events	Delete the old Events according to the active DAP for Events.
24 Done	Done applying database aging policies for Events.

Understanding Alarm Life Cycle

The following diagram shows the Cruz alarm processing flow.

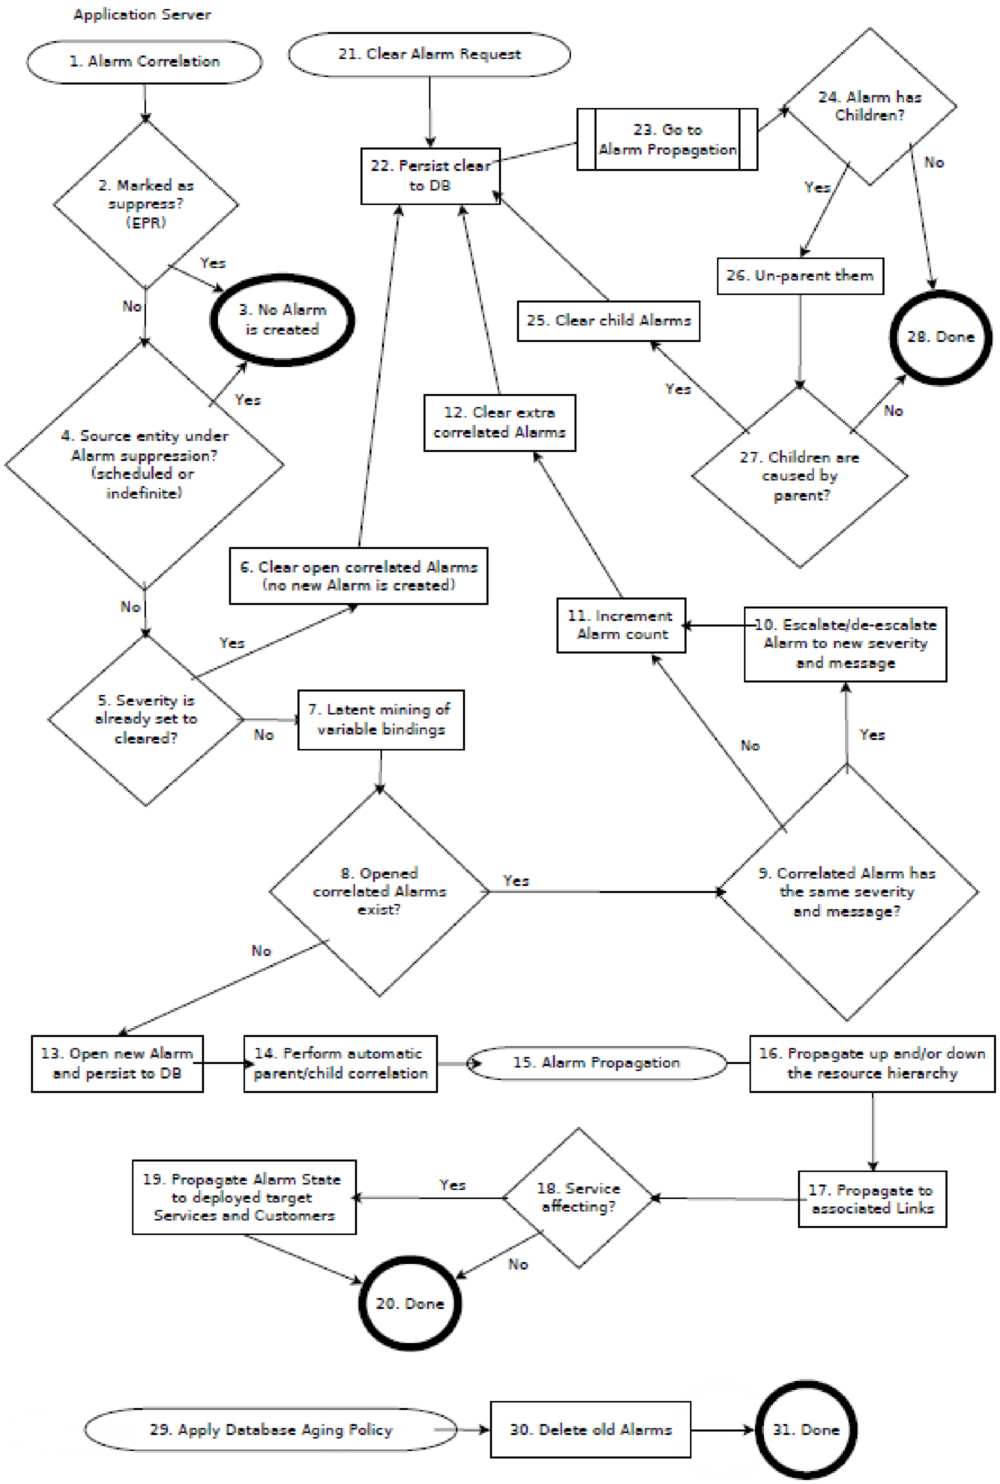
Application Server
1. Alarm Correlation
21. Clear Alarm Request
22. Persist clear to DB
23. Go to Alarm Propagation
24. Alarm has Children?
2. Marked as suppress? (EPR)
Yes
No
3. No Alarm is created
26. Un-parent them
Yes
No
28. Done
25. Clear child Alarms
4. Source entity under Alarm suppression? (scheduled or indefinite)
Yes
12. Clear extra correlated Alarms
Yes
No
27. Children are caused by parent?
6. Clear open correlated Alarms (no new Alarm is created)
No
11. Increment Alarm count
10. Escalate/de-escalate Alarm to new severity and message
Yes
5. Severity is already set to cleared?
No
7. Latent mining of variable bindings
No
Yes
8. Opened correlated Alarms exist?
Yes
9. Correlated Alarm has the same severity and message?
No
13. Open new Alarm and persist to DB
14. Perform automatic parent/child correlation
15. Alarm Propagation
16. Propagate up and/or down the resource hierarchy
19. Propagate Alarm State to deployed target Services and Customers
Yes
18. Service affecting?
17. Propagate to associated Links
No
20. Done
29. Apply Database Aging Policy
30. Delete old Alarms
31. Done

The following items correspond to the numbered processes in the diagram.

Process	Description
1 Alarm correlation	An Event occurred and was not marked as *Reject*.
2 Marked as suppress?	Is the behavior of the Event Suppress? At this point, this might occur because of the default behavior of the Event Definition or possibly because of an override mapping event processing rules (EPR).
3 Event is persisted but creates no Alarm	If Cruz suppresses the Event, then it is inserted to the database but the insertion creates no Alarm
4 Source entity under Alarm suppression?	Is the entity that is the source of the Event under Alarm suppression? This can be scheduled or indefinite.
5 Severity is already set to cleared	Is the severity of the Event set to Cleared?
6 Clear open correlated Alarms	Clear all open Alarms correlated to this Event. Correlated implies more than one meaning for Alarms and/or Events: this correlation refers to rising/clearing correlation. How the Event Definitions associated with these Alarms/Events are correlated to each other is what drives this.
7 Latent mining of variable bindings	Cruz must sometimes return to the device that sent the original trap to “mine” additional variable bindings. Such mining can be either critical or latent, depending on whether this additional information is needed early in the Event life cycle or whether it can wait until after Cruz creates an Alarm. If Cruz needs additional data to associate an Event to an entity (which happens in # 8. Entity lookup in the Understanding Event Life Cycle diagram) then mining is critical, as described below. Critical mining of variable bindings also within the Event Life cycle diagram. On the other hand, if this additional data is only needed to for an alarm message, then it this would be better configured as latent. This drives off the <MineVarBindObjects > tag of the bindobjectdefs.xml files.
8 Opened correlated Alarms exist?	Are there any open Alarms that correlate to this Event? Here, “correlated” refers to the correlation of a single open Alarm to one or more Events. The first of these was the Event that made the Alarm open. If the original Alarm is still open, each subsequent correlated Event does not open a new Alarm, instead incrementing this Alarm’s count. Events correlate to existing open Alarms provided it meets all of the following conditions: 1) It is based on the same Event Definition as the Alarm 2) It has the same values for all key bindings as the Alarm 3) It is associated with the same entity as the Alarm.
9 Correlated Alarmhas the same severity and message?	Are the severity and the message associated with this new Event the same as those of the correlated Alarm?
10 Escalate/de-escalate to new severity and message	Escalate or de-escalate the correlated Alarm by editing its severity and message to match that of the new Event.
11 Increment Alarm count	Increment the count of the correlated Alarm.
12 Clear extra correlated Alarms	Only one open correlated Alarm can exist. It is unusual, although not impossible, for more than one open Alarm to correlate to each other. If this happens it would probably be due to concurrency issues across multiple application servers. If this does happen then one of the Alarms can stay open and Cruz clears the others.
13 Open new Alarm and persist to DB	Open a new Alarm and persist it to the database so that it can be queried later.

Process	Description
14 Perform automatic parent/child correlation	If your configuration designates open Alarms the parent or child of a new Alarm, then Cruz automatically creates the parent/child relationship. This derives from configurations made within the Event Definitions that are the basis of the Alarms. You configure this through the *CorrelatedParentData* tag of the eventdefs.xml files.
15 Alarm Propagation	This is the start of process propagating Alarm States to associated entities.
16 Propagate up and/or down the resource hierarchy	Depending on the Alarm's resource propagation attribute, it might propagate Alarm States to the subcomponents of the alarmed entity and/or to the top-level device.
17 Propagate to associated Links	Propagate the Alarm States to the Links associated with the source entity. Cruz compares the severity of the A and Z endpoints and sets the severity of the Link to the higher of the two.
18 Service affecting?	Is this Alarm service affecting? This will either come from the default behavior of the Event Definition or else a mapping Event Processing Rule (EPR) may override it.
19 Propagate Alarm State to deployed Services and Customers	Propagate the Alarm State to the deployed Services and Customers associated with the source entity. This uses the severity calculator configured for the association type being used to route the propagation. This occurs one association route at a time, where if one routing and calculation results in a change in severity of the target entity then it will find the targets associated with that entity and route to them to do another round of calculation. This propagation is recursive and only stops once there are no more target entities whose severity has changed as a result of the calculation.
20 Done	Processing on the application server is done for this Alarm.
21 Clear Alarm request	The user manually clears an Alarm.
22 Persist clear to the database	The database is updated to make this Alarm cleared.
23 Execute Propagation	Go to # 15. Propagate Alarm States to the associated entities.
24 Alarm has children?	Does the Alarm have children? This includes other Alarms that are caused by or blocked by this one.
25 Clear child Alarms	Clear all of this Alarm's children.
26 Un-parent child Alarms	Update all of this Alarm's children so th at they are top-level Alarms (no longer children of any other Alarm).
27 Children are caused by parents?	Are the child Alarms caused by the parent? This comes from the correlation state of the child Alarms. There is more than one sense of what it means for Alarms and/or Events to be "correlated" so to clarify in this context the "correlation state" refers to Alarm parent/child correlation.
28 Done	Processing on the application server is done for this Alarm.
29 Apply Database Aging Policy	This is the start of the process that applies Database Aging Policies (DAP) to delete old Alarms.
30 Delete old Alarms	Delete the old Alarms according to the active DAPs for Alarms. Note that there can be more than one active DAP, which might include one to delete old cleared Alarms and possibly another to delete very old Alarms that are still open, if this is desired.
31 Done	Done applying database aging policies for Alarms.

7

Performance Monitoring

This section focuses on the following performance monitoring tasks a user performs from the portal:

Performance Monitoring Overview

This section describes Resource Monitors as they appear in the Cruz web portal. The following describes these monitors:

- Application Server Statistics
- Resource Monitors
- Top N [Assets] (pre-configured monitor portlets that come with your installation by default.

Finally, this chapter contains a reminder about scheduling refreshes of monitor target groups. See Scheduling Refresh Monitor Targets on page 412.

If you see a monitor documented here that is unavailable in your installation, you may not have purchased it with your system’s package. Consult your salesrepresentative if you need to have the monitor that appears in documents, but not your installation.

NOTE:

If you configure your installation with multitenancy, you may be unable to edit Monitoring, except in the central, Cruz site. Collected statistics may be visible in the customer's domain as dashboards, but the configuration of the underlying monitor is not. Therefore, do not include a monitor portlet in the customer site template. Multitenancy does limit data dashboards and monitors display to the assets visible to the tenant sites if you create the dashboard on the tenant site.

Understanding Performance Monitoring

This chapter contains the following step-by-step instructions for these features:

- Create a Server Status Monitor Dashboard
- Create an SNMP Interface Monitor
- Create an ICMP Monitor
- Create a Key Metrics Monitor
- Create an Adaptive CLI Monitor
- Create a Monitor for an External Script
- Create a Monitor Report
- Create a Simple Dashboard View
- Create A Performance Template

You can see Performance Options from a variety of locations by right clicking in Cruz. For example:

- Ports in the Ports portlet
- Interfaces
- Ports/Interfaces in the Details panels lets you *Show Performance*
- Right clicking on any of the above within a Reference tree lets you select Performance Options.
- All Top N [Assets] portlets let you right click for Performance options.

Monitoring Strings

Monitors do not directly monitor string attributes, but you can create an extractive Adaptive CLI monitor that responds to string values in devices. See Example 5: Monitor Text Values on page 583 for an example.

Monitor Life Cycle

Diagram 1 and Diagram 2 illustrate the performance monitoring processes.

Diagram 1

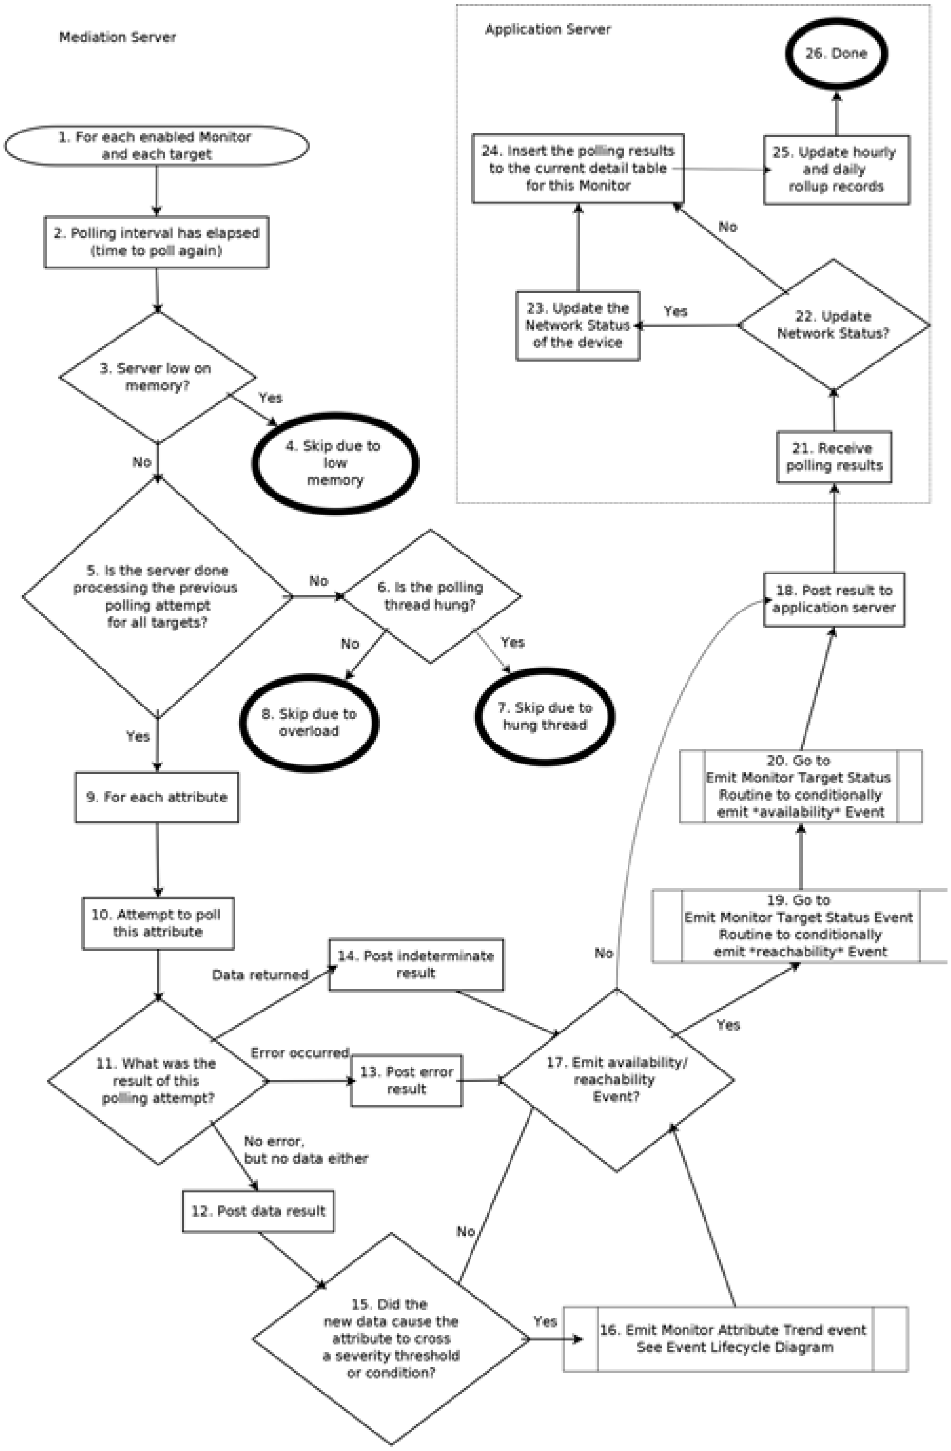
Mediation Server
Application Server
1. For each enabled Monitor and each target
2. Polling interval has elapsed (time to poll again)
3. Server low on memory?
Yes
4. Skip due to low memory
No
5. Is the server done processing the previous polling attempt for all targets?
No
6. Is the polling thread hung?
No
8. Skip due to overload
Yes
7. Skip due to hung thread
Yes
9. For each attribute
10. Attempt to poll this attribute
11. What was the result of this polling attempt?
Data returned
14. Post indeterminate result
Error occurred
13. Post error result
No error, but no data either
12. Post data result
15. Did the new data cause the attribute to cross a severity threshold or condition?
Yes
16. Emit Monitor Attribute Trend event See Event Lifecycle Diagram
No
17. Emit availability/ reachability Event?
Yes
19. Go to Emit Monitor Target Status Event Routine to conditionally emit *reachability* Event
20. Go to Emit Monitor Target Status Routine to conditionally emit *availability* Event
No
18. Post result to application server
21. Receive polling results
22. Update Network Status?
Yes
23. Update the Network Status of the device
No
24. Insert the polling results to the current detail table for this Monitor
25. Update hourly and daily rollup records
26. Done

The following items correspond to the numbered processes in diagram 1.

Process	Description
1 **For each enabled Monitor and each target**	This process executes for each item within the Resource Monitors portlet that is marked as Enabled and for each targeted entity.
2 **Polling interval has elapsed**	This process executes every time the Polling Interval has elapsed. For example, if the Polling Interval for a given Monitor is 15 Minutes then it executes that often.
3 **Server low on memory**	Is the server low on memory? This is determined by the property `lower.threshold.limit` which by default is 5. This means that if the available memory on the server is less than 5% (or whatever value this property is set to) then it meets this condition.
4 **Skip due to low memory**	Numbers 4, 7 and 8 all involve skipping rather than trying to poll the targeted entity. If Cruz finds the system low on memory, then it cannot poll the target(s) and skips polling instead. Any time it skips it increments the associated skip count and also updates the associated last time skipped. You can view these attributes can be viewed through the *JMX console* > *oware* > *service=PollingEngine* . See `SkippedExecutionDueToMemoryCount` and `MostRecentSkippedExecutionDueToMemory`.
5 **Is the server done processing the previous polling attemptforall targets?**	Is the server still working on polling targets from the previous time? The polling process begins every time the polling interval elapses, regardless of whether or not the server is still working on the previous round of polling. This means the previous polling attempt has not produced results so the server will have to skip it (due to overload) rather than trying to perform two (or more) polling attempts simultaneously. If the polling interval is short enough then there will inevitably be skips due to overload because the polling process takes time.
6 **Is the polling thread hung?**	Is the polling thread hung? If so it is not doing any work and will not produce any results. There is a maximum duration that a polling thread has to produce results before the polling engine considers a thread hung. The PollingEngine MBean attribute `ThreadInterruptThreshold` which is set to 90000 (15 minutes) by default determines this.
7 **Skip due to hung thread**	Similar to number 4. If a polling thread is hung then it cannot poll the target(s) and will instead skip. Any time it skips it will increment the associated skip count and also update the associated last time skipped. View these attributes through the JMX console > oware > service =PollingEngine. See `HungThreadInterruptedCount` and `MostRecentHungThreadInterrupted`. When the system identifies a hung polling thread it tries to interrupt the thread and reclaim it.
8 **Skip due to overload**	Like number 4. If the server is overloaded, then it cannot poll the target(s) and will instead skip. Any time it skips it will increment the associated skip count and also update the associated last time skipped. View these attributes through the JMX console > oware > service =PollingEngine. See `SkippedExecutionDueToOverloadCount` and `MostRecentSkippedExecutionDueToOverload`.
9 **For each attribute**	Every Monitor has a set of attributes that it will try to periodically poll from each targeted entities.
10 **Attempt to poll this attribute**	Cruz tries to reach the target device and poll the given attributes.

Process	Description
11 **What was the result of this polling attempt?**	Was this target available and reachable? If so, did data return for this attribute? Certain error conditions result from an error occurring, including if the connection failed or was refused or dropped or if there were bad credentials or if there was a device fault. Some conditions result in posting an indeterminate result, including the data context not being found within the device for the attribute being polled.
12 **Post data result**	Cruz posts attribute data it successfully retrieves from the target device. Once it posts this data to the application server and stores it in the database, because data returned, the Monitor Status Summary of the target device becomes *Available* (with a green checkbox icon) in most cases. The only exception occurs when not all requested attributes for the target returned data. If some data returns but one or more attributes is unavailable or not supported by the device then status of the target says *Partial Results* (with a yellow triangle with an exclamation point icon). Note that this step compiles the polling data that came from the target device, it does not post the returned data to the application server nor store it to the database, but it does identify the polling results as successful, as opposed to error or indeterminate.
13 **Post error result**	If an error occurred during the attempt to poll the target, then Cruz posts information about the error. Possible causes: connection to the target failed, was refused, or was dropped, bad credentials, or another device fault. Once Cruz posts the error information to the application server and stores it in the database, the Monitor Status Summary of the target device becomes *Not Available* (with a red "X" icon). Note that this step exists only to take note of the errors that occurred during the last polling attempt. Like step 12, this step does not post any data to the application server nor store it to the database, but simply compiles the results of the last polling attempt.
14 **Post indeterminate result**	Sometimes the attempt to poll the target device is indeterminate. In such situations, neither an error occurred, nor did any data return from the target. If Cruz did not find the data context on the target device for this attribute, or a timeout occurred while trying to reach the device, this can occur. Once Cruz posts the indeterminate information to the application server and stores it in the database, the Monitor Status Summary of the target device becomes *Not Applicable* (with a gray question mark icon). Like step 12, this step does not post any data to the application server nor store it to the database, but simply compiles the results of the last polling attempt.
15 **Did the attribute to cross a severity threshold or condition because of the new data?**	Each attribute can have one or more severity thresholds and/or conditions. For example, if the CPU usage exceeds 95% then its attribute might be at a critical severity level. If an attribute was previously within a certain severity level and the new polling results show that this attribute crossed the threshold into another severity level (for example minor to major or major to critical) then it meets this condition. If new polling results show that this attribute remained in the same severity level (for example, it was critical and remains critical at least for the moment), then it does not meet this condition.
16 Emit Monitor Attribute Trend Event	Cruz emits a monitorAttributeTrend Event. Go to step *1. RC internal trap occurred* of the Understanding Event Life Cycle diagram.
17 Emit availability/ reachability Event	An *Emit Availability* checkbox appears on the *Editing Monitor* popup screen. When you check it, Cruz creates reachability Events.

Process	Description
18 **Post result to application server**	In distributed environments, much of this processing is going on within the mediation server. Once Cruz produces the polling results, it adds them to a queue that the server then reads. In effect, Cruz has posted the results to the application server. Once the application server receives these results, it inserts them into the database for later querying. Monitor this queue through the JMX console > oware > service = MonitorPollingHandlerMBean.
19 **Go to Emit Monitor Target Status Event Routine to conditionally emit reachability Event**	Go to step 27, the Emit Monitor Target Status Event Routine, to consider the *reachability* of the target device. This routine computes whether to emit the reachability Event.
20 **Go to Emit Monitor Target Status Event Routine to conditionally emit availability Event**	Go to step 27, the Emit Monitor Target Status Event Routine, to consider the *availability* of the target device. This routine computes whether to emit the availability Event.
21 **Receive polling results**	Here, the application server receives polling results from the mediation server.
22 **Update network status?**	The Update Network Status checkbox on the *Editing Monitor* popup screen determines whether this condition is met.
23 **Update the network status of the device**	Cruz updates the network status of the device based on the polling results that were created in steps 12, 13 or 14.
24 **Insert the polling results to the current detail table for this Monitor**	Cruz inserts the polling results into the detail table associated with this Monitor so that they can be queried later.
25 **Update hourly and daily rollup records**	Cruz updates the hourly and daily rollup records for this attribute.
26 **Done**	Processing on the application server is done until more polling results are received from the mediation server.

Diagram 2

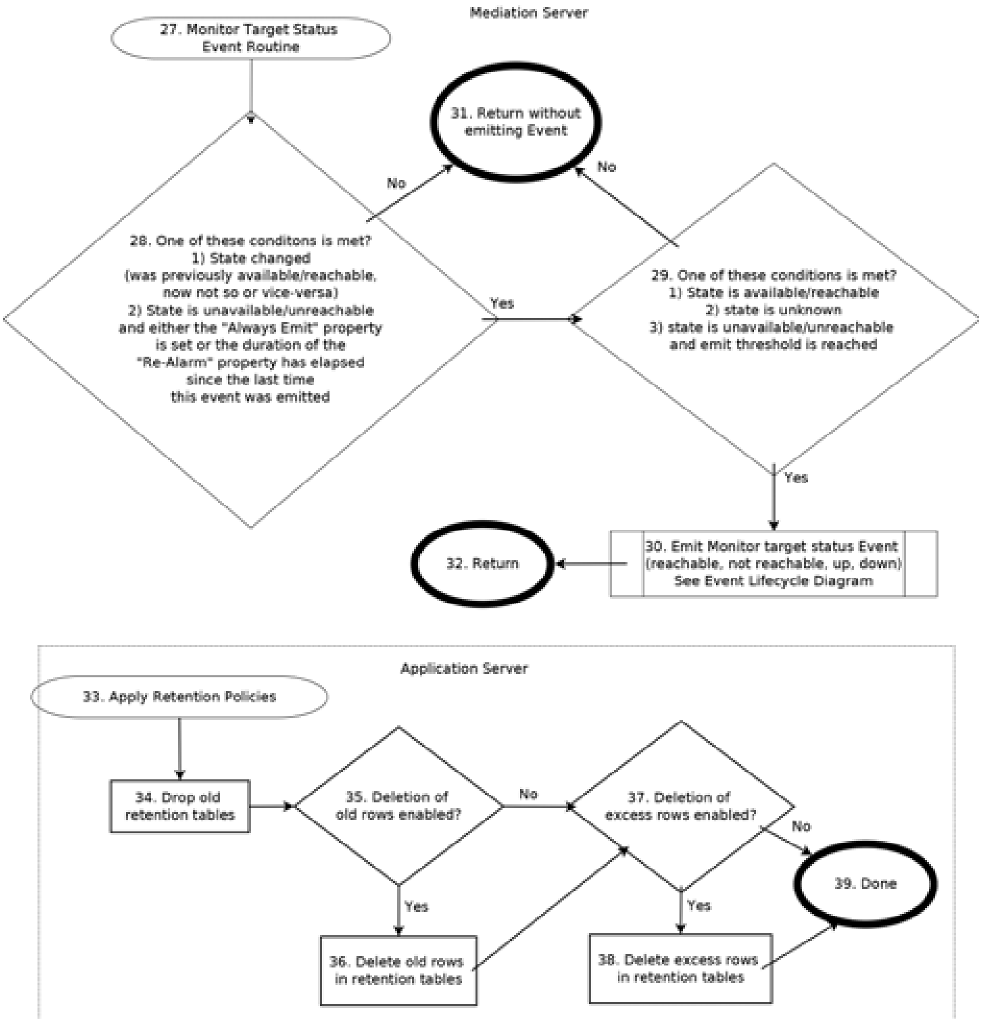
Mediation Server
27. Monitor Target Status Event Routine
31. Return without emitting Event
No
No
28. One of these conditons is met? 1) State changed (was previously available/reachable, now not so or vice-versa) 2) State is unavailable/unreachable and either the "Always Emit" property is set or the duration of the "Re-Alarm" property has elapsed since the last time this event was emitted
Yes
29. One of these conditions is met? 1) State is available/reachable 2) state is unknown 3) state is unavailable/unreachable and emit threshold is reached
Yes
32. Return
30. Emit Monitor target status Event (reachable, not reachable, up, down) See Event Lifecycle Diagram
Application Server
33. Apply Retention Policies
34. Drop old retention tables
35. Deletion of old rows enabled?
No
37. Deletion of excess rows enabled?
No
39. Done
Yes
Yes
36. Delete old rows in retention tables
38. Delete excess rows in retention tables

The following items correspond to the numbered processes in diagram 2.

Process	Description
27 **Monitor Target Status Event Routine**	This routine is executed to compute whether or not to emit the Monitor target status Event (reachabilityEvent or availabilityEvent). There are similar conditions that determine whether these types of Events are emitted.

Process	Description
28 **Is one of these conditions met? State change, always emit, re-emit timeout**	Are any of these conditions met? • The state changed with the last polling attempt. Either the device was previously available/reachable and now it is notavailable or notreachable or vice versa. If the state is the same as before (for example, it was unavailable, and this is still the case) then it does not meet this condition. • The state is unavailable/unreachable and one of the follow two sub-conditions is met: • The `pm.monitor.AlwaysReemitAlarm` propertyis set to true. This property is set to false by default. • The duration of the `pm.monitor.AlarmReemitTimeout` property has elapsed since the last time Cruz emitted a Monitor target state Event for this particular target and Monitor. By default, this property is 30, which means that by default even if a polling target stays down for a period of several hours, it will only emit this Event at most every 30 minutes. Even if you configure the Monitor to poll this target more often than this and even if this target is unavailable every polling cycle, it will still only emit this Event every 30 minutes, unless you change this property.
29 **One of these conditions is met?**	Available, unknown, or threshold reached? (State is unavailable/unreachable and threshold for the number of unreachable/unavailable attempts was reached). This threshold comes from the *# of Unreachable Attempts before update* field on the *Editing Monitor*. Note that this field covers both unreachable and unavailable attempts.
30 **Emit Monitor target status Event**	Emit the Monitor target status Event. If this routine is called to consider the **reachability** of the target, it emits one of these event types: • monitorTargetReachable • monitorTargetUnreachable If this routine is called to consider the **availability** of the target, it emits one of these event types: • monitorTargetUp (available) • monitorTargetDown (not available) See Understanding Event Life Cycle on page 344, the second process (Cruz internal trap occurred) in the event life cycle diagram.
31 **Return without emitting Event**	Return the origin of this routine without emitting an Event.
32 **Return**	Return to the origin calling this routine.
33 **Apply Retention Policies**	Begin the process that applies retention policies to drop tables and/or delete rows from tables.
34 **Drop old retention tables**	Drop the retention tables that are old according to the retention policy for the monitor associated with each table.

Process	Description
35 **Deletion of old rows enabled?**	Most often, the only way old polling data is aged out of the system is by dropping the old tables but there is also a property that controls whether or not to delete old records within the tables without dropping the tables completely. Thispropertyis `pm.retention.DeleteOldDBRecords`and is false by default. This feature was added in an earlier Cruz version to deal with limitations that existed in older database versions that are no longer supported. We strongly recommend keeping this feature disabled because it produces a significant drag on system performance and does not provide any added benefit beyond the dropping of expired tables (and all the data within them), which occurs anyway.
36 **Delete old rows in retention tables**	Deletes the old rows in the retention Cruz tables according to the retention policy for the monitor associated with each table.

37 **Deletion of excess rows enabled?**	Controls whether to delete excess records within the retention tables without dropping the tables completely. This is accomplished using the following porperty, which is false (disabled) by default:
	`pm.retention.PruneRowsInRawDataTablesToMaxCount`
	If you set this property to true, then the value of the the following property is the maximum number of rows that any retention table can have:
	`pm.retention.RawDataMaxCount`
	This feature existed in earlier Cruz versions to deal with limitations that existed in older database versions that are no longer supported. We strongly recommend keeping this feature disabled because it is a significant drag on system performance.

38 **Delete excess rows in retention tables**	Deletes the oldest rows in the Cruz retention tables to get the row count down to the maximum accepted number.
39 **Done**	Done applying retention policies for polling data.

Application Server Statistics

The Application Server Statistics portlet has no expanded view. It displays the statistics for the Cruz application servers and provides access to set logging levels for a variety of categories on application servers.

By default, this portlet is available from the Se ttings page by selecting the Server Configuration menu option.

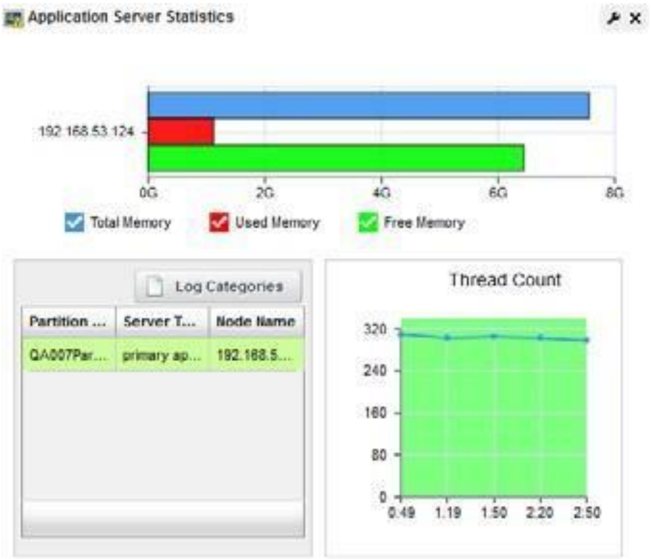

Partition ...	Server T...	Node Name
QA007Par...	primary ap...	192.168.5...

The bar graph displays *Total*, *Used*, and *Free* memory on the server. One such graph appears per server monitored. Hover your cursor over a bar to see its reading in a tooltip. Hover your cursor over the bar graphs related to the server you want to monitor, and its information appears in a tooltip.

The Thread Count graph displays information for as long as this portlet is open, restarting when you revisit it or refresh the page.

Logging Categories

The Application Server Statistics portlet also displaysa table that catalogs servers' *Partition Name*, *Server Type* and *Node Name*. This includes a button the upper right corner where you can access *Log Categories*—`log4j.xml` items—without having to text edit that file. See *Defining a Debug File* on page 140 for more about `log4j.xml`.

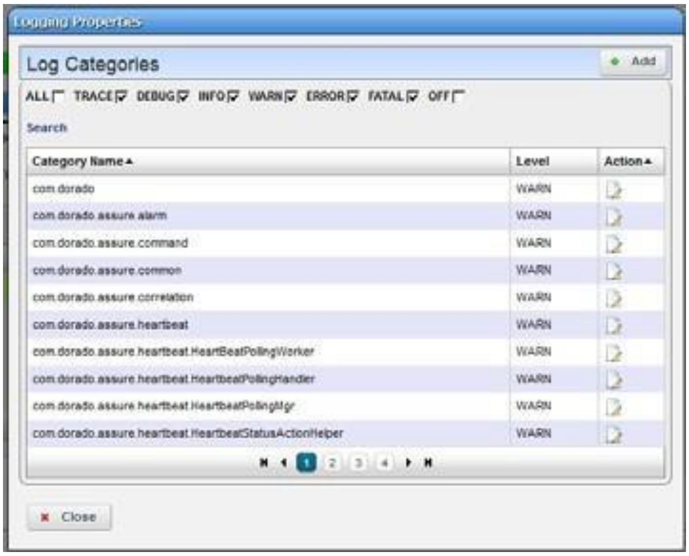

ALL □ TRACE ☑ DEBUG ☑ INFO ☑ WARN ☑ ERROR ☑ FATAL ☑ OFF □

Category Name ▲	Level	Action ▲
com.dorado	WARN	
com.dorado.assure.alarm	WARN	
com.dorado.assure.command	WARN	
com.dorado.assure.common	WARN	
com.dorado.assure.correlation	WARN	
com.dorado.assure.heartbeat	WARN	
com.dorado.assure.heartbeat.HeartBeatPollingWorker	WARN	
com.dorado.assure.heartbeat.HeartbeatPollingHandler	WARN	
com.dorado.assure.heartbeat.HeartbeatPollingMgr	WARN	
com.dorado.assure.heartbeat.HeartbeatStatusActionHelper	WARN	

The `log4j.xml` items appear listed with their default log levels. Altering log levels for the listed items can provide more information for troubleshooting. Log levels determine the detail of server log output.

Notice that you can sort these by clicking the table headings and can look for items with the *Search* link below the checkboxes. You can check or uncheck categories at the top of this screen to confine the display to only desired categories.

These self-monitoring capabilities let you tune Application Server logs to produce meaningful output. Clicking the Edit icon to the right of an item lets you change its log level.

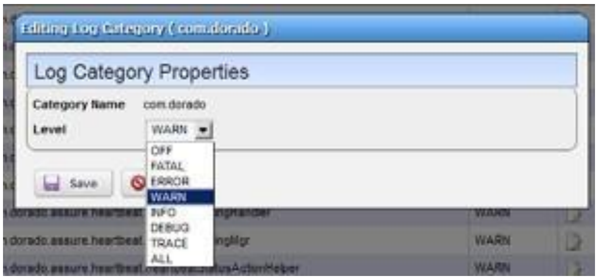

Editing Log Category (com.dorado)

Log Category Properties

Category Name: com.dorado

Level: WARN

Changing log levels in this screen alters log reporting levels for all Application servers, if you have more than one, without restarting them. This simplifies setting log levels and does not require editing the log4j.xml file.

△ CAUTION:
More, and more detailed logging can have a performance impact. See also Understanding Performance Monitoring on page 351.

Resource Monitors

This summary screen displays currently, active performance monitors in brief.

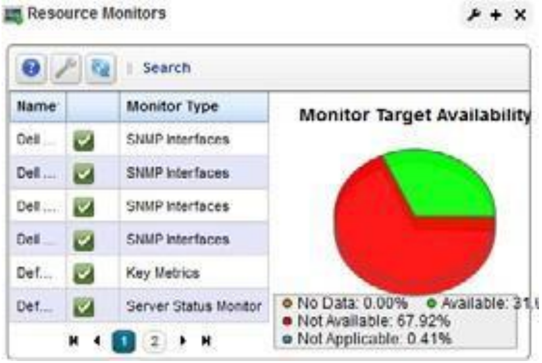

Resource Monitors

Name		Monitor Type
Dell ...	✓	SNMP Interfaces
Dell ...	✓	SNMP Interfaces
Dell ...	✓	SNMP Interfaces
Dell ...	✓	SNMP Interfaces
Def...	✓	Key Metrics
Def...	✓	Server Status Monitor

The *Name* column displays the identifier for each monitor instance, *Enable* displays a green check if it is currently enabled, or a red minus if it is disabled.

The *Monitor Type* column typically displays what the monitor covers. Hover your cursor over this column to see a popup with the selected monitor's properties. The popup that appears after this query displays the relevant information for the monitor, including whether it is *Name, Enabled,* and *Monitor Type*.

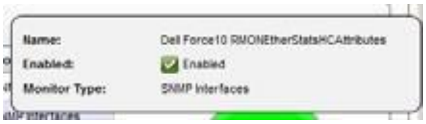

Name: Dell Force10 RMONEtherStatsHCAttributes

Enabled: ✓ Enabled

Monitor Type: SNMP Interfaces

The graph that appears to the right of the monitors displays the aggregate availability information for the enabled monitors. Topics graphed include, *Available, Not Available, No Data* and *Not Applicable.*

Yellowicons mean that not all the data requested was collected. This can occur when MIB attributes have been deprecated. Typically, Cruz monitors include alternative attributes.

Right-click a listed monitor to do the following (not all menu items appear for all types of monitors):

New Monitor —Opens the Select Monitor Type popup, where you can select the type of monitor to create. Pressing continue after doing so will open the Monitor Editor. See Monitor Editor on page 368 for details.

New (from Template)—Opens the Select Template popup, where you can select a template for a new monitor. Pressing continue after doing so will open the Monitor Editor, where you can configure the equipment targets for preconfigured monitor templates of the type(s) selected in the sub-menu. These templates are based on monitors described in Monitor Options Type-Specific Panels on page 388, but have already have selected attributes and calculations. You can examine exactly what these are in the editor that appears when you select one.

Edit Monitor—Opens the Monitor Editor, where you can modify the selected monitor. For a look at the individual monitors' screens, see Monitor Options Type-Specific Panels on page 388.

Details—Opens a Detail panel, with a reference tree, status summary, and general information about the selected monitor.

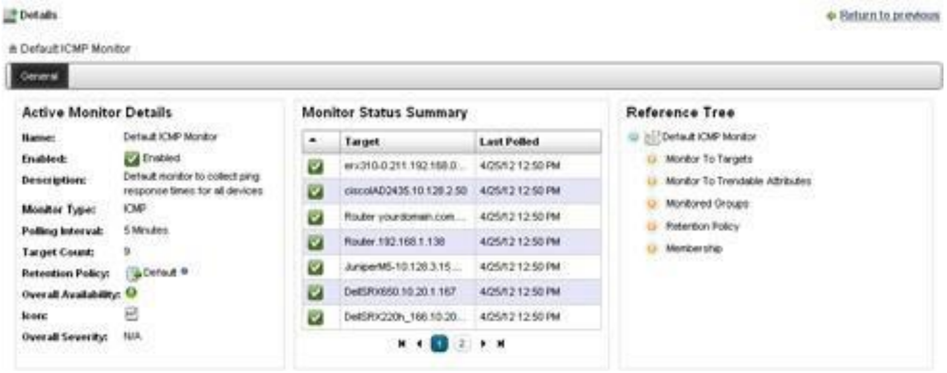

Active Monitor Details

Name: Default ICMP Monitor

Enabled: ✓ Enabled

Description: Default monitor to collect ping response times for all devices

Monitor Type: ICMP

Polling Interval: 5 Minutes

Target Count: 9

Retention Policy: Default

Overall Availability: ●

Icon:

Overall Severity: N/A

Monitor Status Summary

▲	Target	Last Polled
✓	erx310-0.211.192.168.0...	4/25/12 12:50 PM
✓	ciscoIAD2435.10.128.2.50	4/25/12 12:50 PM
✓	Router.yourdomain.com...	4/25/12 12:50 PM
✓	Router.192.168.1.138	4/25/12 12:50 PM
✓	JuniperM5-10.128.3.15...	4/25/12 12:50 PM
✓	DellSRX650.10.20.1.167	4/25/12 12:50 PM
✓	DellSRX220h_166.10.20...	4/25/12 12:50 PM

Reference Tree

Copy Monitor—Copy the selected monitor and its settings to make a new monitor. You must re-name the copy and can change settings selectively.

Enable/Disable Monitor — Enables or disables the monitor. Only one of these options appears. Only enabled monitors report data (and demand resources), while disabled monitors do not.

Refresh Monitor— Re-query to update any targets for the current monitor. See Scheduling Refresh Monitor Targets on page 412 for instructions about automating this.

Manage Retention Policies — Select this to manage the data retention policies for the selected monitor. See Retention Policies on page 366 for details.

Delete — Removes the selected monitor.

View Monitor Data —Viewthe targets' responsiveness to the monitor. Red means unresponsive, green means responsive, and yellow means intermittently responsive.

Viewas PDF —Createsan Acrobat PDF document containing this alarm's contents displayed in the summary portlet.

Import/Export—Export the selected config file to disk or import it from disk. Youcan also import/ export a selected configuration file.

Provides thefollowingactions when availablefor theselectedimage:

- Import retrieves a file containing XML compliance descriptions. Some imports can come from a URL.
- Export Selection exports the selected description to an XML file.
- Export All exports all descriptions to an XML file.

Click Download Export File to specify where to save the file.

The Import/Export option is useful as a backup or to share descriptor s with other projects.

Youmust import data into the correctportlet. Forexample, youcannotimport event data into the Actions portlet.

If one type of data depends on another, you must import the other data before importing the data that depends on it.

Share with User—Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Expanded Resource Monitor

This screen appears when you click the plus in the upper right corner of the summary screen.

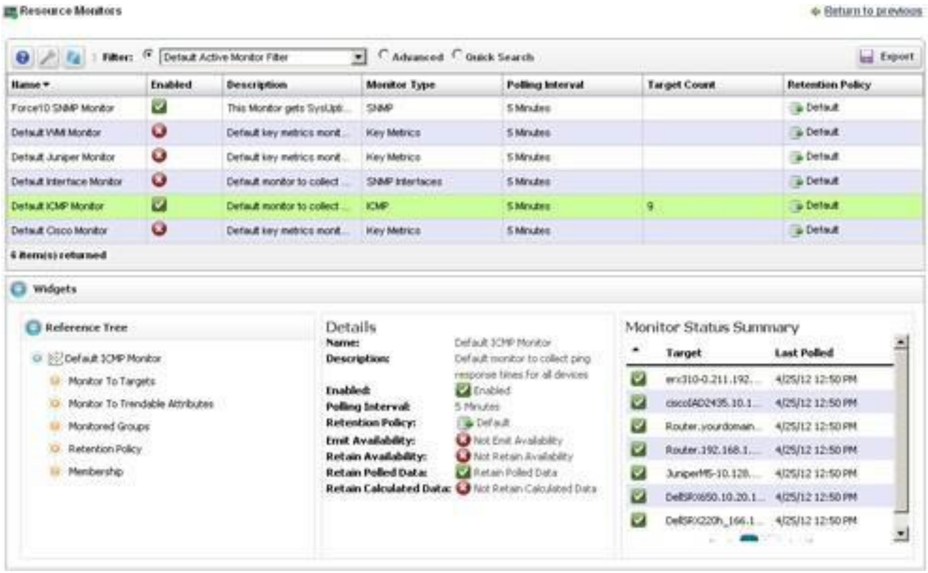

Resource Monitors

Return to previous

Filter: ● Default Active Monitor Filter ○ Advanced ○ Quick Search

Export

Name ▾	Enabled	Description	Monitor Type	Polling Interval	Target Count	Retention Policy
Force10 SNMP Monitor	☑	This Monitor gets SysUpti...	SNMP	5 Minutes		Default
Default VMM Monitor	☒	Default key metrics monit...	Key Metrics	5 Minutes		Default
Default Juniper Monitor	☒	Default key metrics monit...	Key Metrics	5 Minutes		Default
Default Interface Monitor	☒	Default monitor to collect ...	SNMP Interfaces	5 Minutes		Default
Default ICMP Monitor	☑	Default monitor to collect ...	ICMP	5 Minutes	9	Default
Default Cisco Monitor	☒	Default key metrics monit...	Key Metrics	5 Minutes		Default

6 item(s) returned

Widgets

Reference Tree

- Default ICMP Monitor
  - Monitor To Targets
  - Monitor To Trendable Attributes
  - Monitored Groups
  - Retention Policy
  - Membership

Details

Name: Default ICMP Monitor
Description: Default monitor to collect ping response times for all devices
Enabled: ☑ Enabled
Polling Interval: 5 Minutes
Retention Policy: Default
Emit Availability: ☒ Not Emit Availability
Retain Availability: ☒ Not Retain Availability
Retain Polled Data: ☑ Retain Polled Data
Retain Calculated Data: ☒ Not Retain Calculated Data

Monitor Status Summary

▴	Target	Last Polled
☑	erx310-0.211.192...	4/25/12 12:50 PM
☑	ciscoIAD2435.10.1...	4/25/12 12:50 PM
☑	Router.yourdomain...	4/25/12 12:50 PM
☑	Router.192.168.1...	4/25/12 12:50 PM
☑	JuniperM5-10.128....	4/25/12 12:50 PM
☑	DellSRX650.10.20.1...	4/25/12 12:50 PM
☑	DellSRX220h_166.1...	4/25/12 12:50 PM

As in most expanded views, this one displays a list ordered by the *Name* of the monitor. Click *Settings* to configure the column display. Available columns include those on the summary screen (*Name, Enabled, Monitor Type)*as well as *Description, Poling Interval, Target Count* and *Retention Policy*. Menu items are like those described for the summary portlet.

Resource Monitor Snap Panels

When you select a monitor, the Snap Panels at the bottom of the screen display details about it. The *Reference Tree* shows the selected monitor's connection to attributes, groups, retention policies and its membership (the devices monitored).

The *Details* Snap Panel displays the attributes the popup shows when you hover the cursor over the *Monitor Type* column in the summary screen, and adds *Emit Availability* (events), *Retain Availability*, *Retain Polled Data*, and *Retain Calculated Data* parameters.

The *Monitor Status Summary* Snap Panel displays the status ofeach individual member (*Target*) of the monitor, showing the *Last Polled* time and date, and a title bar and icon indicating *Availability* (green is available, red is not).

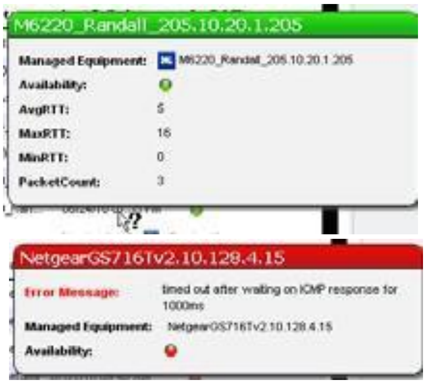

M6220_Randall_205.10.20.1.205

Managed Equipment: M6220_Randall_205.10.20.1.205
Availability: ●
AvgRTT: 5
MaxRTT: 16
MinRTT: 0
PacketCount: 3

NetgearGS716Tv2.10.128.4.15

Error Message: timed out after waiting on ICMP response for 1000ms
Managed Equipment: NetgearGS716Tv2.10.128.4.15
Availability: ●

Hover the cursor over the Availability icon, and a popup appears with details about availability. If the device is available, the *RTT* (round-trip time) for communication appears in *Avg* (average), *Max* (maximum), and *Min* (minimum) amounts, along with the *PacketCount*. If it is not, an *Error Message* appears instead of the *RTT* and *PacketCount* parameters.

To edit more performance settings and targets than are available here, use the features described in Dashboard Views on page 416. You can create and display dashboards by right clicking items in the Managed Resources portlet, selecting *Show Performance*.

Excluding Attributes from Display

The `show.perf.exclude` property in the `portal-ext.properties` file contains a comma delimited list of the attribute display names to exclude from display. Remember, best practice is to override properties as described in Overriding Properties on page 112.

For example,

```
show.perf.exclude=CPU Utilization,AvgRTT
```

If you define this property, the *Show Performance* command creates charts for the listed attributes. This has no impact on manually created dashboards.

NOTE:

You must restart tomcat after changing the properties file for the changes to take effect.

Monitoring Network Availability

In addition to using the standard ICMP monitor, you can create performance monitors that return network availability information displayed on monitor tooltips and reflected in the Network Status column in the Managed Resources portlet.

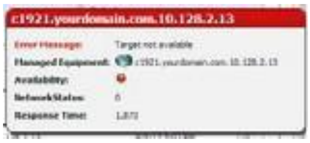

c1921.yourdomain.com.10.128.2.13

Error Message: Target not available

Managed Equipment: c1921.yourdomain.com.10.128.2.13

Availability:

NetworkStatus: 0

Response Time: 1,072

Here are the steps to set that up:

1 Create an Adaptive CLI that monitors some status attribute.
For example, rather than using the built-in ICMP ping on Cruz, you can use a Perl script like this one:

```
use Net::Ping;
$hostname = '10.128.7.11';          # host to check
$timeout = 10;                 # how long to wait for a response
print "-1" if pingecho($hostname, $timeout);
```

2 Makesure the script's network-monitoring attribute return maps to an Adaptive CLI integer attribute named or containing *NetworkStatus*. (NetworkStatus can be just part of the attribute name. It is case insensitive.) You must add this attribute to the Adaptive CLI.

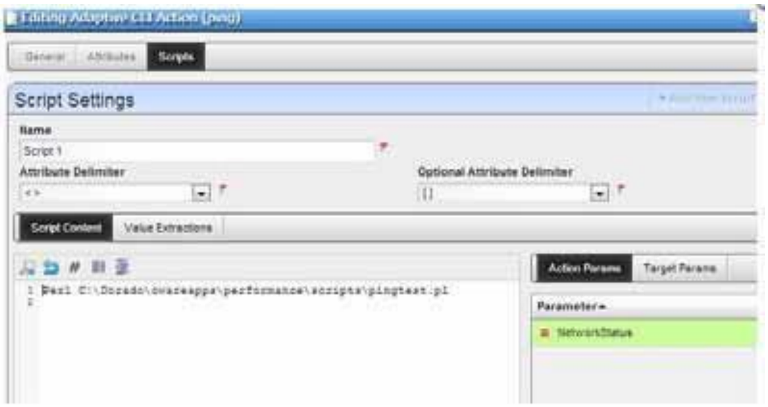

Editing Adaptive CLI Action (ping)

Script Settings

Name: Script 1

Attribute Delimiter: <>

Optional Attribute Delimiter: []

```
Perl C:\Dorado\owareapps\performance\scripts\pingtest.pl
```

Parameter

NetworkStatus

3 Configure the data extraction as follows:

Attribute Name: NetworkStatus

Parse Algorithm: Extract

Parse Expression: (/-*/d)$

4 Create an Adaptive CLI monitor referring to your Ping script Adaptive CLI.

5 Configure it to monitor the attribute in the selected Adaptive CLI as follows:

Attribute Name: Network Status (or your name containing that)

Attribute Type: Integer

Enabled: [check]

Metric Type: Gauge

Also make sure to check the *Update Network Status* checkbox on the monitor.

6 If this attribute exists and the value is 1, polling sets the device’s network availability to *Available*. If its value is 0 it sets availability to *Not Available* and if it is -1 it sets availability to *Indeterminate*.

Retention Policies

The basis of all reporting and dashboard presentations is data retained from monitors. In other words, each monitor provides a simple schema from which you can produce a chart, graph or report.

All monitors relyon a polling engine which execut es device interactions at regular intervals. To reduce resource impacts, you can exclude some of the collected data from what Cruz retains. A monitor could even have no retained data, only emitting events based on transient results in the execution/calculation. Another example: theapplication can derivea metricfromseveralcollected values, and you may opt to retain only the derived result. Monitors may share a retention policy.

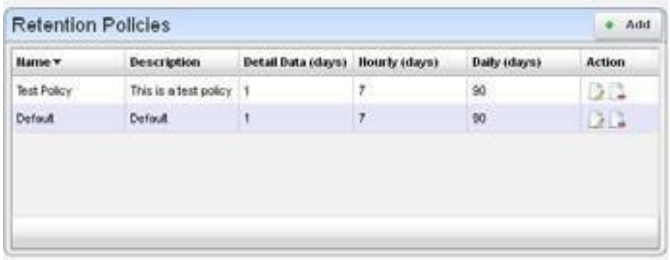

Name ▾	Description	Detail Data (days)	Hourly (days)	Daily (days)	Action
Test Policy	This is a test policy	1	7	90	
Default	Default	1	7	90	

Cruz rolls up data hourly and daily into aggregations. See the Aggregate Data on page 367 for more details. The retention policy controls how long Cruz holds data per aggregation period. You must select the correct period if you want to review what has been collected.

When you select *Manage Retention Policies* in the Monitors portlet, first a list of available policies appears. Clicking the *Add* button at the top of the screen lets you create a new policy, while clicking the *Edit* button to the right of selected, listed policies lets you modify existing policies. The *Delete* button to the right of listed policies removes them from the list.

Editor

Monitors may share a retention policy. The retention policy controls howlong data is held per roll-up period. The editor for Retention policies lets you assign characteristics and monitors to them.

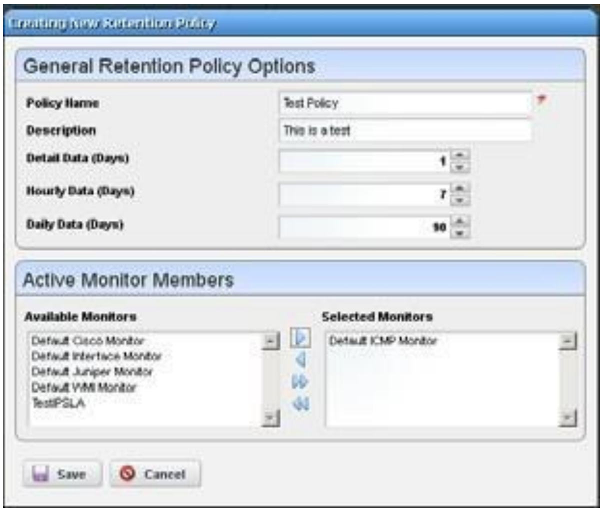
Creating New Retention Policy
General Retention Policy Options
Active Monitor Members
Save
Cancel

Policy Name: Test Policy

Description: This is a test

Detail Data (Days): 1

Hourly Data (Days): 7

Daily Data (Days): 90

Available Monitors:
Default Cisco Monitor
Default Interface Monitor
Default Juniper Monitor
Default WMI Monitor
TestPSLA

Selected Monitors:
Default ICMP Monitor

The editor contains the following fields:

General Retention Policy Options

Policy Name — A text identifier for the policy.

Description — An optional description for the policy.

Detail/Hourly/Daily Data (Days) — How many days to retain the selected data.

> The amount retained has both a performance and data storage impact. For example, retaining day's information from an active performance SNMP monitor configured with one target's worth of data, retrieved on one-minute intervals can consume 0.7 G of database, and require 21 insertions per second. See Understanding Performance Monitoring on page 351 for more about retention policies, monitors and how they impact performance.

Active Monitor Members

Select from *Available Monitors* on the left, and click arrows to move the desired monitor(s) to the *Selected Monitors* on the right.

Click *Save* to preserve your edits, and include the monitor as listed among existing Retention Policies, or click *Cancel* to abandon any changes.

Aggregate Data

Cruz uses detail polling data to produce aggregate values for larger intervals of time, including hourly and daily periods. Here is how this works: at the end of each hour, the detail polling data for the previous hour is aggregated (or as we can say"rolled up") into hourly data. The resulting hourly records are always given the polled time that is the start of the hour, but these data points represent the aggregation across the entire hour. For example, if there is a monitor that polls its targets every 5 minutes, then during the 10 am hour it will poll 12 times and it will have results for each target and for each attribute for each time it polls. Then when the 11 am hour begins, Cruz will aggregate all of the detail polling data collected between 10:00 am and 10:59 am and these new hourly data points will be given the polled time 10:00 am. Likewise, at the end of each day, Cruz will aggregate all hourly data into daily data, and all new daily data points will be given the polled time of midnight of the previous data. These daily data points should be understood to represent an aggregation across the entire day.

Youcanalsoreportonlongeraggregationperiods,including Weekly,Monthly, Quarterly,and Yearly, but these values are computed on-the-fly and are not stored in the database. See Create a Monitor Report for more information about how to generate reports for monitoring data.

Deployment and Polling of Monitor Targets

For each active monitor, the polling targets are deployed and Cruz polls the appropriate attributes every interval. There are two ways of defining the targets for a given monitor: explicit and implicit. Explicit targets are defined by selecting the specific devices that should be polled (equipment managers, subcomponents, etc.) Implicit targets are defined by configuring a monitor to poll a group ofdevices, where the group implicitlyincludes certain devicesbased oncertainfilter criteria. Explicit targets can only be added or removed if you edit the monitor. Implicit targets can be added and removed any time a device is added or removed from inventory or some attribute of a device changes. When implicit targets are removed from a monitor, they are not deleted but they are changed from active to inactive. Only active targets are deployed. Inactive targets are not deployed, but you can still report on the polling data that is associated with inactive targets.

Also, for each monitor target, the Equipment Manager that the target is associated with has a Device Management State, and this data is copied over into an attribute of the monitor target itself. Only targets whose Device Management State is Normal will be deployed. Forexample, if a monitor target is an interface of a device in inventory whose Management State is Decommissioned, then this target will not be deployed. If the Management State is something other than Normal, this target will never be deployed and Cruz will never poll the device, even if the target is active. See Create a Monitor Target Report on page 387 for more information.

Monitor Editor

This editor lets you fine-tune the monitor you selected and right clicked to open the editor. It includes the following panels and fields:

- General
- Monitor Options
- Calculated Metrics
- Thresholds
- Inventory Mappings
- Conditions

General

The General panel is common to all different monitor types.

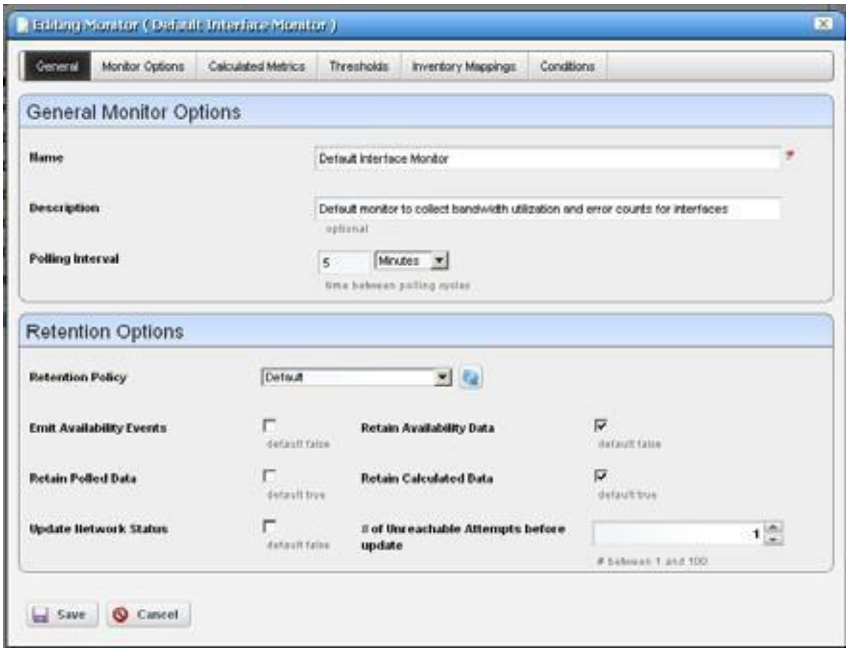

Editing Monitor (Default Interface Monitor)

General Monitor Options

Name: Default Interface Monitor

Description: Default monitor to collect bandwidth utilization and error counts for interfaces (optional)

Polling Interval: 5 Minutes (time between polling cycles)

Retention Options

Retention Policy: Default

Emit Availability Events: □ (default false)

Retain Availability Data: ☑ (default false)

Retain Polled Data: □ (default true)

Retain Calculated Data: ☑ (default true)

Update Network Status: □ (default false)

# of Unreachable Attempts before update: 1 (# between 1 and 100)

General Monitor Options

Name — The identifier for this monitor.

Description — A text description for this monitor.

Polling Interval — Use these fields to configure how often the monitor polls its target(s).

Retention Options

Retention Policy — This configures how long Cruz retains the monitor’s data. Manage these by right clicking in the Resource Monitors portal and selecting *Retention Policies*. You must make retention policies before you can select them here. See also Retention Policies on page 366.

Enabled — Check to enable.

Emit Availability Events —Check to activate emitting availabilityevents. The monitor does not emit an event until the monitored entity’s state has changed. All monitors can generate events on failure to contact the monitored device, port, and so on. For example, by default ICMP monitor updates the network status after a selected number of consecutive failures.

You can configure the monitor to generate an event in addition to updating network status, but Cruz does not like the polling interval to be very small especially when monitoring many devices.

Example: poll every 10 secs for 10,000 devices with Packet Size = 64 bytes, Packet Count = 3 Timeout (secs) = 1 and configure Unreachable attempts = 1 with polling interval = 10 seconds. This polls the device every 10 seconds and emits a “down” event on the first failed attempt.

Retain Availability Data —Check to activate. You must Retain availability data to enable alarms. If you define thresholds, you should retain availability data. *Retainavailability data* stores the Boolean values of whether availability data was in the range your defined metrics.

Retain Polled Data — Check to activate. If you uncheck *Retain polled data* only calculated data remains, you cannot view data retrieve d from monitored entities. Turning off *Retain polled data* discards the data as it arrives from the device.

Retain Calculated Data — Check to activate. *Retain calculated data* complements *Retain polled data*. If checked, it stores the calculated results which came from the raw poll data received from the device.

Update Network Status — Check to activate reporting the network status of the target device(s). The results of this monitor's activity then appear in the Network Status column of the Managed Resources portlet. Only one monitor—and no monitors on interfaces or child components—should ever update networks status. Any monitors on child components or interfaces are rolled up to the top-level device,so status may be erroneously reported. For example, the top-level device is not necessarily down if the interlace is down.

If two monitors report the network status of a single device on different intervals, they must both agree it is down before that state appears in Managed Resources. If one monitor says a device is *Responding*, then that is the state displayed.

If ping fails (an endpoint is down) and update network status is configured, then Cruz tries to ping the switch/router in front of the endpoint to determine if that device is reachable. If that device also failed, then the endpoint's status becomes *indeterminate.*

NOTE:

For clarity's sake, best practice has only one monitor per device updating network status. By default, ICMP monitoring enables *Update Network Status*, and monitors all discovered devices.

Migratingfromprevious Cruzversionsautomaticallyreplacesanyconfigured Heartbeatswith ICMP monitors with *Update Network Status* enabled. If your previous system had HTTP or SNMP heartbeats, you must manually configure monitors to provide equivalent monitoring in this version.

# of Unreachable Attempts before update —The number of attempts to reach the device before Cruz updates the displayed network status of the device.(1-100)

Click *Save* to preserve any edits you make or *Cancel* to abandon them.

Monitor Options

Monitor options contains two panels. The entity panel lets you select the monitor targets. The types of monitor entities allowed varies depending on the type of monitor. The second panel contains options specific to the monitor type being edited.

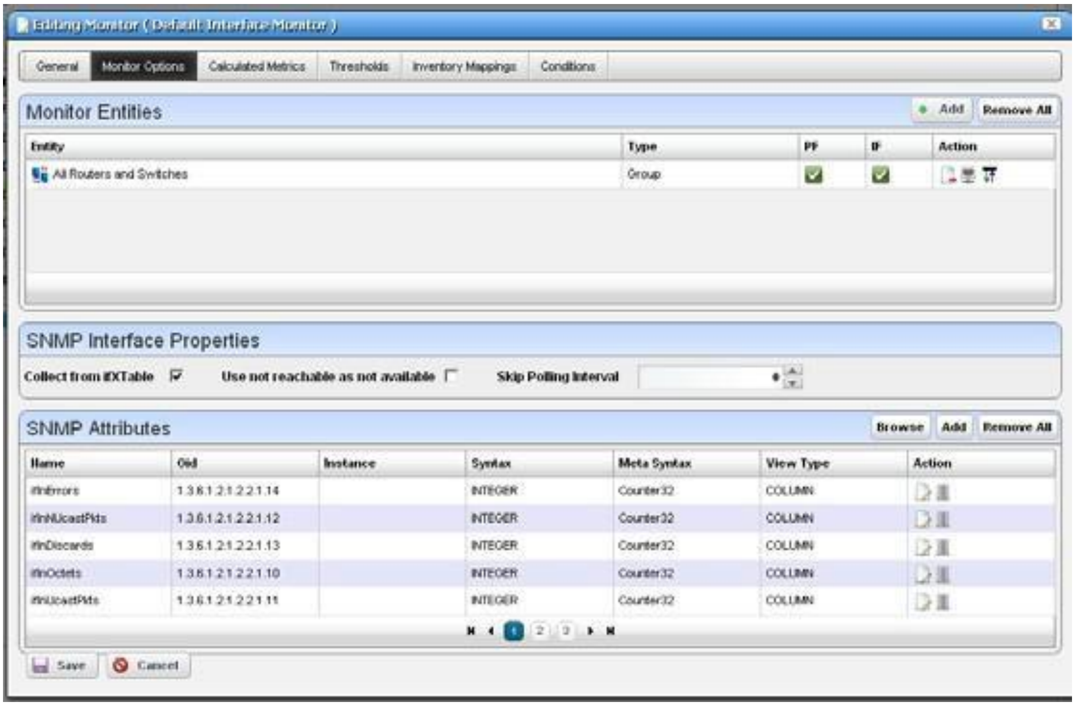

Entity	Type	PF	IF	Action
All Routers and Switches	Group	☑	☑	

Collect from ifXTable ☑ Use not reachable as not available □ Skip Polling Interval 0

Name	Oid	Instance	Syntax	Meta Syntax	View Type	Action
ifInErrors	1.3.6.1.2.1.2.2.1.14		INTEGER	Counter32	COLUMN	
ifInNUcastPkts	1.3.6.1.2.1.2.2.1.12		INTEGER	Counter32	COLUMN	
ifInDiscards	1.3.6.1.2.1.2.2.1.13		INTEGER	Counter32	COLUMN	
ifInOctets	1.3.6.1.2.1.2.2.1.10		INTEGER	Counter32	COLUMN	
ifInUcastPkts	1.3.6.1.2.1.2.2.1.11		INTEGER	Counter32	COLUMN	

The entity and options panels for the various types of monitors appear below in Monitor Options Type-Specific Panels on page 388.

CAUTION:

Have no more than 20,000 targets on a single monitor. Your system may not keep up with polling if you exceed the recommended target limit. Best practice is to poll important devices in shorter intervals and less important devices over longer intervals.

NOTE:

You can elect to monitor the same attribute in multiple, different monitors. This has a performance impact. Best practice is to monitor an attribute only once.

Calculated Metrics

The calculated metrics panel lets you create attributes that are calculated from existing monitor attributes. The metric attribute legend assigns a letter value to each monitor attribute. The *Reassign* button reassigns the letters. This is useful if some attributes have been deleted and their letters are no longer used.

The *Configured Metrics* table lists the calculated metrics. An edit and delete action appear to the right of each row. The *Add* button creates a new calculated metric and the *Remove All* button deletes all the calculated metrics.

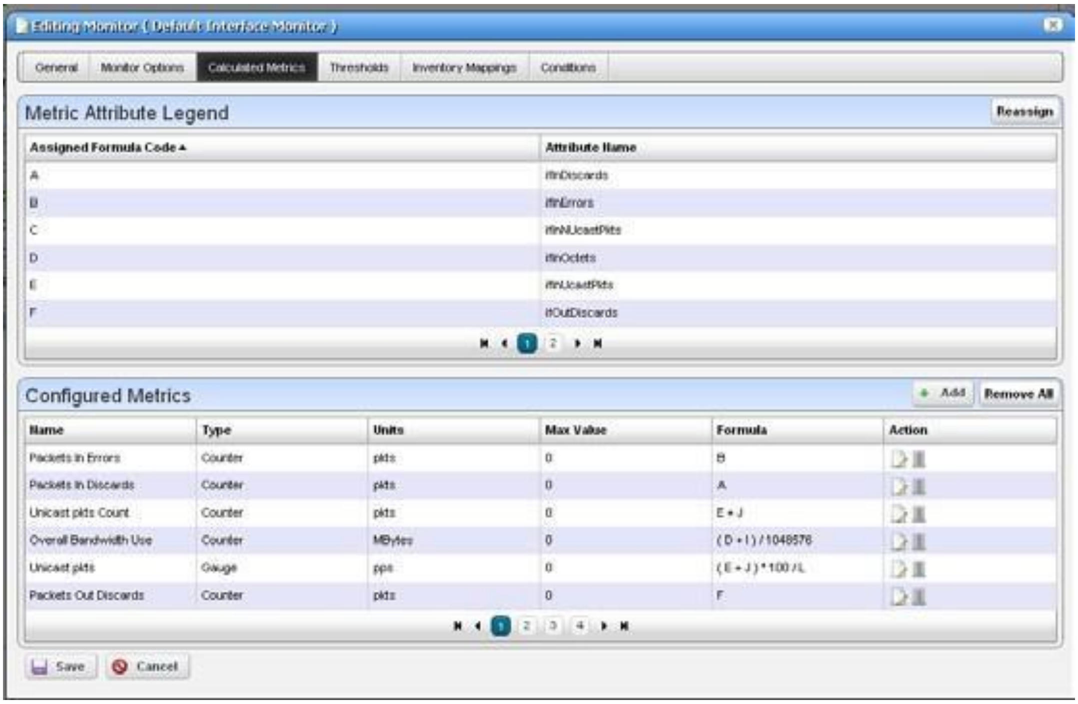

Metric Attribute Legend

Assigned Formula Code ▲	Attribute Name
A	ifInDiscards
B	ifInErrors
C	ifInNUcastPkts
D	ifInOctets
E	ifInUcastPkts
F	ifOutDiscards

Configured Metrics

Name	Type	Units	Max Value	Formula	Action
Packets In Errors	Counter	pkts	0	B	
Packets In Discards	Counter	pkts	0	A	
Unicast pkts Count	Counter	pkts	0	E + J	
Overall Bandwidth Use	Counter	MBytes	0	(D + I) / 1048576	
Unicast pkts	Gauge	pps	0	(E + J) * 100 / L	
Packets Out Discards	Counter	pkts	0	F	

Clicking on the Add button or edit button displays the calculation editor.

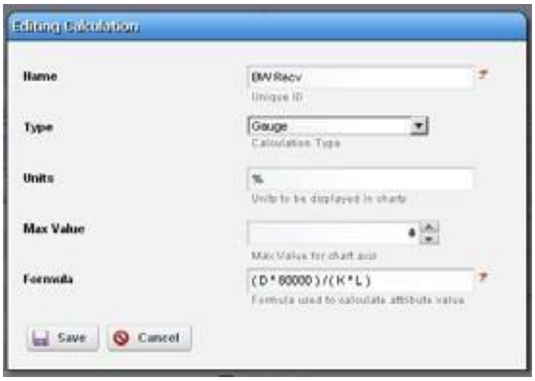

Name: BW Recv
Unique ID

Type: Gauge
Calculation Type

Units: %
Units to be displayed in charts

Max Value:
Max Value for chart axis

Formula: (D * 80000) / (K * L)
Formula used to calculate attribute value

This panel contains the following properties:

Name — The attribute name to be displayed for the calculation

Type — Calculation Type - Gauge or Counter

Units — Units string to appear in graphs. Units do not appear in dashboards with a single attribute.

Max Value — Maximum value to be used in graphing (0 = no max)

Formula—The formula for the calculation using the assigned formula codes from the metric attribute legend.

Thresholds

The thresholds panel allows the user to set threshold intervals on attributes in the monitor. The table lists the attributes for which attributes have been configured. Each row has an edit action and delete action. The Add button allows thresholds to be specified for another attribute. If all monitor attributes have thresholds defined for them the Add button will be disabled.

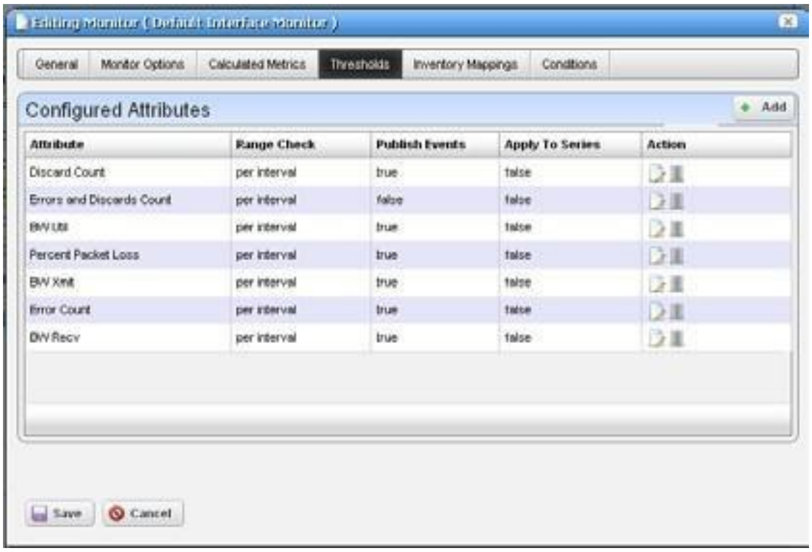

Attribute	Range Check	Publish Events	Apply To Series	Action
Discard Count	per interval	true	false	
Errors and Discards Count	per interval	false	false	
BW Util	per interval	true	false	
Percent Packet Loss	per interval	true	false	
BW Xmit	per interval	true	false	
Error Count	per interval	true	false	
BW Recv	per interval	true	false	

The *Add* or *Edit* buttons open a threshold editor (blank or with existing, configured thresholds, respectively).

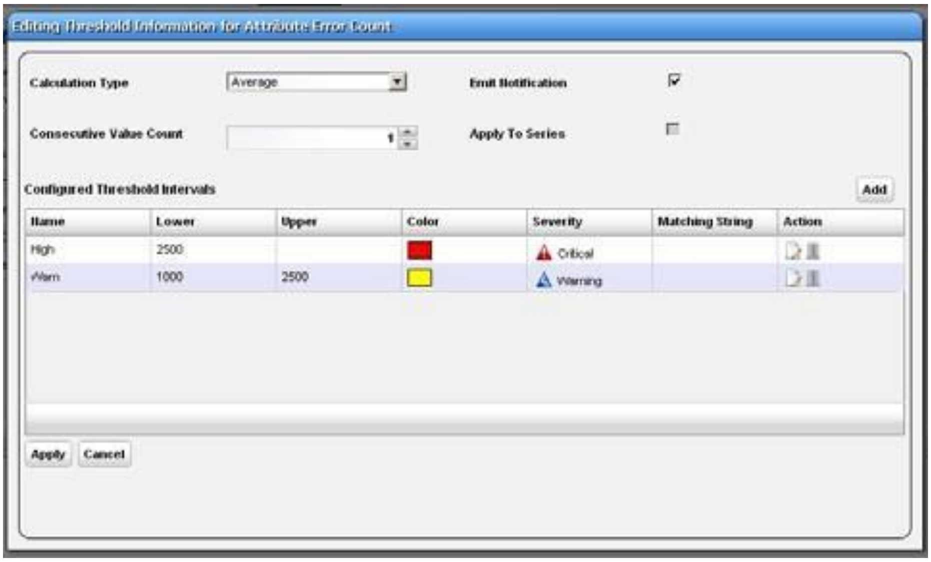

Calculation Type: Average

Emit Notification: ☑

Consecutive Value Count: 1

Apply To Series: ☐

Configured Threshold Intervals

Name	Lower	Upper	Color	Severity	Matching String	Action
High	2500		🟥	Critical		
Warn	1000	2500	🟨	Warning		

Configure threshold intervals you *Add* at in the editor screen according to the following parameters.

Attribute Name— Appears when you click *Add* rather than *Edit*ing a selected threshold. Use the pick list that appears in this screen to select the attribute for which you are specifying threshold information. When you *Edit*, the name of the attribute appears as a title within the editor screen.

Calculation Type—Select from the pick list. Specifies whether the range calculation is to be done based on *Average* or *Consecutive* values.

Consecutive Value Count—Select how many consecutive values to consider at once for a range calculation. Typically, the larger the number here, the less “flutter” in reporting threshold crossings.

Emit Notification—Check to emit an event if the device crosses the configured threshold(s). The notification event contains the threshold-crossing value, as well as which threshold was crossed, and is an alarm at the severity selected when you configure the threshold.

You can make a set of thresholds for each monitored attribute, so a single monitor can throw different alarms for different attributes. Tosee available events and their descriptions, view the contents of the RedcellMonitor-MIBin `\owareapps\performance\mibs`.

Apply to Series—Check to enable on composite attributes only. Checking this applies the threshold to individual elements within the series. When it is unchecked, the threshold applies only to aggregate measurements (the overall value of the series), not individual elements within the series.

For example; a Key Metric monitor for CPU utilization on a device with two CPUs monitors both CPUs. When unchecked, the threshold applies to the average of both CPUs, when checked, the threshold applies to each individual CPU.

Youcan also apply thresholds to regular expressions. This is useful to monitor components within components, for example cores within a CPU.

Click *Apply* to preserve your edits or *Cancel* to abandon them.

The threshold interval editor pops up when you select the *Add* button or the *Edit* icon to the right of a threshold’s row in the threshold attribute editor.

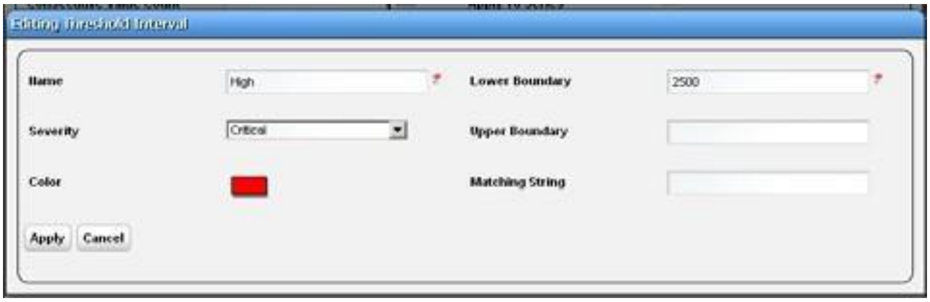

Editing Threshold Interval

Name: High

Severity: Critical

Color: 🟥

Lower Boundary: 2500

Upper Boundary:

Matching String:

This screen contains the following fields:

Name — The identifier for the threshold interval.

Severity—The event severityfor crossing this threshold interval (*informational/indeterminate/warning/minor/major/critical*)

Color — The color to display threshold interval on graphs.

Lower Boundary — The interval’s lower boundary.

Upper Boundary — The interval’s upper boundary. May be blank.

Matching String — A Regex matching string.

NOTE:

You can configure a response to threshold crossing with an Event Processing Rule. Create your thresholds within the monitor and then create an Event Processing Rule whose filter conditions respond to `monitorAttributeTrend` and other conditions such as severity, and so on. You can even use specific values of the event varbinds in the filter conditions too.

Threshold Graph Background

If you configure a set of thresholds, the dashboard graph displaying the data monitored displays the threshold colors (and text label) in the background. When an upper or lower threshold has no upper or lower bound, then those background colors may appear as white.

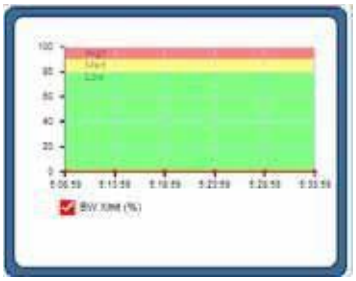

Inventory Mappings

The inventory mappings panel lets you associate predefined inventory metrics with a monitored attribute to normalize the attribute if a device does not report metrics in a way that matches the monitored attribute's name or format. Available metrics include *CPU Utilization %, Memory Utilization %, ICMP Round Trip Time, ICMP packet errors,* and *Bandwidth utilization %.*

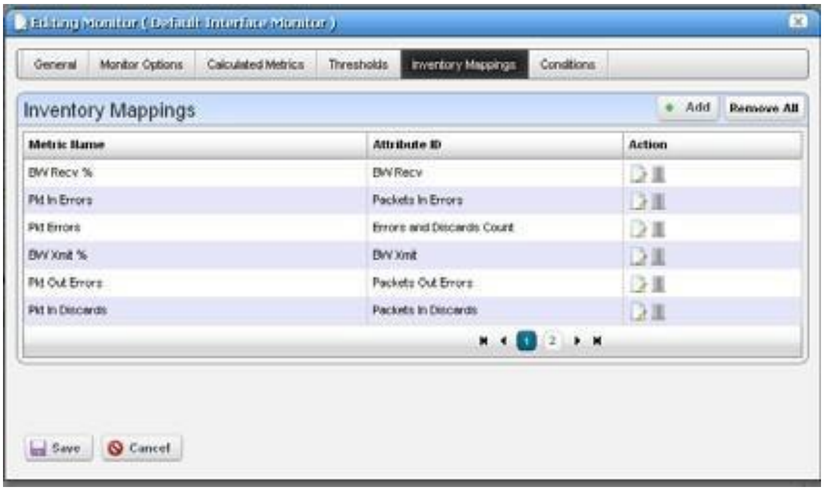

Inventory Mappings

Metric Name	Attribute ID	Action
BW Recv %	BW Recv	
Pkt In Errors	Packets In Errors	
Pkt Errors	Errors and Discards Count	
BW Xmit %	BW Xmit	
Pkt Out Errors	Packets Out Errors	
Pkt In Discards	Packets In Discards	

Common attributes include those for Top N. For example, service A may call it "Disk % Utility" and Service B may call it "% Disk Utility". We can map them to a common name and can display them as Top N.

Youcan *Add* a new mapping with that button or *Remove All* listed mappings with that button. You can also edit or delete listed mappings with the *Action* icons to the right of each row. Adding or editing opens the Inventory Mapping Editor.

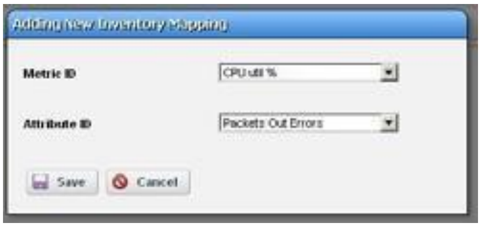

Adding New Inventory Mapping

Metric ID: CPU util %

Attribute ID: Packets Out Errors

This lets you con figure the following:

Metric ID — Inventory metric name

Attribute ID — Attribute to associate with the inventory metric

Conditions

This panel lets you add multiple conditions to the monitor you are editing.

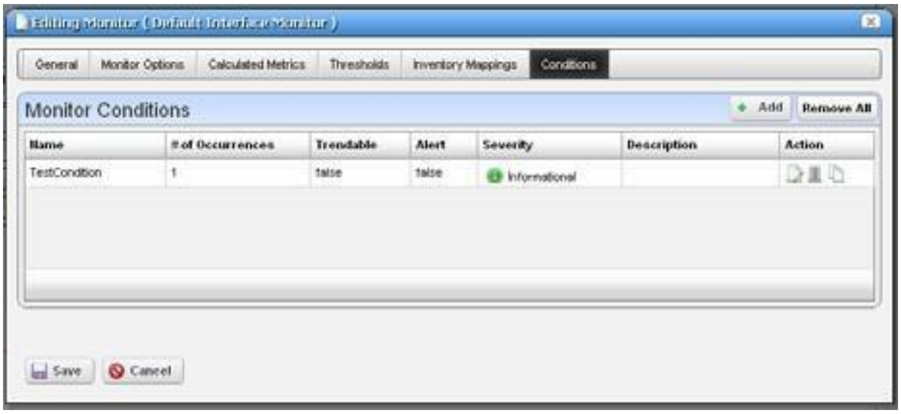

Editing Monitor (Default Interface Monitor)

Monitor Conditions

Name	# of Occurrences	Trendable	Alert	Severity	Description	Action
TestCondition	1	false	false	Informational		

Click the *Add* button to enter a new set of conditions or click the *Edit this entry* button to the right of a listed Monitor Condition to open the editor. Click the *Delete* button to remove a listed set of conditions. Click the *Copy* icon to duplicate the listed condition.

The editor has the following fields and settings to configure:

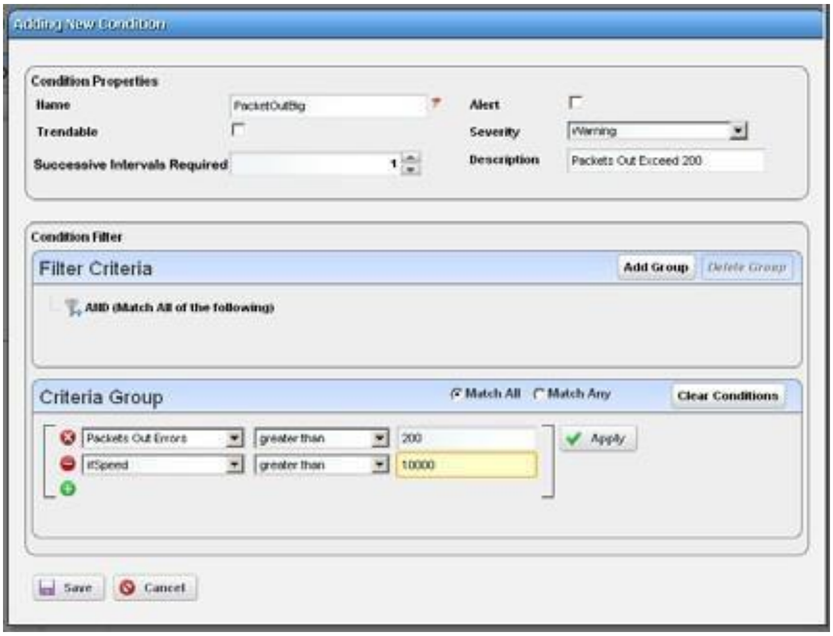

Adding New Condition

Condition Properties

Name: PacketOutBig

Alert: □

Trendable: □

Severity: Warning

Successive Intervals Required: 1

Description: Packets Out Exceed 200

Condition Filter

Filter Criteria

AND (Match All of the following)

Criteria Group

● Match All ○ Match Any

Packets Out Errors	greater than	200
ifSpeed	greater than	10000

Condition Properties

Name — Enter a text identifier for the conditions.

Alert — Check this if you want Cruz to emit an alert when the monitor satisfies the conditions.

Trendable—Check if the conditions specified are trendable. If this is true, the database retains qualifying conditions (or thresholds) for later reporting/dashboards.

Severity — Specify the severity of the emitted alert, if any.

Successive Intervals Required—Enter the number of occurrences of what is specified in the Condition Filter to satisfy the Conditions.

Description — A text description for the conditions.

Condition Filter

Minimally, use this panel to select a condition, an operator and a value. If you want to use the logical AND or OR operators with a second condition, click the green plus (+), and select a second condition, operator and value. For example, *Packet Out Errors greater than 200* AND *ifSpeed greater than 10000* can be a set of conditions that only has to occur once to satisfy this monitor’s condition.

Click *Save* to accept your edits or *Cancel* to abandon them.

Self -Management/Self- Monitoring: Default Server Status Monitor

Cruz also includes a Default Server Status Monitor that monitors its own server(s). Even clustered application and mediation servers are automatically added to this monitor. You can edit this monitor to alter polling intervals and make different calculations for the monitored attributes.

Those attributes include TotalMemory, FreeMemory, MemoryInUse, ThreadCount and TrapCount for Application Server and Mediation Server processes. You cannot modify the targets for this monitor.

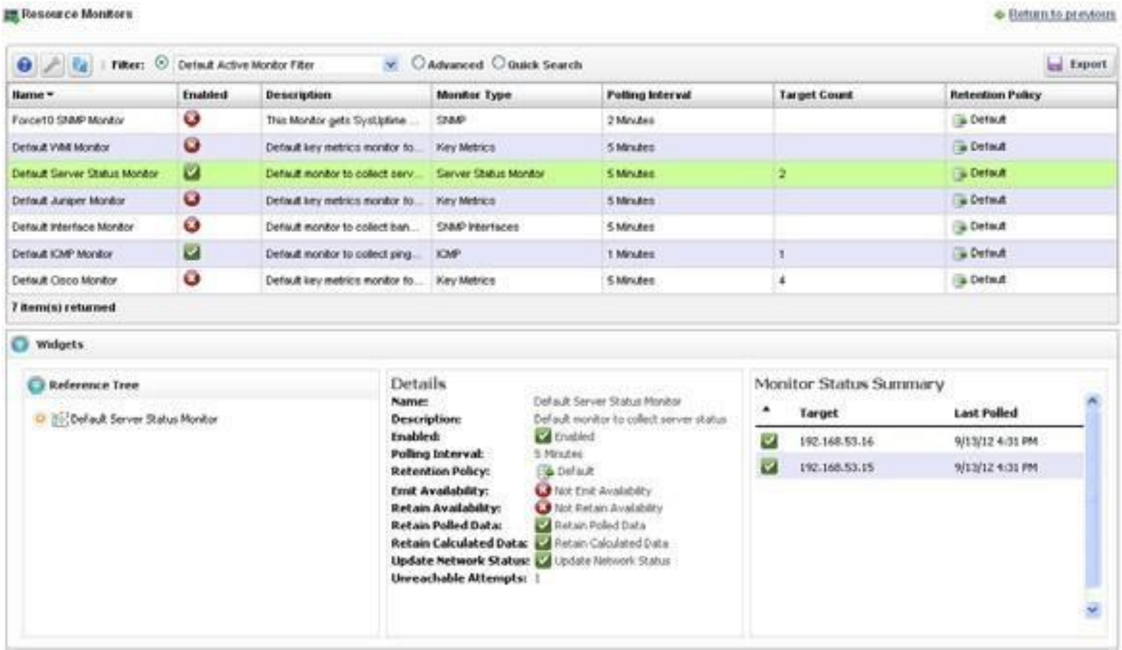
Resource Monitors
Return to previous
Filter: ● Default Active Monitor Filter ○ Advanced ○ Quick Search
Export
Widgets
Reference Tree
Default Server Status Monitor

Name	Enabled	Description	Monitor Type	Polling Interval	Target Count	Retention Policy
Force10 SNMP Monitor	☒	This Monitor gets SysUptime ...	SNMP	2 Minutes		Default
Default WMI Monitor	☒	Default key metrics monitor fo...	Key Metrics	5 Minutes		Default
Default Server Status Monitor	☑	Default monitor to collect serv...	Server Status Monitor	5 Minutes	2	Default
Default Juniper Monitor	☒	Default key metrics monitor fo...	Key Metrics	5 Minutes		Default
Default Interface Monitor	☒	Default monitor to collect ban...	SNMP Interfaces	5 Minutes		Default
Default ICMP Monitor	☑	Default monitor to collect ping...	ICMP	1 Minutes	1	Default
Default Cisco Monitor	☒	Default key metrics monitor fo...	Key Metrics	5 Minutes	4	Default

7 item(s) returned

Details

Name: Default Server Status Monitor
Description: Default monitor to collect server status
Enabled: ☑ Enabled
Polling Interval: 5 Minutes
Retention Policy: Default
Emit Availability: ☒ Not Emit Availability
Retain Availability: ☒ Not Retain Availability
Retain Polled Data: ☑ Retain Polled Data
Retain Calculated Data: ☑ Retain Calculated Data
Update Network Status: ☑ Update Network Status
Unreachable Attempts: 1

Monitor Status Summary

	Target	Last Polled
☑	192.168.53.16	9/13/12 4:31 PM
☑	192.168.53.15	9/13/12 4:31 PM

You must create your own Dashboard to view the data in this monitor. Create a custom dashboard for this as in described below.

How To:
Create a Server Status Monitor Dashboard

1 Create a custom dashboard as described in How to: Create a Custom Dashboard View on page 421.

2 Click the edit icon on one of the dashboard components and set the data source as the Default Server Status Monitor, and the target as the server(s) monitored.

3 Save the monitor

See Dashboard Views on page 416 for more about configuring dashboards.

How To:
Create an SNMP Interface Monitor

To set up a typical performance monitor, follow these steps:

1 In the *Resource Monitors* portlet, and create a new monitor by right clicking and selecting *New*.

2 Select the type of monitor from the submenu—for this example, an *SNMPInterfaces* monitor.

NOTE:

Some devices have ports rather than interfaces. This monitor works for them too, even though it is an “interface” monitor.

3 In the *General* screen, enter a polling interval (5 minutes is the default). For this example, check *Retain polled data* and accept the remaining defaults for checkboxes and the retention policy.

4 Select an entity to monitor by clicking the *Add* button in the top portion of the *Monitor Options* screen. For an interface monitor, select *Interface* as the Type at the top of the screen. You can also filter the list of interfaces that appear further by selecting *Interface Type* as *ge* (gigabit ethernet), for example.

NOTE:

Notice that you can add refinements like filtering on *Administrative State* and *IP Address* to the filter.

5 Select interfaces (Ctrl +click to add more than one), then click *Add Selection* then *Done* to confirm your entity. Hover your cursor over a line describing an interface to have a more complete description appear as a popup.

6 Click *Browse* to display the MIB Browser. For the sake of this example, we elect to monitor ifInErrors (in RFC Standard MIBs, RFC1213-MIB > Nodes > mib-2 > interfaces > ifTable > ifEntry >ifInErrors).

7 In the *Thresholds* screen, configure thresholds by first clicking *Add.*

8 Click *Add* above the threshold levels list for each threshold you want to add.

9 In the threshold editor, enter a name (Examples: *Low, Medium, Overload*), an upper and lower boundary, (0 - 10, 10 - 100, 100 +), a severity (*Informational, Warning, Critical*) and color (BLUE, YELLOW, RED). In this case, no string matching is necessary. When the data crosses thresholds, the monitor reacts.

Attributes available depend on the type of monitor you are creating. Notice that you can also check to make crossing this threshold emit a notification (an alarm that would appear on the Alarm panel). You can also configure the type of calculation, and so on. You can even alter existing thresholds by selecting one then clicking *Edit* to the right of the selected threshold.

10 Click *Apply* for each threshold interval you configure, then *Apply* for the entire threshold configuration.

If a threshold’s counter is an SNMP Counter32 (a 32-bit counter) monitoring can exceed its capacity with a fully utilized gigabit interface in a relativelyshort period of time. The defaults configured in this monitor account for this, but if you know that this is an issue, you can probably configure the monitor to account for it too.

After taking a look at Thresholds no more configuration is required. Notice, however, that you can also configure *Calculated Metrics*, *Inventory Mappings* and *Conditions* on other screens in this editor to calculate additional values based on the monitored attributes, to map them, and to make conditional properties based on monitored behavior.

NOTE:

Calculated Metrics isparticularlyvaluableifyouwanttomonitoracompositelikeifInErrors + ifOutErrors orwanttocalculate aparameter likeerrors perminute when themonitor’s interval is 5 minutes.

11 Click *Save* and the monitor is now active.

Notice that the *Availability* icon appears at the top of a *Monitor Status Summary* snap panel in the Expanded Resource Monitor next to a time/date stamp of its last polling. Right-click the monitor and select *Refresh Monitor* to manually initiate polling.

Values displayed in the Overall Availability column of the Monitor Manager do not automatically refresh and may be out of date. The *Reference Tree* snap panel maps the monitor's relationship to its target(s) attribute(s) and other elements. The *Details* snap panel summarizes the monitor's configuration.

12. For information about having the monitor's results appear in the *Dashboard* portlet, see Dashboard Views on page 416.

How To:

Create an ICMP Monitor

The following steps create an ICMP (ping) monitor.

1. In the *Resource Monitors* portlet, and create a new monitor by right clicking and selecting *New*.
2. Select the type of monitor from the submenu—for this example, an *ICMP* monitor.
3. In the *General* screen, enter a name (Test ICMP Monitor), and a polling interval (5 minutes is the default). For this example, check *Retain polled data* and accept the remaining defaults for checkboxes and the retention policy.
4. Select an entity to monitor by clicking the *Add* button in the top portion of the *Monitor Options* screen.
5. Select devices you want to ping, (Ctrl +click to add more than one), then click *Add Selection* then *Done* to confirm your entity.
6. Define packetsin the ICMP Monitor Options panel, including Packet Size, Packet Count and timeout. You can accept the defaults here, too.
7. In the Thresholds tab, select an attribute (MaxRTT, or maximum round trip time) and add the following thresholds by clicking *Add*:

   Name *High* color red, Lower Boundary 15 and Upper Boundary [blank] Severity *Critical*

   Name *Fine* color green, Lower Boundary 0 and Upper Boundary 15 Severity *Cleared*.

   Notice that this example does not emit a notification. If you checked that checkbox, an alarm of the configured severity would accompany crossing the threshold.
8. Accept the other defaults and click *Apply*
9. Click *Save*.
10. Test ICMP Monitor now appears in the portlet.

How To:

Create a Key Metrics Monitor

Follow these steps to create a Key Metrics Monitor (also, see Key Metric Editor on page 428).

1. In the *Resource Monitors* portlet, and create a new monitor by right clicking and selecting *New*.
2. Select the type of monitor from the submenu—for this example, a *Key Metrics* monitor.

3. In the *General* screen, enter a name (Test Key Metrics Monitor), and a polling interval (5 minutes is the default). For this example, check *Retain polled data* and accept the remaining defaults for checkboxes and the retention policy.
4. Select an entity to monitor by clicking the *Add* button in the top portion of the *Monitor Options* screen.
5. Select devices on which you want to monitor Key Metrics.
6. Select from the available metrics that appear at the bottom of the screen in Key Metric Properties by selecting a category with the pick list at the top of the screen, then click on an *Available* metric, and click the right arrow to make it a *Selected* metric.
7. Click *Save* to retain your new Monitor.
8. Test Key Metrics Monitor appears in the Resource Monitorsportlet.

How To:

Create an Adaptive CLI Monitor

You can create monitors that track Adaptive CLI responses. The following outlines the steps:

1. Determine the Show Command that you want to run.
2. Create ACLI to extract the data as described below.
3. Create the Monitor that uses the data from the ACLI.
4. Create any threshold crossing events/actions (see Thresholds on page 373).
5. Create dashboards (see Dashboard Views on page 416) to view results and reports (see How to: Create a Monitor Report) to preserve or display the data.

Monitor Reports in Multitenant Environments

Reports for the master site can target all available de vices, however, in tenant sites, only devices to which the tenant site has access are visible in reports.

If, for example, a tenant wants to make a report about a monitor shared by all tenants, then the tenant can create the report only for data from devices assigned to its site.

Show Command

For this example, use the Cisco `show ip traffic` command. Run the command so you can see the data you want to extract. Here, we want to know the number of dropped packets due to adjacency and no route issues. Here is some example output:

```
c1720-1.30#show ip traffic
IP statistics:
  Rcvd: 2072045 total, 1995503 local destination
         0 format errors, 0 checksum errors, 0 bad hop count
         0 unknown protocol, 0 not a gateway
         0 security failures, 0 bad options, 0 with options
  Opts:  0 end, 0 nop, 0 basic security, 0 loose source route
         0 timestamp, 0 extended security, 0 record route
         0 stream ID, 0 strict source route, 0 alert, 0 cipso, 0 ump
         0 other
  Frags: 0 reassembled, 0 timeouts, 0 couldn't reassemble
```

```
       0 fragmented, 0 couldn't fragment
Bcast: 1952255 received, 4 sent
Mcast: 0 received, 0 sent
Sent: 86915 generated, 0 forwarded
Drop:  18 encapsulation failed, 0 unresolved, 0 no adjacency
       0 no route, 0 unicast RPF, 0 forced drop
```

This command has more output, our only concern is extracting the number of no routes and no adjacencies in the first section, underlined, above.

Create ACLI

Create a *show* ACLI for the `show ip traffic` command. This ACLI executes the command and it will extract the appropriate data using RegEx.

1 Create a new ACLI (or modify an existing)

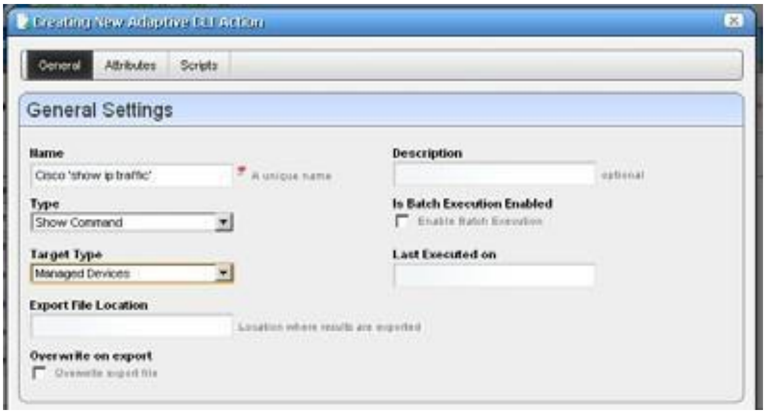

Name: Cisco 'show ip traffic'

Type: Show Command

Target Type: Managed Devices

Export File Location:

Overwrite on export: □

Description:

Is Batch Execution Enabled: □

Last Executed on:

Create an attribute in ACLI for each attribute that you want to monitor. The Date Type must be Integer.

2 Here, there is an attribute for `no_adjacency`.

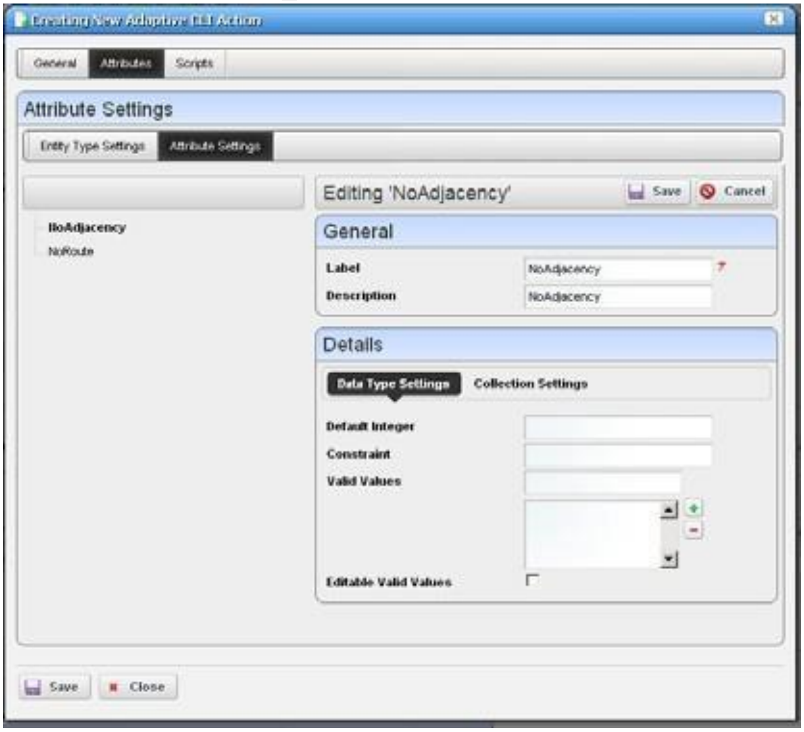

Creating New Adaptive CLI Action

General | Attributes | Scripts

Attribute Settings

Entity Type Settings | Attribute Settings

NoAdjacency
NoRoute

Editing 'NoAdjacency'

General

Label: NoAdjacency

Description: NoAdjacency

Details

Data Type Settings | Collection Settings

Default Integer:

Constraint:

Valid Values:

Editable Valid Values: □

3 Create an attribute for `no_route,` too.

4 Create the script which contains the command you want to run. Notice the the attributes appear listed on the right panel. Do not refer to the attributes in the script as you would in a configuration script. The next step contains attribute references.

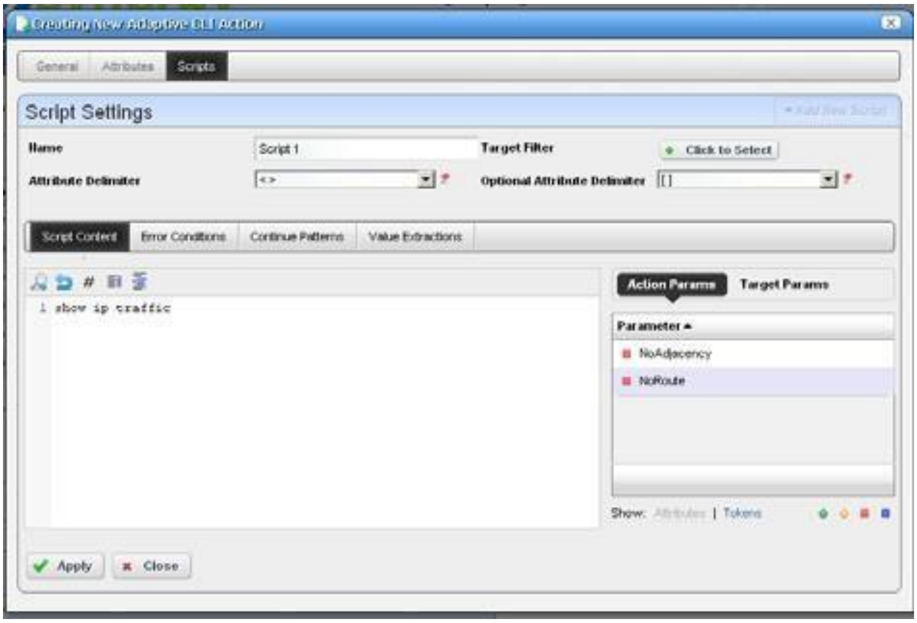

Creating New Adaptive CLI Action

General | Attributes | Scripts

Script Settings

Name: Script 1

Target Filter: Click to Select

Attribute Delimiter: <>

Optional Attribute Delimiter: []

Script Content | Error Conditions | Continue Patterns | Value Extractions

```
1 show ip traffic
```

Action Params | Target Params

Parameter ▲
NoAdjacency
NoRoute

Show: Attributes | Tokens

5 In the Value Extractions Tab Click *Add new attribute extraction* and then pick the attribute. The Parse Expression is the RegEx that extracts the correct data. For `no_route` use this RegEx: `.*(\d+) no route,.*`

6. *Add* the next attribute. For `no_adjacency` use the RegEx: `.*unresolved, (\d+) no adjacency.*`

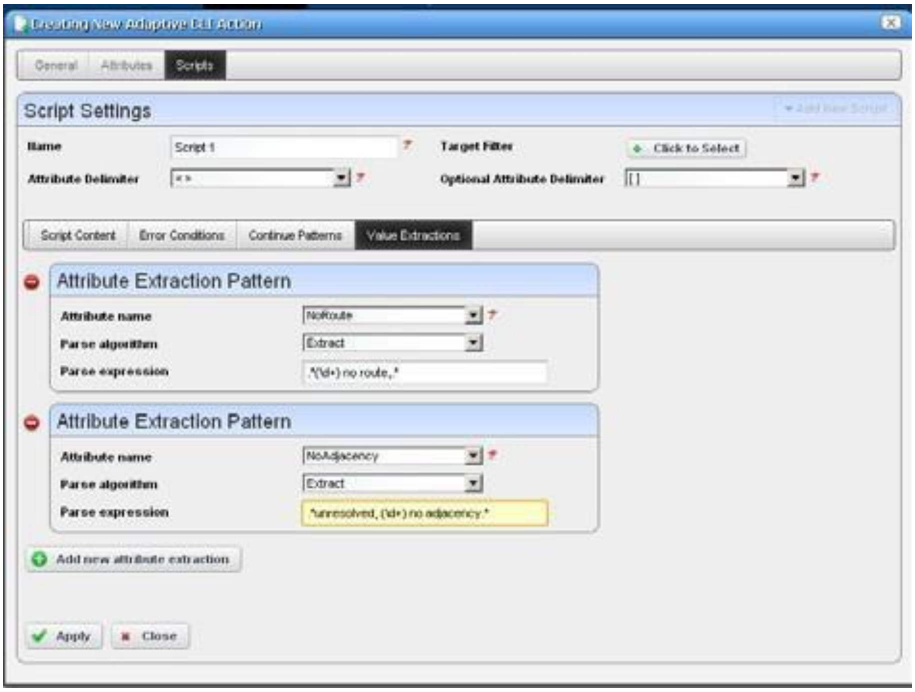

Name: Script 1
Target Filter: Click to Select
Attribute Delimiter: <>
Optional Attribute Delimiter: []

Attribute Extraction Pattern
Attribute name: NoRoute
Parse algorithm: Extract
Parse expression: .*(\d+) no route,.*

Attribute Extraction Pattern
Attribute name: NoAdjacency
Parse algorithm: Extract
Parse expression: .*unresolved, (\d+) no adjacency.*

7. Apply and Save the ACLI.
8. Select the ACLI and execute it, selecting the devices you want to monitor.

   The audit panel catalogs the progress of the job, and the Execution History snap panel in the expanded Actionsportletdisplaystheexecution, listing multipleexecutionsbytimeand date. Right-click an execution listed and select *Results* to see the results.

Create the Monitor

Follow these steps to monitor the Adaptive CLI created in the previous section.

1. Create a new monitor, select *Adaptive CLI* as thetype.
2. Name the monitor, set the polling interval.
3. Select device(s) to monitor and then select the ACLI you just created. Notice that the attributes appear in the *input parameters* tab.
4. Under the Monitor Attributes tab, use the defaults.
5. Set any thresholds you like in the *Thresholds* node. This example monitors normal functionality, so it includes no thresholds.
6. When you save the monitor, it begins working and executes the ACLI every polling cycle, extracting the data.

How To:

Create a Monitor for an External Script

The following steps describe creating a monitor for an external command configured as an Adaptive CLI (ACLI).

Create the Adaptive CLI

1 Right click in the Actions portlet, and create a new *External Command* ACLI

2 Make a new attribute schema with attribute: Status (integer)

3 In Scripts, enter the follo wing as Script Content:

```
perl "C:\[installation path] \owareapps\performance\scripts\
 http_test.pl"[_EquipmentManager_IP_Address]
```

The path that precedes the owareapps directory may differ if you have installed Cruz elsewhere, and by default on Linux.

Several Perlscripts appear in this performance\scripts directorybydefault. Youcan tryothers in addition to the http_test.pl script.

4 In the Value Extraction panel enter the following:

```
^\{(\d+)\}.*
```

5 Click Apply

6 Click Save

7 Right click and *Execute* the ACLI to test it.

8 Look in Job Viewer for the results.

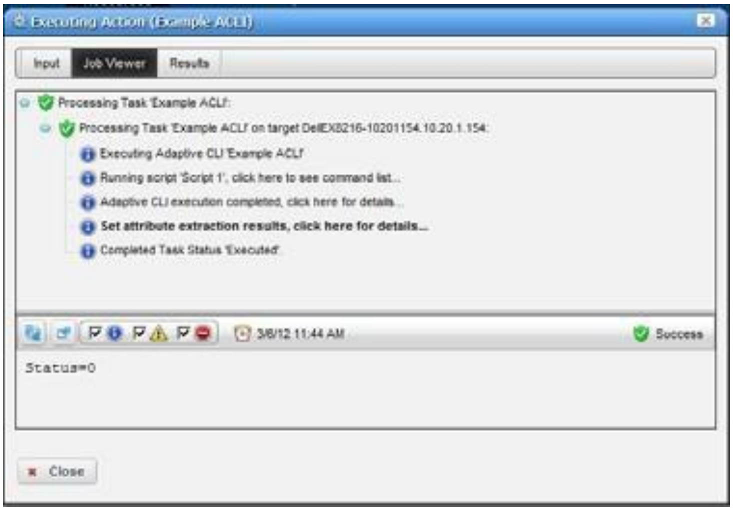

Processing Task 'Example ACLI':
Processing Task 'Example ACLI' on target DellEX8216-10201154.10.20.1.154:
Executing Adaptive CLI 'Example ACLI'
Running script 'Script 1', click here to see command list...
Adaptive CLI execution completed, click here for details...
Set attribute extraction results, click here for details...
Completed Task Status 'Executed'

```
Status=0
```

Click *Set attribute extraction results, click here* to see the results appear in the bottom panel. Notice also that you must check informational messages for all these to appear, and that several additional sets of messages besides the extraction results appear.

Create a Monitor for the External Script Adaptive ACLI

Now that you have verified the script is working, you can create a monitor to see how this attribute is doing.

1 In the Monitors portlet, create a new ACLI Monitor

2 Uncheck *Update Network Status* (recommended since the ICMP monitor is already doing this)

3 In *Monitor Options* select your example monitor configured previously.

4 Confirm that *Monitor Attributes* displays the Status attribute configured previously.

5. In the *Conditions* tab of the Monitor Editor, create “Status Up” condition, with the severity of *Informational*, and check *Alert*.
6. Create a criterion which is Status = 0.
7. Save this condition
8. Create a new Condition called “Status Down”
9. The criterion is Status = 1
10. Apply and Save
11. Save your monitor

NOTE:

You may want to test your monitor, in which case, you may want to change the interval to 30 seconds.

12. Right-click to select View Monitor Data, and you can see the results of your efforts.

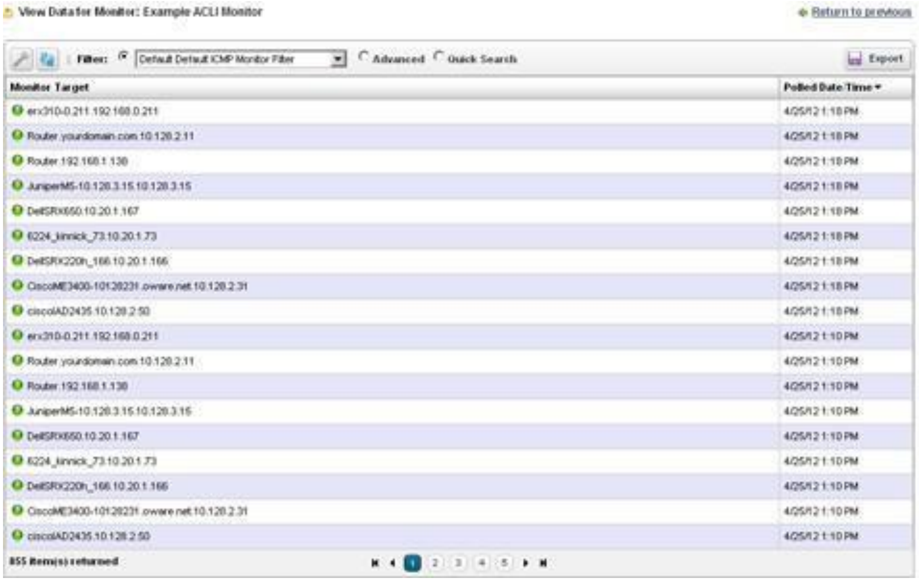

View Data for Monitor: Example ACLI Monitor

Filter: ● Default Default ICMP Monitor Filter ○ Advanced ○ Quick Search

Monitor Target	Polled Date Time
erx310-0.211.192.168.0.211	4/25/12 1:18 PM
Router.yourdomain.com.10.128.2.11	4/25/12 1:18 PM
Router.192.168.1.130	4/25/12 1:18 PM
JuniperM5-10.128.3.15.10.128.3.15	4/25/12 1:18 PM
DellSRX650.10.20.1.167	4/25/12 1:18 PM
6224_kinnick_73.10.20.1.73	4/25/12 1:18 PM
DellSRX220h_166.10.20.1.166	4/25/12 1:18 PM
CiscoME3400-10128231.oware.net.10.128.2.31	4/25/12 1:18 PM
ciscoIAD2435.10.128.2.50	4/25/12 1:18 PM
erx310-0.211.192.168.0.211	4/25/12 1:10 PM
Router.yourdomain.com.10.128.2.11	4/25/12 1:10 PM
Router.192.168.1.130	4/25/12 1:10 PM
JuniperM5-10.128.3.15.10.128.3.15	4/25/12 1:10 PM
DellSRX650.10.20.1.167	4/25/12 1:10 PM
6224_kinnick_73.10.20.1.73	4/25/12 1:10 PM
DellSRX220h_166.10.20.1.166	4/25/12 1:10 PM
CiscoME3400-10128231.oware.net.10.128.2.31	4/25/12 1:10 PM
ciscoIAD2435.10.128.2.50	4/25/12 1:10 PM

855 item(s) returned

How To:

Create a Monitor Report

Youcan create reports based on your monitors. The following example creates a report based on How to: Create an SNMP Interface Monitor above.

1. Create a new Report Template by right clicking the Report Templates portlet, selecting *New > Table Template* .
2. Name the report (here: Test SNMP Interface Report).
3. Select a source in the Source tab. Here: *Active Monitoring > SNMP Interfaces* .
4. Notice that the *Select your inventory columns* panel displays the attributes available based on your monitor selection.

5 Select *Available* columns and click the right arrow to move them to *Selected.* In this case we select SNMP Interfaces: Monitor Ta rget, Polled Date/Time, ifInErrors.

6 Arrange the columns and fonts as you like in the *Layout* tab.

7 *Save* the template.

8 Right-click and select *New* in the Reports portlet.

9 Enter a *Name* and *Title* for the report.

10 Notice that since this is the first report created since you made the Test SNMP Interface Report template, that it is the *Report Template* already selected.

11 Since the monitor already filters devices, we add no filter in the Report, although you could add one to further filter the monitored devices.

12 Test SNMP Interface Report should appearin the Reports portlet.

13 Right-click and select *Execute* (noticing that you can also schedule such reports, even repeatedly).

14 Click the magnifying glass to the right of the Report Completed message in My Alerts to see the report.

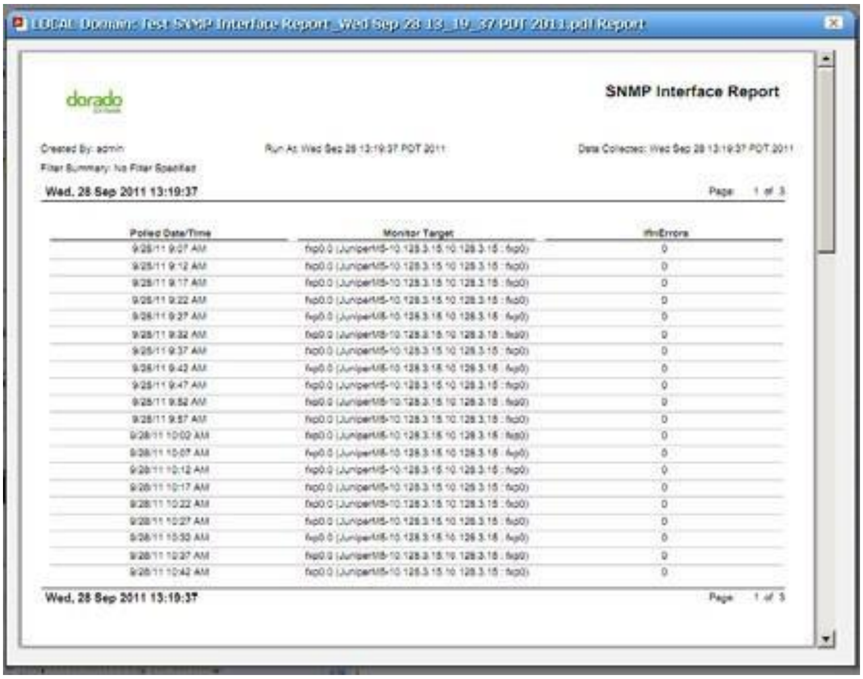

LOCAL Domain: Test SNMP Interface Report _Wed Sep 28 13_19_37 PDT 2011.pdf Report

SNMP Interface Report

Created By: admin
Run At: Wed Sep 28 13:19:37 PDT 2011
Data Collected: Wed Sep 28 13:19:37 PDT 2011
Filter Summary: No Filter Specified

Wed, 28 Sep 2011 13:19:37 Page 1 of 3

Polled Date/Time	Monitor Target	ifInErrors
9/28/11 9:07 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 9:12 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 9:17 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 9:22 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 9:27 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 9:32 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 9:37 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 9:42 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 9:47 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 9:52 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 9:57 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 10:02 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 10:07 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 10:12 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 10:17 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 10:22 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 10:27 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 10:32 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 10:37 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0
9/28/11 10:42 AM	fxp0.0 (JuniperM5-10.128.3.15 10.128.3.15 : fxp0)	0

Wed, 28 Sep 2011 13:19:37 Page 1 of 3

15 Hover your cursor over the lower right corner of the report to see a set of icons that let you expand, zoom out and in, save, or print the report.

How To:

Createa MonitorTargetReport

Similar to Monitor Reports, you can create reports based on the targets associated with your monitors. The following example creates a report based on How to: Create an SNMP Interface Monitor above.

> Create a new Report Template by right clicking the Report Templates portlet, selecting New > Table Template.

> Name the report (here: Test SNMP Interface Targets Report).

> Select a source in the Source tab. Here: Active Monitoring > SNMP InterfacesTargets (there should be two entity types for each monitor: one for the actual monitor data, and one for the targets).

> Notice that the Select your inventorycolumns panel displays the attributes associated with the monitor targets. These attributes are the same for each monitor target entity type. This is unlike the monitor data entity types, where the list of available attributes is different for each monitor.

> Select Available columns and click the right arrow to move them to Selected. In this case we select Monitor Target, Equipment, Availability, Last Polled, Enabled, Device Management State.

> Arrange thecolumns and fonts as you like in the Layout tab.

> Save the template.

> Right-click and select New in the Reports portlet.

> Enter a Name and Title for the report.

> Notice that since this is the first report created since you made the Test SNMP Interface Targets Report template, that it is the Report Template already selected.

> Since the monitor already filters devices, we add no filter in the Report, although you could add one to further filter the target devices and/or subcomponents.

> Test SNMP Interface Target Report should appear in the Reports portlet.

> Right-click and select Execute (noticing that you can also schedule such reports, even repeatedly).

> Click the magnifying glass to the right of the Report Completed message in My Alerts to see the report.

Monitor Options Type-Specific Panels

The following describes the panels associated with the following Monitor Options types.

- Adaptive CLI
- Cisco IPSLA
- Cisco Metro Ethernet SLA Monitors
- Cisco QoS Monitors
- ICMP
- Juniper CoS
- Juniper RPM
- Key Metrics
- ProScan
- SNMP
- SNMP Interfaces
- SNMP Table Monitor
- VRF

Adaptive CLI

For this monitor, see How to Create an Adaptive CLI Monitor on page 381.

Select Monitor Entities in the top panel, and an Adaptive CLI to monitor at the top of the bottom panel. The *Input Parameters* for the Adaptive CLI appear in that tab, and you can edit the *Monitor Attributes* in that tab.

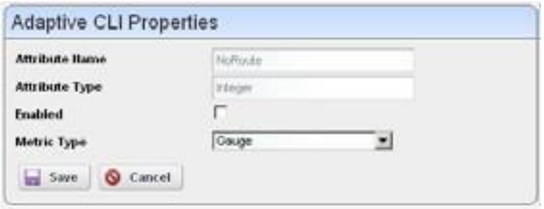

Adaptive CLI Properties

Attribute Name	NoRoute
Attribute Type	Integer
Enabled	☐
Metric Type	Gauge

The *Name* and *Type*, and whether the attribute is *Enabled* appear in this editor. You can also select whether the attribute is a Counter, Gauge or Boolean. For Counter types, the monitor computes change from previous readings, and for Gauges it does not. Boolean attributes are either true or false

Cisco IPSLA

This screen configures options for Cisco IPSLA monitoring.

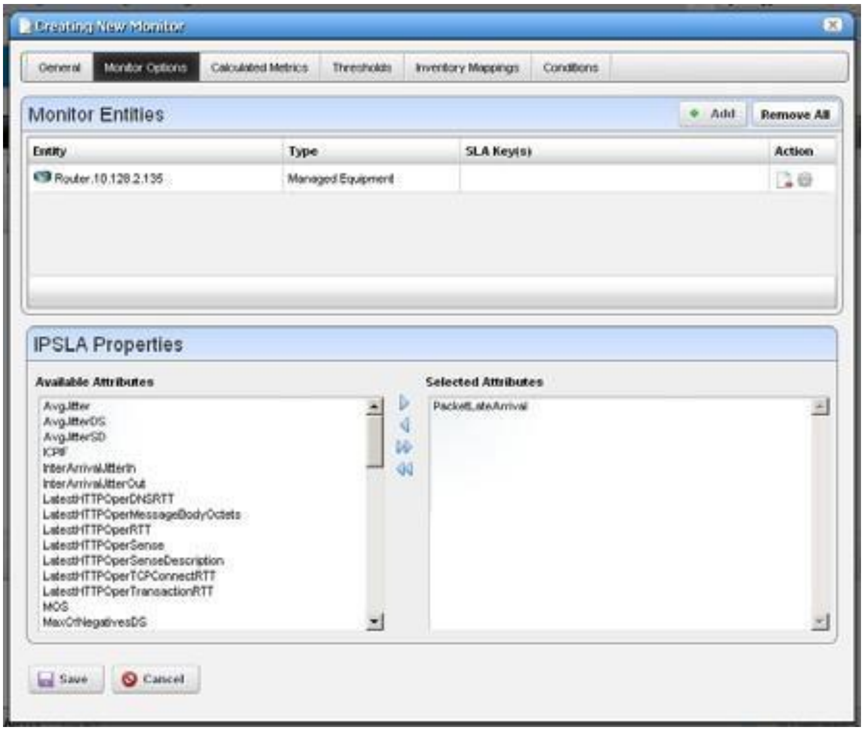

Monitor Entities

Entity	Type	SLA Key(s)	Action
Router.10.128.2.135	Managed Equipment		

IPSLA Properties

Available Attributes

- AvgJitter
- AvgJitterDS
- AvgJitterSD
- ICPIF
- InterArrivalJitterIn
- InterArrivalJitterOut
- LatestHTTPOperDNSRTT
- LatestHTTPOperMessageBodyOctets
- LatestHTTPOperRTT
- LatestHTTPOperSense
- LatestHTTPOperSenseDescription
- LatestHTTPOperTCPConnectRTT
- LatestHTTPOperTransactionRTT
- MOS
- MaxOfNegativesDS

Selected Attributes

- PacketLateArrival

△ CAUTION:

Target equipment for this monitor must be an IPSLA transmitter, not an IPSLA responder.

Click to select from *Available Attributes* and use the arrows to move such attributes to *Selected Attributes* that you want to monitor.

NOTE:

This Monitor provides end-to-end serviceverification. Alarms appear in the Service Details Panel and service topology. You mustconfigure the monitor to emitavailability events for this tooccur.

IPSLA OIDS

The following are the object IDs for IPSLA, all found in CISCO-RTTMON-MIB

MonitorAttribute Name	Mib Attribute name	OID
NumOfPositvesDS	rttMonEchoAdminNumPackets	1.3.6.1.4.1.9.9.42.1.2.2.1.18
NumOfRTT	rttMonLatestJitterOperNumOfRTT	1.3.6.1.4.1.9.9.42.1.5.2.1.1
RTTSum	rttMonLatestJitterOperRTTSum	1.3.6.1.4.1.9.9.42.1.5.2.1.2
RTTSum2	rttMonLatestJitterOperRTTSum2	1.3.6.1.4.1.9.9.42.1.5.2.1.3
MinRTT	rttMonLatestJitterOperRTTMin	1.3.6.1.4.1.9.9.42.1.5.2.1.4
MaxRTT	rttMonLatestJitterOperRTTMax	1.3.6.1.4.1.9.9.42.1.5.2.1.5
MinOfPositivesSD	rttMonLatestJitterOperMinOfPositivesSD	1.3.6.1.4.1.9.9.42.1.5.2.1.6
MaxOfPositvesSD	rttMonLatestJitterOperMaxOfPositivesSD	1.3.6.1.4.1.9.9.42.1.5.2.1.7
NumOfPositivesSD	rttMonLatestJitterOperNumOfPositivesSD	1.3.6.1.4.1.9.9.42.1.5.2.1.8
NumOfPositivesDS	rttMonLatestJitterOperSumOfPositivesSD	1.3.6.1.4.1.9.9.42.1.5.2.1.9
Sum2PositivesSD	rttMonLatestJitterOperSum2PositivesSD	1.3.6.1.4.1.9.9.42.1.5.2.1.10
MinOfNegativesSD	rttMonLatestJitterOperMinOfNegativesSD	1.3.6.1.4.1.9.9.42.1.5.2.1.11
MaxOfNegativesSD	rttMonLatestJitterOperMaxOfNegativesSD	1.3.6.1.4.1.9.9.42.1.5.2.1.12
NumOfNegativesSD	rttMonLatestJitterOperNumOfNegativesSD	1.3.6.1.4.1.9.9.42.1.5.2.1.13
SumOfNegativesSD	rttMonLatestJitterOperSumOfNegativesSD	1.3.6.1.4.1.9.9.42.1.5.2.1.14
Sum2NegativesSD	rttMonLatestJitterOperSum2NegativesSD	1.3.6.1.4.1.9.9.42.1.5.2.1.15
MinOfPositivesDS	rttMonLatestJitterOperMinOfPositivesDS	1.3.6.1.4.1.9.9.42.1.5.2.1.16
MaxOfPositivesDS	rttMonLatestJitterOperMaxOfPositivesDS	1.3.6.1.4.1.9.9.42.1.5.2.1.17
NumOfPositivesDS	rttMonLatestJitterOperNumOfPositivesDS	1.3.6.1.4.1.9.9.42.1.5.2.1.18
SumOfPositivesDS	rttMonLatestJitterOperSumOfPositivesDS	1.3.6.1.4.1.9.9.42.1.5.2.1.19
Sum2PositivesDS	rttMonLatestJitterOperSum2PositivesDS	1.3.6.1.4.1.9.9.42.1.5.2.1.20

Monitor Attribute Name	Mib Attribute name	OID
MinOfNegativesDS	rttMonLatestJitterOperMinOfNegativesDS	1.3.6.1.4.1.9.9.42.1.5.2.1.21
MaxOfNegativesDS	rttMonLatestJitterOperMaxOfNegativesDS	1.3.6.1.4.1.9.9.42.1.5.2.1.22
NumOfNegativesDS	rttMonLatestJitterOperNumOfNegativesDS	1.3.6.1.4.1.9.9.42.1.5.2.1.23
SumOfNegativesDS	rttMonLatestJitterOperSumOfNegativesDS	1.3.6.1.4.1.9.9.42.1.5.2.1.24
Sum2NegativesDS	rttMonLatestJitterOperSum2NegativesDS	1.3.6.1.4.1.9.9.42.1.5.2.1.25
PacketLossSD	rttMonLatestJitterOperPacketLossSD	.3.6.1.4.1.9.9.42.1.5.2.1.26
PacketLossDS	rttMonLatestJitterOperPacketLossDS	.3.6.1.4.1.9.9.42.1.5.2.1.27
PacketOutOfSequence	rttMonLatestJitterOperPacketOutOfSequence	1.3.6.1.4.1.9.9.42.1.5.2.1.28
PacketMIA	rttMonLatestJitterOperPacketMIA	.3.6.1.4.1.9.9.42.1.5.2.1.29
PacketLateArrival	rttMonLatestJitterOperPacketLateArriv al	1.3.6.1.4.1.9.9.42.1.5.2.1.30
OWSumSD	rttMonLatestJitterOperOWSumSD	1.3.6.1.4.1.9.9.42.1.5.2.1.33
OWSum2SD	rttMonLatestJitterOperOWSum2SD	1.3.6.1.4.1.9.9.42.1.5.2.1.34
OWMinSD	rttMonLatestJitterOperOWMinSD	1.3.6.1.4.1.9.9.42.1.5.2.1.35
OWMaxSD	rttMonLatestJitterOperOWMaxSD	1.3.6.1.4.1.9.9.42.1.5.2.1.36
OWSumDS	rttMonLatestJitterOperOWSumDS	1.3.6.1.4.1.9.9.42.1.5.2.1.37
OWSum2DS	rttMonLatestJitterOperOWSum2DS	1.3.6.1.4.1.9.9.42.1.5.2.1.38
OWMinDS	rttMonLatestJitterOperOWMinDS	1.3.6.1.4.1.9.9.42.1.5.2.1.39
OWMaxDS	rttMonLatestJitterOperOWMaxDS	1.3.6.1.4.1.9.9.42.1.5.2.1.40
NumOfOW	rttMonLatestJitterOperNumOfOW	1.3.6.1.4.1.9.9.42.1.5.2.1.41
MOS	rttMonLatestJitterOperMOS	1.3.6.1.4.1.9.9.42.1.5.2.1.42
ICPIF	rttMonLatestJitterOperICPIF	.3.6.1.4.1.9.9.42.1.5.2.1.43
InterArrivalJitterOut	rttMonLatestJitterOperIAJOut	1.3.6.1.4.1.9.9.42.1.5.2.1.44
InterArrivalJitterIn	rttMonLatestJitterO perIAJIn	1.3.6.1.4.1.9.9.42.1.5.2.1.45

Monitor Attribute Name	Mib Attribute name	OID
AvgJitter	rttMonLatestJitterOperAvgJitter	1.3.6.1.4.1.9.9.42.1.5.2.1.46
AvgJitterSD	rttMonLatestJitterOperAvgSDJ	1.3.6.1.4.1.9.9.42.1.5.2.1.47
AvgJitterDS	rttMonLatestJitterOperAvgDSJ	1.3.6.1.4.1.9.9.42.1.5.2.1.48
OWAvgSD	rttMonLatestJitterOperOWAvgSD	1.3.6.1.4.1.9.9.42.1.5.2.1.49
OWAvgDS	rttMonLatestJitterOperOWAvgDS	1.3.6.1.4.1.9.9.42.1.5.2.1.50
LatestHTTPOperRT	rttMonLatestHTTPOp erRTT	1.3.6.1.4.1.9.9.42.1.5.1.1.1
LatestHTTPOperDNSRTT	rttMonLatestHTTPOperDNSRTT	1.3.6.1.4.1.9.9.42.1.5.1.1.2
LatestHTTPOperTCPConnectRTT	rttMonLatestHTTPOperTCPConnectRTT	1.3.6.1.4.1.9.9.42.1.5.1.1.3
LatestHTTPOperTransactionRTT	rttMonLatestHTTPOperTransactionRTT	1.3.6.1.4.1.9.9.42.1.5.1.1.4
LatestHTTPOperMessageBodyOctets	rttMonLatestHTTPOperMessageBodyOctets	1.3.6.1.4.1.9.9.42.1.5.1.1.5
LatestHTTPOperSense	rttMonLatestHTTPOp erSense	1.3.6.1.4.1.9.9.42.1.5.1.1.6
LatestHTTPErrorSenseDescription	rttMonLatestHTTPErrorSenseDescription	1.3.6.1.4.1.9.9.42.1.5.1.1.7
LatestRttOperCompletionTime r	ttMonLatestRttOperCompletionTime	1.3.6.1.4.1.9.9.42.1.2.10.1.1

Cisco Metro Ethernet SLA Monitors

This monitor Cisco's CISCO-IPSLA-ETHERNET-MIB. It collects performance metrics against individual RTT probe operations for the target device. This monitor works like the existing Cisco IPSLA monitor for the CISCO-RTTMON-MIB. It indexes statistics using the probe index as defined in the rttMonCtrlAdminTable in the CISCO-RTTMON-MIB. The Ethernet jitter probes have a rttMonCtrlAdminRttType of ethernetJitter.

Cruz does not provide configuring of Ethernet jitter SLA operations. Information about configuring CFM can be found in the following documents found on the Cisco web site: *Configuring Ethernet CFM and OAM* and *Configuring IP SLAs for Metro-Ethernet*.

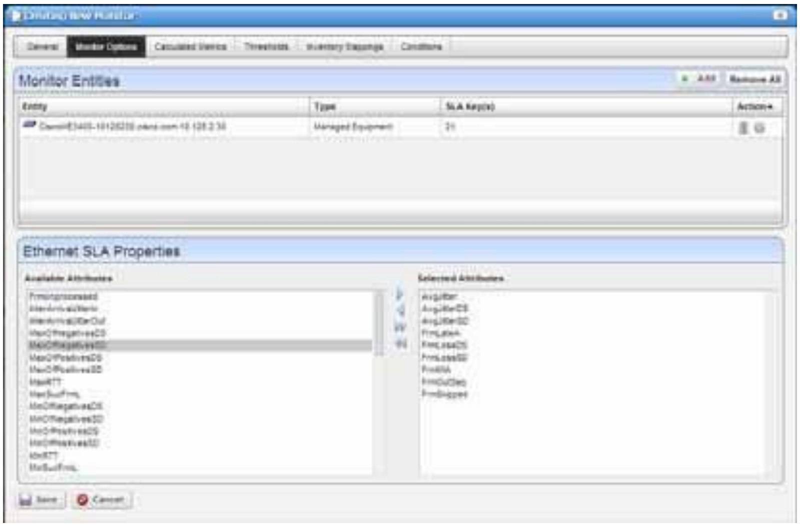

To create a Cisco Metro Ethernet monitor, select New Monitor > Cisco Ethernet SLA from the right-click menu in the Resource Monitors portlet. The Monitor Options screen lets you configure monitored attributes and devices. As with the IPSLA monitor, you can select either equipment groups or individual managed equipment objects. If you select a managed equipment object you can click on the gear icon in the Action column to see a list of the Ethernet Jitter SLA probe IDs configured for that device and select specific probes which then appear in the SLA Key(s) column. Without such a selection, the monitor tracks all Ethernet Jitter SLA probes for the device.

The Ethernet SLA Properties panel allows selection of the specific attributes you wish to collect data on. The attributes correspond directly to the oids in the ipslaEtherJitterAggStatsTable in the ipslaEthernetStats section of the CISCO-IPSLA-ETHERNET-MIB.

Cisco QoS Monitors

This monitors values for Cisco QoS from the Cisco Class-Based QOS MIB. The following screenshots come from the Class Map monitor, but Cruz's Cisco QoS monitoring capabilities include more than just this monitor. See Additional QoS Monitors on page 395.

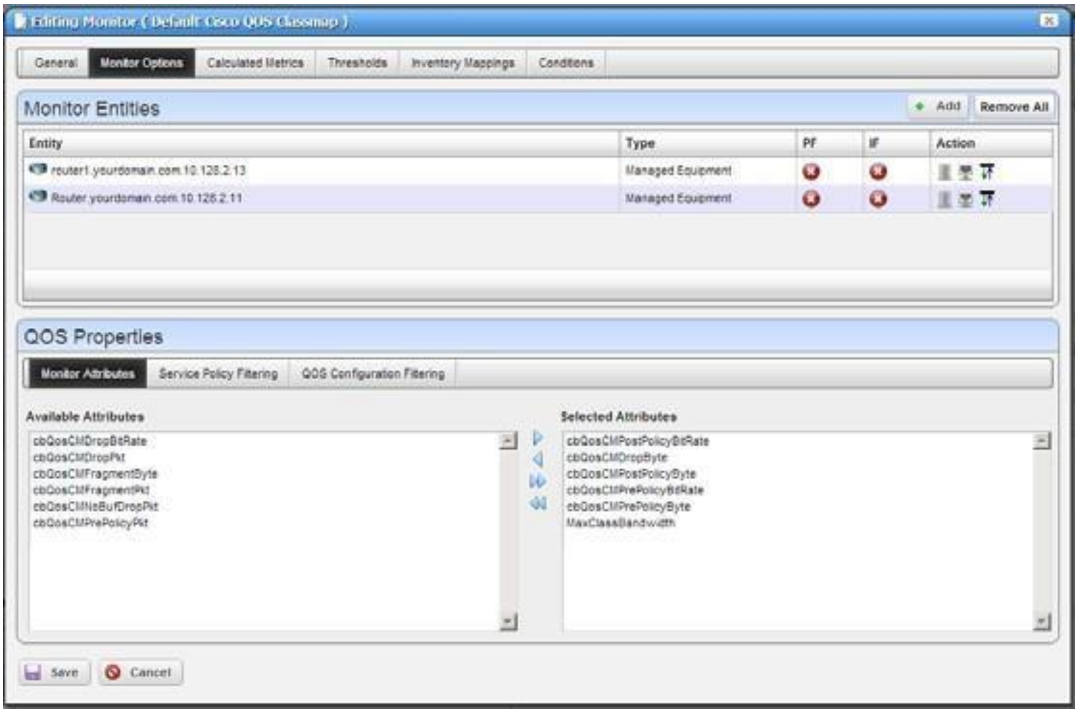

Entity	Type	PF	IF	Action
router1.yourdomain.com.10.128.2.13	Managed Equipment	✖	✖	
Router.yourdomain.com.10.128.2.11	Managed Equipment	✖	✖	

Available Attributes

cbQosCMDropBitRate
cbQosCMDropPkt
cbQosCMFragmentByte
cbQosCMFragmentPkt
cbQosCMNoBufDropPkt
cbQosCMPrePolicyPkt

Selected Attributes

cbQosCMPostPolicyBitRate
cbQosCMDropByte
cbQosCMPostPolicyByte
cbQosCMPrePolicyBitRate
cbQosCMPrePolicyByte
MaxClassBandwidth

Service Policies

A Service Policy is a policy map attached to a logical interface. Because a policy map can also be a part of the hierarchical structure (inside a classmap), Cruz considers only a policy map directly attached to a logical interface as a service policy.

Class Map—A user-defined traffic class that contains one or many match statements that classify packets into different categories.

Match Statement — Specifies specific match criteria to identify packets for classification purposes. Match statements exist within a class map.

Policy Map—A user-defined policy that associates QoS actions to the user-defined traffic class - ClassMap.

Qos Actions—These include: Queueing, Random Detect (WRED), Traffic Shaping, Police, Set (Packet Marking), Compression (IP header), Account(C3pl).

See Additional QoS Monitors on page 395 for attributes you can monitor related to

these.

Monitor Entitles

Select the equipment to monitor in this screen. Notice the *PF* and *IF* columns that indicate whether a port filter or interface filter is active to further limit what parts of the selected device is monitored. Delete the device or configure port and interface filters with the icons to the right of listed equipment.

QOS Properties - Monitored Attributes

This panel lets you select attributes monitored with right/left selection arrows. Move the desired attributes from the *Available* to the *Selected* side of this panel. Notice that, by default, you can monitor two additional calculated attributes. You can also edit the monitor to create additional calculated attributes, which would also appear here.

NOTE:

If you change the bandwidth (speed) on a port/interface, you must resync the device to update the port speed in Cruz. Youmust refresh targets on any QOS monitor after resync for theapplication to reflect the correct MaxClassBandwidth value. See also Bandwidth Calculation on page 410.

Additional QoS Monitors

The sections below list the monitored attributes for possible Monitor types that include the following:

- Qos Class Map Monitor
- Qos Match Statement Monitor
- Qos Police Monitor
- Qos Queuing Monitor
- Qos Traffic Shaping Monitor
- Qos RED Monitor
- Qos IPHC Monitor
- Qos Packet Marking Monitor
- QoS Police Monitor
- Qos Estimate Bandwidth Monitor
- QoS C3pl Account Monitor

Qos Class Map Monitor

Qos Class Map Monitor collects metrics from the cbQosCMStatsTable. This table specifies ClassMap related Statistical information.

Available Metrics:

```
DropBitRate
PrePolicyPkt
PrePolicyByte
PrePolicyBitRate
PostPolicyByte
PostPolicyBitRate
DropPkt
DropByte
NoBufDropPkt
FragmentPkt
FragmentByte
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

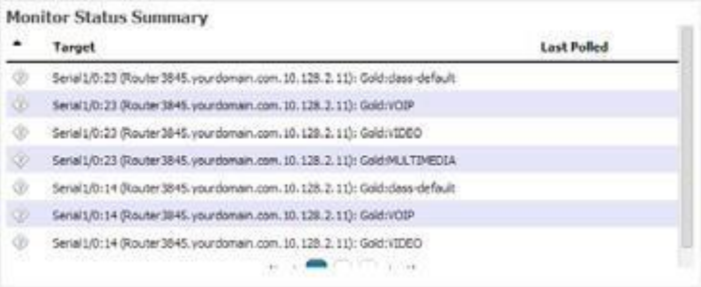

Monitor Status Summary

▲	Target	Last Polled
	Serial1/0:23 (Router3845.yourdomain.com.10.128.2.11): Gold:class-default	
	Serial1/0:23 (Router3845.yourdomain.com.10.128.2.11): Gold:VOIP	
	Serial1/0:23 (Router3845.yourdomain.com.10.128.2.11): Gold:VIDEO	
	Serial1/0:23 (Router3845.yourdomain.com.10.128.2.11): Gold:MULTIMEDIA	
	Serial1/0:14 (Router3845.yourdomain.com.10.128.2.11): Gold:class-default	
	Serial1/0:14 (Router3845.yourdomain.com.10.128.2.11): Gold:VOIP	
	Serial1/0:14 (Router3845.yourdomain.com.10.128.2.11): Gold:VIDEO	

```
<PolicyMapName> : <Name>
```

Qos Match Statement Monitor

Qos Match Statement Monitor collects metrics from the cbQosMatchStmtStatsTable. This table specifies Match Statement related statistical information.

Available Metrics

```
PrePolicyPkt
PrePolicyByte
PrePolicyBitRate
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

```
<PolicyMapName> : <Name> : <StmtName>
```

Qos Police Monitor

Qos Police Monitor collects metrics from the cb QosPoliceStatsTable. This table specifies Police Action related statistical information.

Available Metrics

```
ConformedPkt
ConformedByte
ConformedBitRate
ExceededPkt
ExceededByte
ExceededBitRate
ViolatedPkt
ViolatedByte
ViolatedBitRate
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

```
<PolicyMapName> : <CMName> : <CfgRate>
```

Qos Queuing Monitor

Qos Queuing Monitor collects metrics fromthe cbQosQueueingStatsTable. This table specifies Queueing Action-related statistical information.

Available Metrics

```
CurrentQDepth
MaxQDepth
DiscardByte
DiscardPkt
```

QoS Configuration Filtering Attributes

```
CfgBandwidth - specified in kbps or percentage
CfgBandwidthUnits - enumeration
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

```
<PolicyMapName> : <CMName> : <CfgBandwidth>
```

Qos Traffic Shaping Monitor

Qos Traffic Shaping Monitor collects metrics from the cbQosTSStatsTable. This table specifies traffic-shaping Action rela tedstatistical information.

Available Metrics

```
DelayedByte
DelayedPkt
DropByte
DropPkt
Active
QSize
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

```
<PolicyMapName> : <CMName> :<TSCfgRate>
```

Qos RED Monitor

Qos RED Monitor collects metrics from the cbQo sREDClassStatsTable. This table specifies per Precedence WRED (wait random early detection) Action-related statistical information.

Available Metrics

```
RandomDropPkt
RandomDropByte
TailDropPkt
TailDropByte
TransmitPkt
TransmitByte
ECNMarkPkt
ECNMarkByte
```

```
MeanQSizeUnits
MeanQSize
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

```
<PolicyMapName> : <CMName> : < CfgDscpPrec>
```

Qos IPHC Monitor

Qos IPHC Monitor collects metricsfrom the cbQo sIPHCStatsTable. Thistablespecifies IP Header Compression statistical information.

Available Metrics

```
RtpSentPkt
RtpCmprsOutPkt
RtpSavedByte
RtpSentByte
RtpSentByteRate
TcpSentPkt
TcpCmprsOutPkt
TcpSavedByte
TcpSentByte
TcpSentByteRate
RtpFullHdrSentPkt
TcpFullHdrSentPkt
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

```
<cbQosPolicyMapName> : <cbQosCMName> :<CfgOption>
```

Qos Packet Marking Monitor

Qos Packet Marking Monitor collects metrics from the cbQosSetStatsTable. This table specifies packet marking statistical information.

Available Metrics

```
DscpPkt
PrecedencePkt
QosGroupPkt
FrDePkt
AtmClpPkt
L2CosPkt
MplsExpImpositionPkt
DiscardClassPkt
MplsExpTopMostPkt
SrpPriorityPkt
FrFecnBecnPkt
DscpTunnelPkt
```

```
PrecedenceTunnelPkt
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

```
<cbQosPolicyMapName> : <cbQosCMName> :<CfgFeature>
```

QoS Police Monitor

Qos Police Monitor collects metrics from the cb QosPoliceColorStatsTable. This table specifies Police Action-related statistical inform ation for two rate coloraware marker.

Available Metrics

```
ConformedBitRate
ConfirmedByte
ConformedPkt
ExceededBitRate
ExceededByte
ExceededPkt
ViolatedBitRate
ViolatedByte
ViolatedPkt
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

```
<cbQosPolicyMapName> : <cbQosCMName> :<cbQosPoliceCfgRate>
```

Qos Estimate Bandwidth Monitor

Qos Estimate Bandwidth Monitor willcollect metrics fromthe cbQosEBStatsTable. Thistable specifies Estimate Bandwidth related statistical information.

Available Metrics

```
StatsCorvilEBValue
StatsCorvilEBStatus
StatsCorvilCTD
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

```
<cbQosPolicyMapName> : <cbQosCMName> :<???>
```

QoS C3pl Account Monitor

Qos C3pl Account MonitorcollectsmetricsfromthecbQosC3plAccountStatsTable. Thistable specifies C3pl Account Action-rel ated statistics information.

Available Metrics

```
cbQosC3plAccountDropPkt
cbQosC3plAccountDropByte
cbQosC3plAccountTailDropPkt
cbQosC3plAccountTailDropByte
```

Target Summary Pattern

This pattern appears in the expanded Resource Monitors portlet when you select the monitor in the top panel

```
<cbQosPolicyMapName> : <cbQosCMName> : <cbQosC3plAccountFeatureType>
```

ICMP

The ICMP Monitor Options panel contains the following properties:

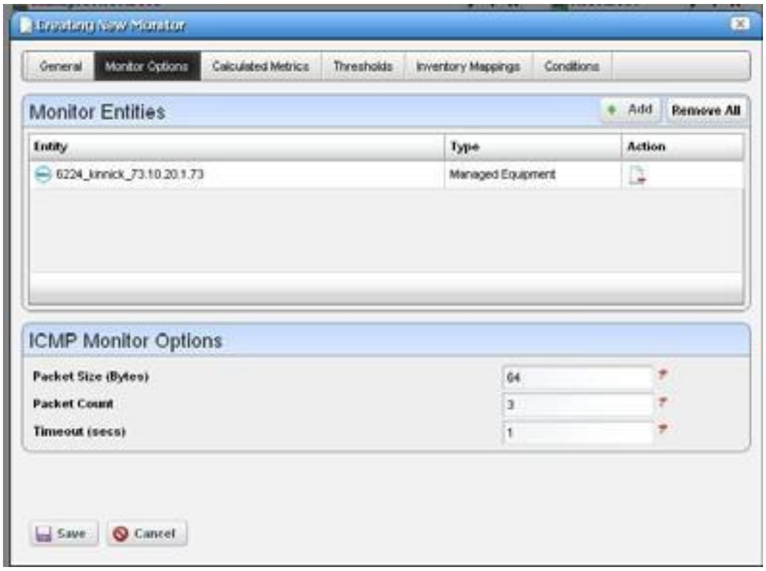

Monitor Entities

Entity	Type	Action
6224_kinnick_73.10.20.1.73	Managed Equipment	

ICMP Monitor Options

Packet Size (Bytes)	64
Packet Count	3
Timeout (secs)	1

Packet Size — Size of packet for ICMP transmission

Packet Count —Number of packets to send.

Timeout — Number of seconds without a response before a timeout is issued

The ICMP Entity Panel lets you select resource groups and Resource manager objects. Clicking *Add* button displays a selector panel for these.

Select the type of entity you want to add, then select any desired filter attributes, then click *Apply Filter*. Select from the entities that appear and add them to the monitor.

NOTE:

Migrating from previous versions updates the Network Status check box to true and redeploys the monitor.

Juniper CoS

This (optional) monitor uses the fields describe d below and lets you track CoS attributes for Juniper equipment. It appears only on systems with a Juniper device driver installed.

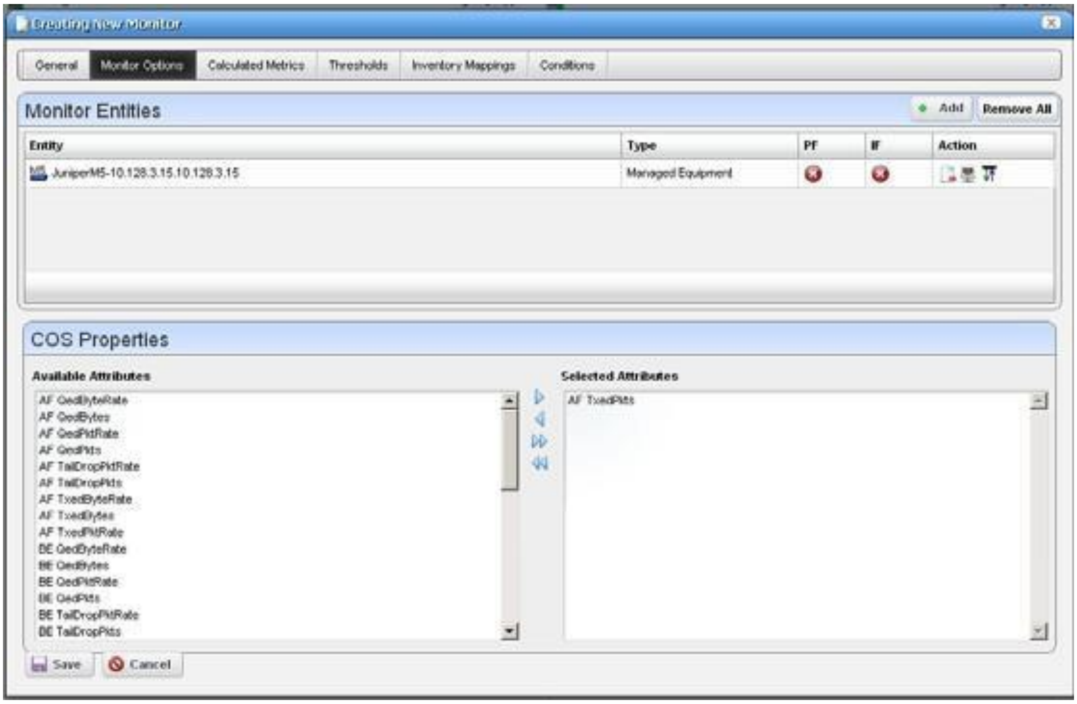

Monitor Entities

Entity	Type	PF	IF	Action
JuniperM5-10.128.3.15.10.128.3.15	Managed Equipment	✖	✖	

COS Properties

Available Attributes

AF QedByteRate
AF QedBytes
AF QedPktRate
AF QedPkts
AF TailDropPktRate
AF TailDropPkts
AF TxedByteRate
AF TxedBytes
AF TxedPktRate
BE QedByteRate
BE QedBytes
BE QedPktRate
BE QedPkts
BE TailDropPktRate
BE TailDropPkts

Selected Attributes

AF TxedPkts

Monitor Entities

Click *Add* to configure monitored devices in a subsequent selector screen. This is the typical selector with a filter to help you find discovered devices.

At the right you can see *PF* (Port Filter) and *IF* (Interface Filter) columns, which display green icons if such filters are active. Click the *Configure Port/Interface Filter* icon at the far right to configure such filters. These contain the standard filtering mechanism visible throughout Cruz. (See Defining Advanced Filters on page 145, for example). Notice that for port and interface filters, the editor also lets you delete the filter. The *Delete* button on the right of listed Monitor Entities lets you delete equipment.

COS Properties

Click to select *Available Attributes*, and use the arrows between columns to move these to the *Selected Attributes* column to select the monitored CoS properties.

Juniper RPM

This (optional) monitor uses the fields describe d below and lets you track RPM attributes for Juniper equipment. It appears only on systems with a Juniper device driver installed.

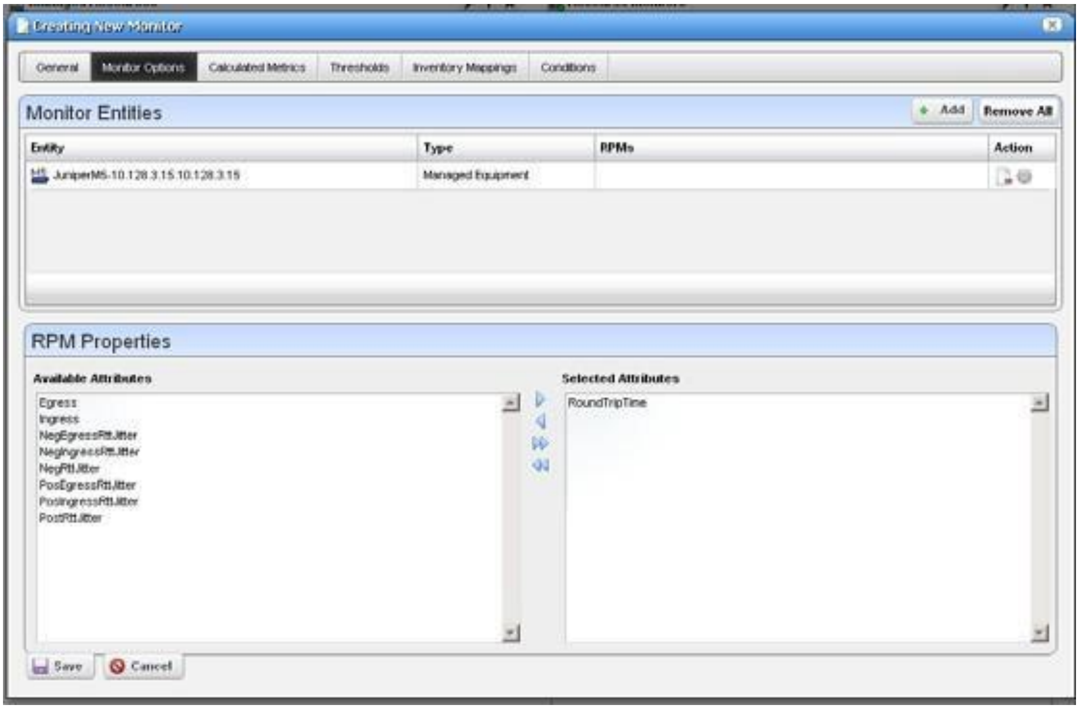

Entity	Type	RPMs	Action
JuniperM5-10.128.3.15.10.128.3.15	Managed Equipment		

Available Attributes

Egress
Ingress
NegEgressRttJitter
NegIngressRttJitter
NegRttJitter
PosEgressRttJitter
PosIngressRttJitter
PostRttJitter

Selected Attributes

RoundTripTime

Monitor Entities

Click *Add* to configure monitored devices in a subsequent selector screen. This is the typical selector with a filter to help you find discovered devices.

At the right you can see the *RPMs* column. This displays information about RPM probes. Click the *Configure RPM Probes* icon (a gear) at the far right to select and configure such probes. The *Delete* button on the right of listed Monitor Entities lets you delete equipment and probe combinations.

RPM Properties

Click to select *Available Attributes*, and use the arrows between columns to move these to the *Selected Attributes* column to select the monitored RPM properties.

These monitors collect data for all tests for the selected probe(s), and co llect only attributes assigned to them. If all attributes are assigned in the monitor, but only a handful of actual attributes are being tested, then the monitor collects only data from attributes running tests.

Forexample, if you select Egress to monitor, and tests occur for Egress, then the monitor collects Egress. However, if you select all attributes, and only Egress and PostRttJitter are tested, and the monitor collects only Egress and PostRttJitter.

Another example: If you select all attributes for the monitor, then slowly add more tests and attributes on the device, the monitor picks up these changed attributes as you add them to the tests.

The default refresh rate is 30 minutes, but you can configure the refresh byoverriding the attribute pm.monitor.rpm.refresh_rate. This property determines how often the monitor fetches the list of tests and attributes for the probe. Disable/Enable-ing the monitor also refreshes the list. The monitor can stay up to date with the device without much user intervention.

Youcan alsochange the attribute names (like Egress) too through the `pmmsgs.properties`file. Search for RPM and modify the nine attribute names. Remember, best practice is to override properties as described in Overriding Properties on page112.

Click *Save* to preserve your edits or *Cancel* to abandon them.

Key Metrics

The Key Metrics Properties panel contains a list ofkey metrics you can add to the monitor. They are grouped by category.

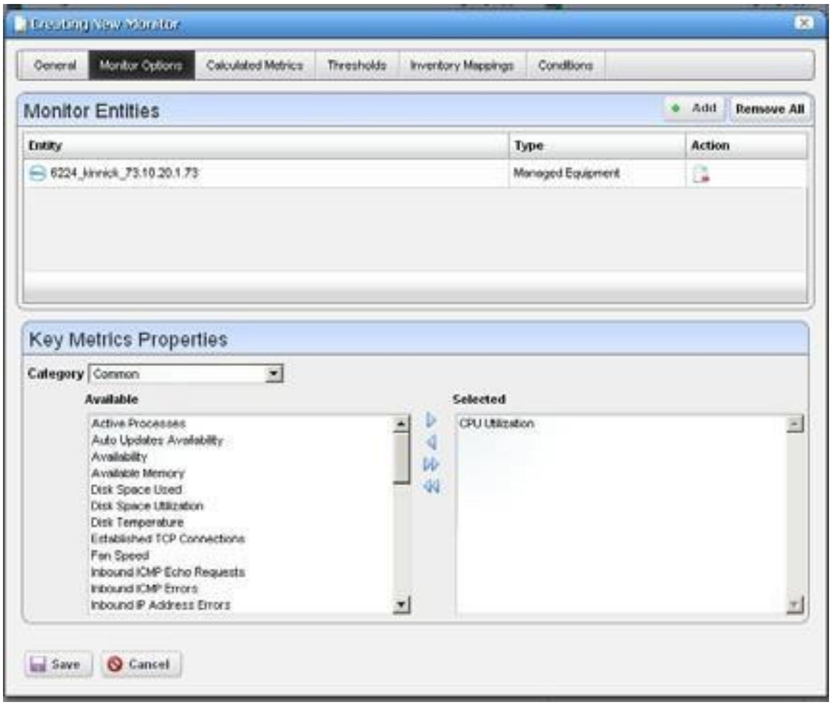

Creating New Monitor

Monitor Entities

Entity	Type	Action
6224_kinnick_73.10.20.1.73	Managed Equipment	

Key Metrics Properties

Category: Common

Available:
- Active Processes
- Auto Updates Availability
- Availability
- Available Memory
- Disk Space Used
- Disk Space Utilization
- Disk Temperature
- Established TCP Connections
- Fan Speed
- Inbound ICMP Echo Requests
- Inbound ICMP Errors
- Inbound IP Address Errors

Selected:
- CPU Utilization

The Monitor Entities Panel lets you select equipment group and equipment manager objects (as described in ICMP on page 400, above).

The Key Metrics Properties panel at the bottom of this screen uses a pre-defined list of key metrics. It does not check if the key metrics selected are supported by the devices and groups selected in the monitor.

ProScan

In this screen, you simply select the Proscan (also known as Compliance Policy) to monitor. In the Thresholds tab, you can set thresholds for both in and out of compliance numbers.

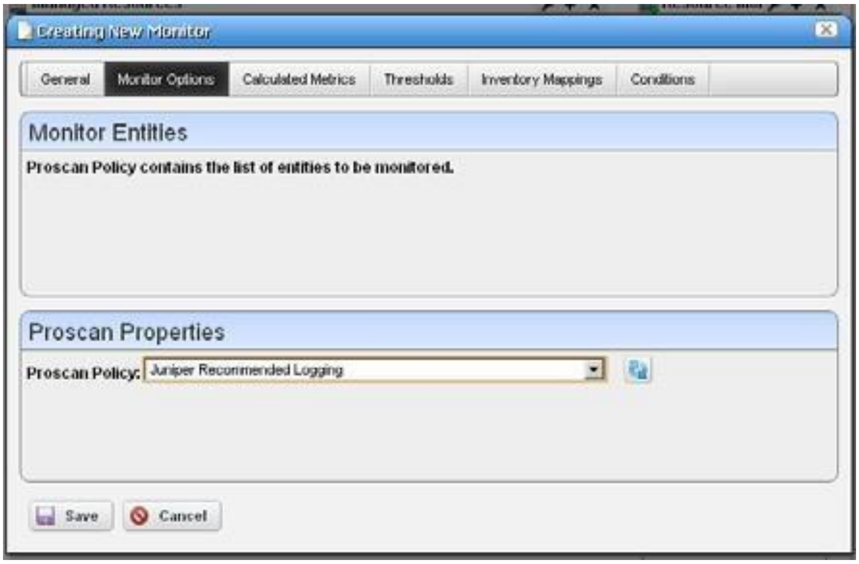

Proscan Policy contains the list of entities to be monitored.

Proscan Policy: Juniper Recommended Logging

The Proscan policy contains the target network assets.

Execute the Proscan only *after* creating the monitor. The Proscan monitor displays data when you create it and its supporting Proscan policy in the following order:

> Create Proscan policy X that has explicit targets.

> Create a Proscan monitor referring to Proscan policy X and modify polling to the desired interval.

> Execute Proscan X.

SNMP

The SNMP attributes panel lets you specify which SNMP attributes are to be monitored.

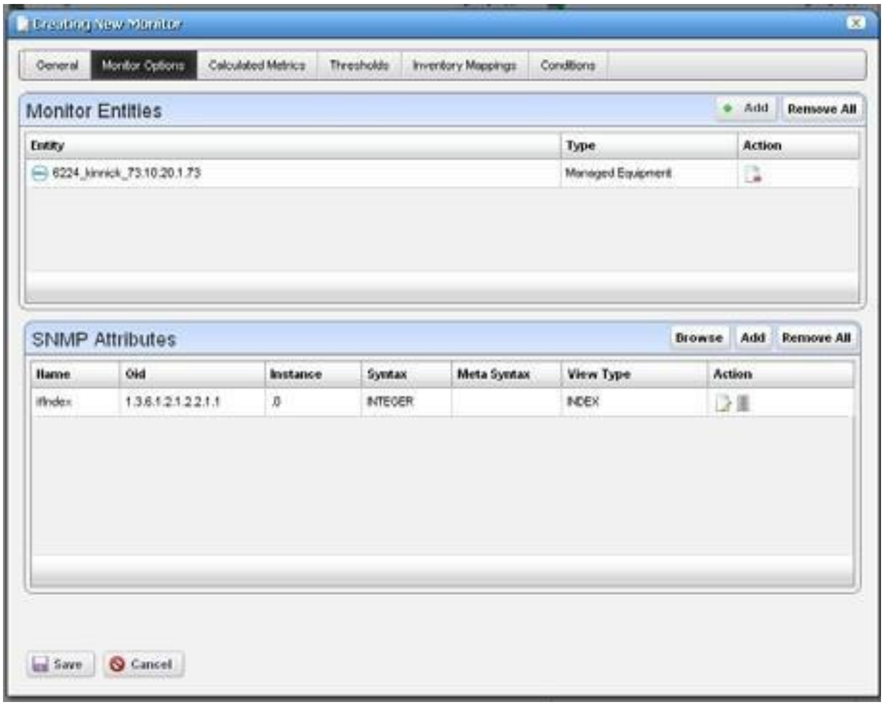

Monitor Entities

Entity	Type	Action
6224_kinnick_73.10.20.1.73	Managed Equipment	

SNMP Attributes

Name	Oid	Instance	Syntax	Meta Syntax	View Type	Action
ifIndex	1.3.6.1.2.1.2.2.1.1	.0	INTEGER		INDEX	

You can specify the SNMP attributes the following ways:

- With the SNMP browser, or
- Entering the SNMP attribute properties explicitly.

The *Browse* button launches the SNMP MIB browser. (See MIB Browser Tool on page 53) Youcan also click the *Device Results* tab to open an SNMP authentication screen and log into any device you specify, even undiscovered de vices. Specify the IP address, SNMP Read Community, port, SNMP version, timeout and retries.

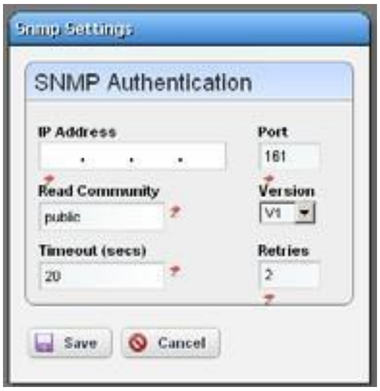

SNMP Authentication

IP Address: . . .

Port: 161

Read Community: public

Version: V1

Timeout (secs): 20

Retries: 2

Click on the desired SNMP nodes and then click on the *Add Selection* button to add an SNMP attribute. When done selecting, click the *Done* button to add selected attributes to the monitor or *Cancel* to abandon the operation and close the browser.

The Add and Edit buttons in the SNMP attrib ute panel launch the SNMP Attribute editor.

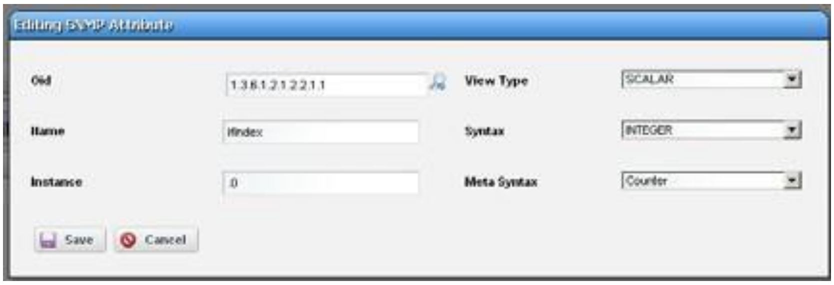

Editing SNMP Attribute

Oid	1.3.6.1.2.1.2.2.1.1	View Type	SCALAR
Name	ifIndex	Syntax	INTEGER
Instance	0	Meta Syntax	Counter

This panel contains the following properties:

Oid—The object identifier for this attribute

Name — This attribute’s name

Instance — SNMP instance. 0 for scalar or the ifIndex value for an SNMP column.

View Type — *Scalar* or *Column.*

Syntax — *Integer, Boolean, DisplayString,* and so on.

Meta Syntax — *Counter, Gauge,* and so on.

If you type in an OID and click the search button next to the OID field, the browser searches the MIB for the OID and fills in the other values if it finds the OID.

NOTE:

For all counter types, the polled data stored in database reflect the changes between two polled data points from an SNMP table.

SNMP Interfaces

The SNMP Interface Monitor Entity editor supports the followingentity types: group, equipment manager, port and interface. It also supports port and interface filters on groups and equipment manager objects.

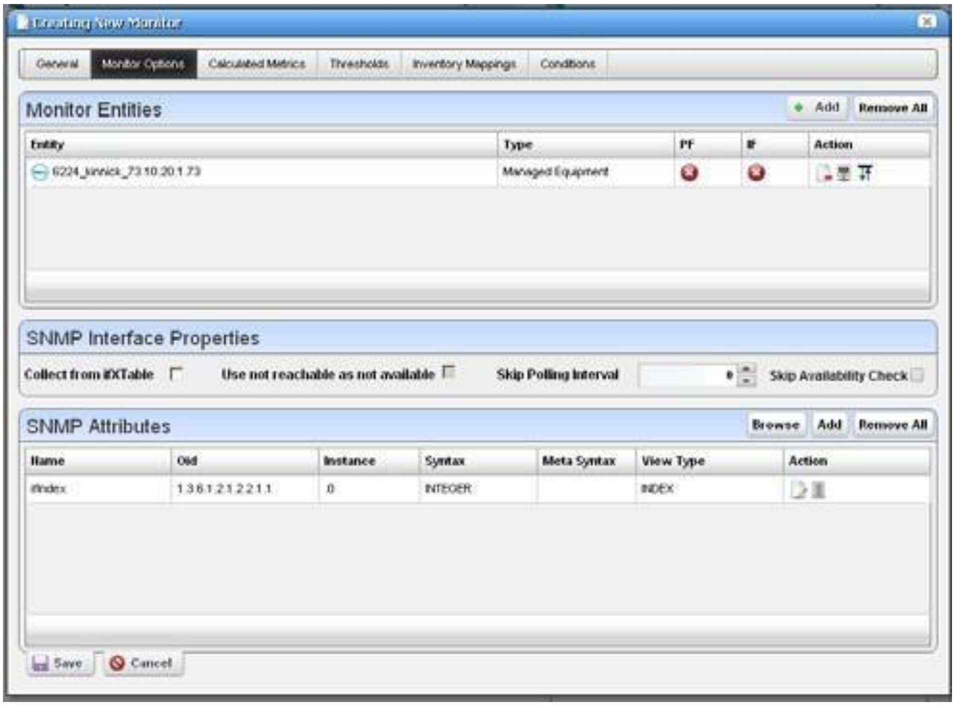

Entity	Type	PF	IF	Action
6224_kinnick_73.10.20.1.73	Managed Equipment	☒	☒	

Collect from ifXTable □ Use not reachable as not available □ Skip Polling Interval 0 Skip Availability Check □

Name	Oid	Instance	Syntax	Meta Syntax	View Type	Action
ifIndex	1.3.6.1.2.1.2.2.1.1	0	INTEGER		INDEX	

If you check the *Collect from ifXTable* checkbox, then Cruz attempts to fetch attributes from the ifXTable. These attributes are ifHighSpeed, ifHCInOctets, ifHCInUcastPkts, ifHCOutOctets and ifHCOutUcastPkts. If any of these attributes are not available, then it fetches from ifTable.

If an interface does not support ifxTable, SNMP ge t typically retrieves an error and Cruz uses the ifTable instead. Some ATM ports do not send errors from the ifxTable oids, so Cruz also uses the ifTable values if ifHighSpeed is 0.

NOTE:

The SNMP V1 protocol does not support 64bit counters located in ifXtable. This means Cruz monitors only collect performance data from ifTable when a device is discovered using the SNMP V1 protocol. Best practice: Discover devices using snmpV2c orsnmpV3 protocols to collect performance data located in ifXtable.

Even with this checked, Cruz defaulting to 32-bit counters if 64-bit is not available.

Cruz now supports multiple indexes in the SNMP Interface monitor. Specify them in the instance field, separated by dots. For `sfpTxPowerValue 1.3.6.1.4.1.28458.7.2.4.6.7.1.22.Y.Z`, where `Y` is the slot and `Z` is the @ifindex, specify `@slotNumber.@ifIndex` as the instance. Youcan also specify a constant string. For `sonetLineIntervalUASs 1.3.6.1.2.1.10.39.1.3.2.1.5. X.16` Where `X` is the `@ifindex` and `16` is the last record, specify `@ifIndex.16`.

The variable name following the “@” must correspond to an attribute in the port or interface bean.

When determining the “not available” status of a device, SNMP AdminStatus and OperationalStatus messages both have to indicate adevice is Available before a monitor determines it is available.

Certain devices that do not support ifTable availability indicators. For the sake of these devices, a *Skip availability check* checkbox appears.

The *Skip Polling Interval* configures skipped availability chec ks when polling, so you can check availability, forexample, everyfourthpollinginte rval(skippingthree). Thishelpsthe monitor avoid flutter artifacts

The PF and IF table columns indicate if a port filter or interface filter is configured for the entity. Click the icons on the right side of the list of Monitor Entities to configure filters. Clicking these buttons displays an interface configuration panel.

This panel lets you specify filter at tributes for the port or interface filters you want to monitor. For example, if you select a device but only want to monitor active interfaces created by a particular user, then these filters do the job.

The SNMP Attributes panel is the same as described in SNMP on page 405.

SNMP Table Monitor

This panel appears if you are editing an SNMP Table monitor. The application stores not absolute numbers from counters but the counter’s change since its last measurement.

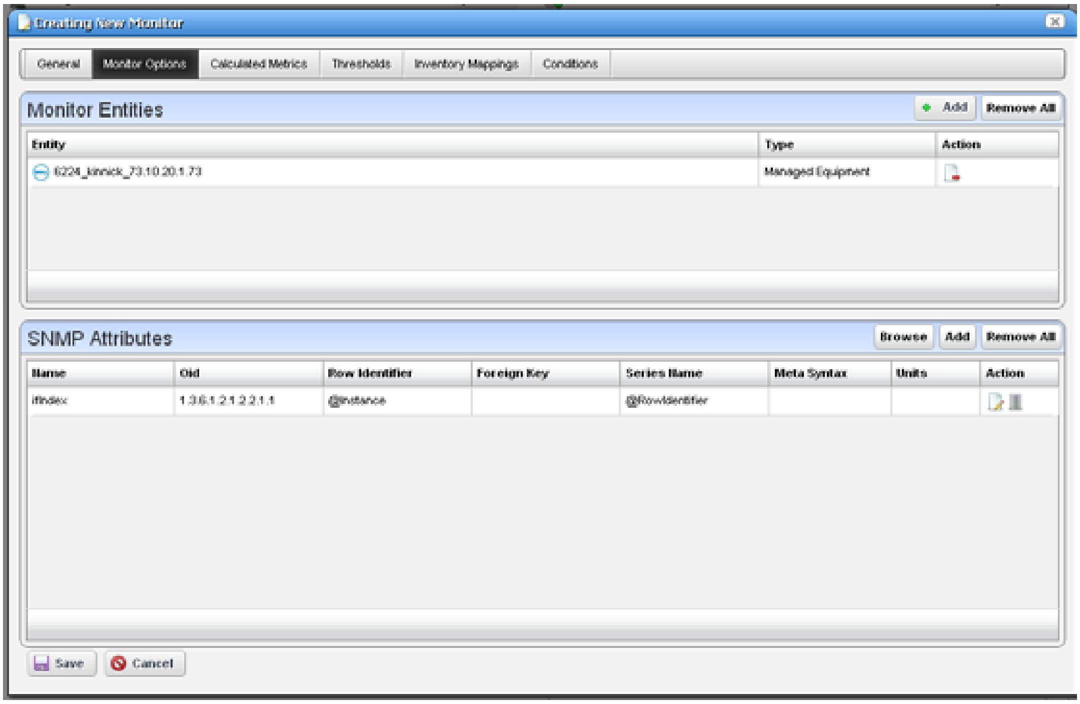

Monitor Entities

Entity	Type	Action
6224_kinnick_73.10.20.1.73	Managed Equipment	

SNMP Attributes

Name	Oid	Row Identifier	Foreign Key	Series Name	Meta Syntax	Units	Action
ifIndex	1.3.6.1.2.1.2.2.1.1	@instance		@RowIdentifier			

Columnsincludethe SNMPAttribute Name, OID, RowIdentifier, Foreign Key,Series Name, Meta Syntax, Units, and Action.

NOTE:

If you select one of the 64-bit counters in ifXTable, make sure the Meta Syntax is 64-bit.

Clicking the *Add* or the *Edit* button to the right opens either a MIB Browser where you can retrieve these attributes, or an Add/Edit SNMP Attributes editor at the bottom of the screen, See the following sections for details.

MIB Browser

This lets you select attributes to monitor as described in MIB Browser Tool on page 53. The SNMP table monitor lets you pick a table column, not the entire table.

Add/Edit SNMP Attributes

This screen lets you specify individual attributes.

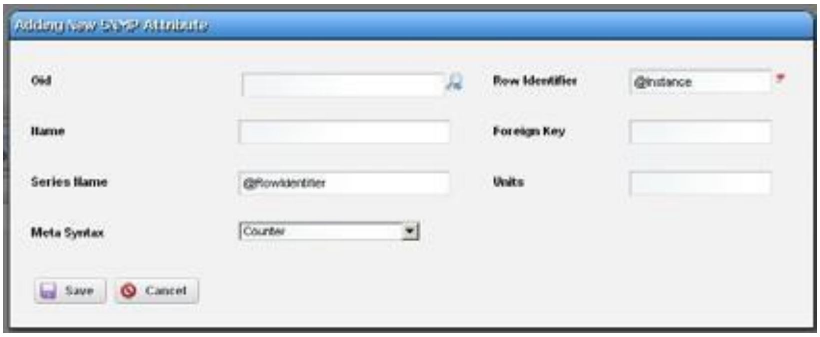

Row Identifier: @instance

Series Name: @RowIdentifier

Meta Syntax: Counter

It has the following fields:

Oid— A field where you can enter the object identifier. Thisalso hasan integrated search function. Click the magnifying glass icon on the right to activate it. A successful search populates the rest of the fields for the object identifier.

Row Identifier — This mandatory field defaults to @instance (The OID instance).

Name — The text identifier for the OID

Foreign Key — Enter the foreign key, if any.

Series Name—This defaults to @RowIdentifier.

Units—Enter the units of measurement.

Meta Syntax—Further refine the variable type with the pick list. For example, you can select *Counter32* (a 32-bit counter). For Counter types, the monitor computes change from previous readings, and for Gauges it does not.

NOTE:

If a message appears saying: "Device fault: Return packet too big" in the Monitor Status Summary, then you have selected too many SNMP attributes to poll in a single request. Please modify your monitor to request smaller numbers of attributes

VRF

Select the device with the VRF you want to monitor in this screen.

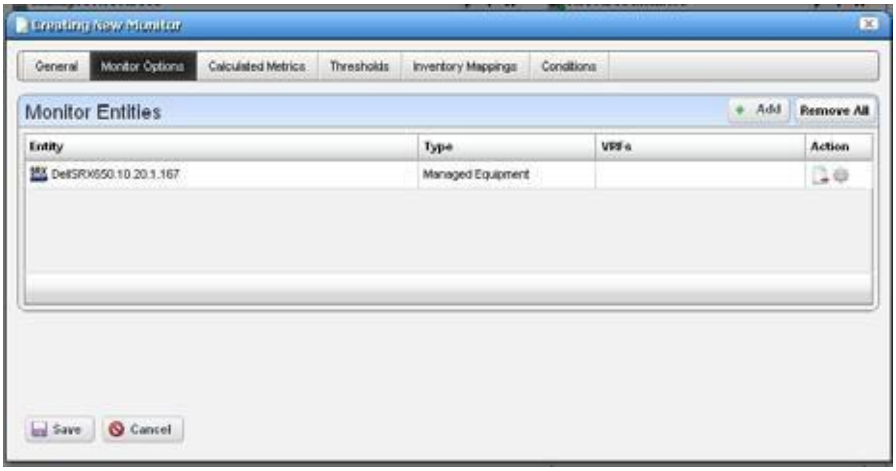

Entity	Type	VRFs	Action
DellSRX650.10.20.1.167	Managed Equipment		

The monitor calculates the average interfaces utilization associated with the VRF on the selected target device, based on the current VRF configuration. Click *Save* to preserve your choice.

Bandwidth Calculation

Cisco devicesrunning IOS or IOS-XR can calculate bandwidth. Tosupport this functionality, all ports and interfaces now have four new filterable attributes:

- Ingress Bandwidth (the ingress bandwidth in bps)
- Ingress Bandwidth Type (the type of calculation used to determine the ingress bandwidth)
- Egress Bandwidth (the egress bandwidth in bps), and
- Egress Bandwidth Type (the type of calculation used to determine the egress bandwidth).

Types of Bandwidth Calculation

The ways to arrive at a bandwidth number appear here in order, from highest priority to lowest priority. If Cruz can calculate bandwidth in more than one way for a particular port/interface, it uses the highest priority calculation type.

NOTE:

If a port or interface's Administrative State is not Up, its bandwidth will always be 0, regardless of the calculation type!

CONFIGURED — This means Cruz has a QoS policy configured directly against this port or interface. Forexample: a policer, shaper, or queuing policy. Policies applied to *input* affect the ingress bandwidth, *output* affect the egress bandwidth.

TRUNK AGGREGATION—Thiscalculation means a port is in trunking mode and calculatesits bandwidth from its access ports. Essentially, if a trunk shares a VLAN with any access port, that access port adds its bandwidth values to the trunk port's totals. This reverses ingress and egress values. The total ingress bandwidth of the access ports becomes the egress bandwidth of the trunk port, and the total egress bandwidth of the access ports becomes the ingress bandwidth of the trunk port. If a port in this configuration is no longer in trunking mode, it reverts to UNCONFIGURED.

INTERFACEAGGREGATION —Ifaport orinterface hassub-interfaces, the total bandwidth of the parent is the sum of the bandwidth of its children. Forexample: If GigabitEthernet0/1 has four subinterfaces GigabitEthernet0/1.1, GigabitEthernet0/1.5, GigabitEthernet0/1.8 and GigabitEthernet0/1.9 and each sub-interface has bandwidth of 1G, the total bandwidth of GigabitEthernet0/1 will be 4G.

Unlike trunk aggregation, this does not reverse ingress and egress. The total ingress of the children becomes the total ingress of the parent, and the total egress of the children becomes the total egress of the parent. If a port or interface in this configuration loses its children (i.e. the interfaces all get deleted), it reverts to UNCONFIGURED.

ASSOCIATION—If a port is currently UNCONFIGURED and has a physical link to another port that isn't UNCONFIGURED, it takes the bandwidth of the linked port. This reverses Ingress and Egress—the linked port's ingress becomes the other port's egress and vice-versa. If someone deletes the link, the port reverts to UNCONFIGURED.

UNCONFIGURED — The default setting for bandwidth. This sets the ingress and egress bandwidth of a port or interface to the IfSpeed of the port. This means a 10G port with an IfSpeed of 10G registers an Ingress and Egress Bandwidth setting of 10G.

Triggering BandwidthCalculations

On resync, Cruz’s rules check for configured QoS policies and update the port and interface bandwidth as needed (the CONFIGURED bandwidth calculation type). Cruz adds any ports and interfacesregisteringachangeintheirbandwidthvalues(oranynewly created portsand interfaces, as in initial discovery) to a list to be processed after the resync is over. Cruz also queues the device for recalculation of TRUNK AGGREGATION every time it collects VLAN data (for example, during resync or during network data collection).

After the resync finishes, the Cruz bandwidth processor processes the list of ports and interfaces whose bandwidth values changed and re-check its calculations to see if INTERFACE AGGREGATIONorASSOCIATIONcalculationtypesareapplicable, and thencalculate themif necessary.

Link creation, modificationordeletion also trigger recalculationsof ASSOCIATION calculation type for the endpoints of the link in question.

Cruz adds any ports/interfaces that change in the bandwidth processor back to the processor, so that changes can propagate throughout the network. For example, if Cruz discovers a link and it changes the linked port's bandwidth (type ASSOCIATED), it needs to recalculate the TRUNK AGGREGATION for all trunk ports on that device in case the port wasan access port and the trunk port bandwidth values need to be updated.

Another example: If an interface's bandwidth value changed, then Cruz adds its parent for INTERFACE AGGREGATION reprocessingin caseit has aparent usingthat calculation type that now needs to be updated.

Cruz uses a strict priority to determine which calculation method to use if multiple is applicable: CONFIGURED > TRUNK AGGREGATION > INTERFACE AGGREGATION > ASSOCIATED > UNCONFIGURED. This means that if a port has a direct QoS configuration against it, it doesn't matter if it also has child interfaces or links. Cruz uses the QoS configuration's value. Likewise, if Cruz calculates aport’sbandwidth using INTERFACE AGGREGATION and it has a link, the link does not matter for the purpose of bandwidth calculation since ASSOCIATED is lower priority.

Scheduling Refresh Monitor Targets

Becausemonitorscanaddresstargetsthatare members ofdynamic groups, refreshing these ensures that group memberships are up to date. Todo this, you can create oralter the schedule for Monitor Target Refresh (in most packages, such a scheduled item appears by default). When executed, this updates monitors with groups as targets based on current memberships. This removes targets no longer members of a monitored group and adds new group members. A seeded schedule refreshes every six hours, by default.

Refresh Monitor manually by right clicking in the Resource Monitors table.

Refresh Monitor Targets for Newly Discovered Devices

If you discover a new device that is part of a monitored dynamic group, it may take some time before monitoring includes that device. To provide immediate monitoring, as soon as discovery finds the device, add the *Refresh Monitor* action to the discovery profile. See *Actions* on page 102 for more about that Discovery Profile capability.

Tomake sure this refresh occurs, do *not* override the following in `redcell.properties` (This section defines the actions executed when no default discovery profile exists):

```
#This shows the default order of Task Activities within Resource Discovery
 Options
#The TaskDefOid will be used for identification and true/false will
 determine if they are
#on by default
#format: <TaskDefOid>&&true,<TaskDefOid2>&&false,<TaskDefOid3>&&true
redcell.discovery.taskactivity.order=Resync&&true,\
DataCollectionForGroupOfDevices&&false,\
Discover_Links_for_a_Group_of_Devices&&false,\
Scheduled_Resync&&false,\
Refresh_Monitor_Targets&&true
```

Topological Correlation

A device can appear unresponsive in monitors if devices through which Cruz must access it are down, even though it may be active independen tly of the condition of its access. Topological correlation takes this into account, and produces fewer false MonitorTargetDown events because Cruz attempts to communicate with adjacent devices to the first device in the network topology. Cruz calculates adjacency according link configuration. If adjacent devices are unreachable, then Cruz doesnot generate a MonitorTargetDown event because this first device is simply unreachable from Cruz now.

Updating Polling Subscriptions

Polling subscriptions on the mediation agent process can get out ofsync with the application server process. When the application server and mediation agent start running, whichever one comes online last triggers applicationserver sending the polling subscription and target information to the mediation server. Enabling ordisablingperformance monitors also sends this information fromthe application server to the mediation server.

Any time data goes from one machine to another, temporary connectivity issues can arise, along with potential for data loss, so Cruz accommodates this possibility too. If a server is running for a long time (weeks) and the performance monitors have been frequentlyenabled and disabled, some targetsmaynotbepolledbecause theappserver/medserver information hasbecomeoutofsync. It’s possible, but less likely, that polling, and target information could be out ofsync even on standalone systems, where the application server also serves as a mediation agent. Therefore, the following feature is also available on standalone systems.

Youcan schedule periodic resyncs of polling subscription and target information. The scheduled item that runs this process is disabled by default, but if enabled it typically runs every 30 minutes. You can also enable this item and schedule it often than every 30 minutes.

If you enable this scheduled item, or run it a single time manually, then it ensures the mediation agent has all of information it needs for polling. This includes the polling subscription information associated with each active performance monitor and all active targets associated with each active monitor.

Top N [Assets]

The Cruz uses seeded, default Active Performance Monitors (APM) to displayperformance data in severalcategories. The Top *Asset* portlets displaydevice monitoring summaryresults. Forexample, the Top Ping Response (Slowest) portlet displays the devices that are the slowest to respond to a ping.

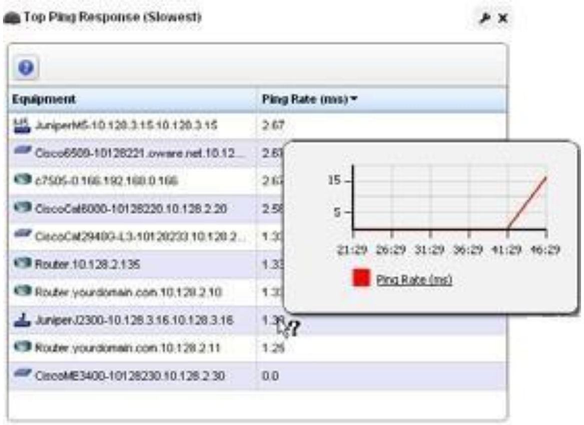

Top Ping Response (Slowest)

Equipment	Ping Rate (ms)
JuniperM5-10.128.3.15.10.128.3.15	2.67
Cisco6509-10128221.oware.net.10.12...	2.6
c7505-0.166.192.168.0.166	2.6
CiscoCat6000-10128220.10.128.2.20	2.5
CiscoCat2940G-L3-10128233.10.128.2...	1.3
Router.10.128.2.135	1.3
Router.yourdomain.com.10.128.2.10	1.3
JuniperJ2300-10.128.3.16.10.128.3.16	1.3
Router.yourdomain.com.10.128.2.11	1.25
CiscoME3400-10128230.10.128.2.30	0.0

Devices listed are ranked by the monitored parameter. Hover over the Ping Rate column to view recent activity in a pop-up graph.

If you right-click a monitored item, you can select from menu items like those that appear in the portlet described in Managed Resources on page 177.

Forsomeportlets (forexample Top CPU/Disk/Memory Utilization, Top Interface Bandwidth/ Errors), the right-click Performance menu it ems include Key Metrics. The menu can include *Performance*, which displays Dashboard Views related to the selected monitor.

For some packages, these can also include IP SL A statistics like the following: Top Bandwidth Received/Transmitted, Top CPU/Disk Utilization, Top Ingress/Egress Packet Loss, Top Jitter, and Top RT Delay. To see all available *Top* portlets, click *Add* > *Applications* and look below *Top N* on the subsequent panel.

Top RT Delay maps to the AvgRTDelay inventory metric. When no metric for Average RT Delay existsin the MIB, Cruz calculates it the average RT Delay using two MIB attributes: RTTSum (the total time taken for all round trips) and NumOfRTT (the number of round trips taken). The calculation of AvgRTT is the value of RTTSum divided by the value of NumOfRTT. Cruz maps this attribute to the AvgRTDelay inventory metric.

NOTE:

An alternative way toprovidethis kind of performance information isto use the Traffic Flow Analyzer. For systems generating large amounts of information that strain the limits of processing capacity, see Best Practices: Performance Tuning Traffic Flow Analysis onpage 536 as a possible solution.

Calculations within Top N Portlets

The following Top N portlets are potentially available. Only those with monitored parameters displaydata. The data comes from monitor data using the monitor inventory mappings specified in each monitor.

The tooltip graph shows the values for the attribute over the last 30 minutes. The Errors and Discards attributes are counter values that show the change in value since the previous polling cycle. The other attributes are gauges which displaya rate or percentage. The value displayed for each entry is the average over the last 30 minutes for the gauge attributes and the sum of values over the last 30 minutes for counter attributes.

These top asset portlets are available from the Performance page or can be added using the Add > Application menu option as needed.

The following portlets display data based on equipment targets:

Portlet	Description
Top CPU Utilization	Percentage of CPU used. Located on the Performance Summary page by default.
Top Disk Utilization	Disk use.
Top Memory Utilization	Memory use. Located on the Performance Summary page by default.
Top Ping Response (Slowest)	Lowest ping response. Located on the Performance Summary page by default.

The following portlets display data based on port or interface targets:

Portlet	Description
Top Interface Bandwidth	Most interface bandwidth use. Located on the Performance Summary page by default.
Top Interface Errors	Most interface errors. Located on the Performance Summary page by default.
Top Input Discards/Errors	Most input discards/errors. The tooltip/graph that appears when you hover your cursor over a row in these portlets shows the change in discards or errors, then the Cruz system adds changes to the base value and that sum appears within the table. Top Input Errors portlet is located on the Performance Summary page by default.

Portlet	Description
Top Output Discards/Errors	Largest number of output discards/errors. Top Output Errors portlet is located on the Performance Summary page by default.
Top Bandwidth Received/Transmitted	Displays percentage of bandwidth use received or transmitted.
Top Bandwidth Received (bps)/Transmitted (bps)	Displays bandwidth use in bytes per second. Located on the Performance Summary page by default.

The following portlets display data based on SLA or VRF targets:

Portlet	Description
Top Egress/Ingress Packet Loss	Most egress/ingress packets lost. Located on the Performance Summary page by default.
Top Jitter	Highest jitter rates
Top MOS	Highest MOS (a network performance measurement).
Top Packet Loss	Greatest packet loss. Located on the Performance Summary page by default.
Top RT Delay	Longest round trip (RT) ping delay

The following portlets are not monitor-based:

Portlet	Description
Top Problem Nodes	Devices with the highest alarm state
Top Configuration Backups	The most recent backups

Displaying Tenant Domains in Top N Portlets

If you have implemented a multitenant (MSP) system, but want the master domain to display Top N for just a tenant domain, the key icon in the portlet's toolbar lets you select different sub-domains.

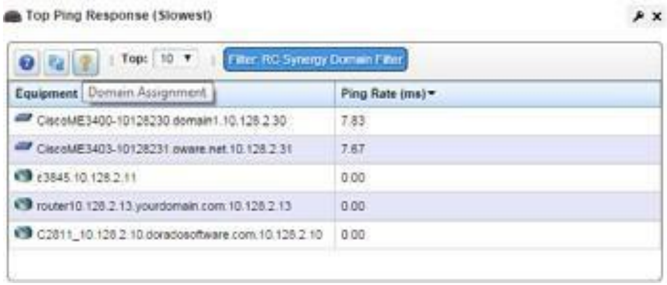

Top Ping Response (Slowest)

Top: 10

Filter: RC Synergy Domain Filter

Equipment	Ping Rate (ms)
CiscoME3400-10128230.domain1.10.128.2.30	7.83
CiscoME3403-10128231.oware.net.10.128.2.31	7.67
c3845.10.128.2.11	0.00
router10.128.2.13.yourdomain.com.10.128.2.13	0.00
C2811_10.128.2.10.doradosoftware.com.10.128.2.10	0.00

The filter label in the toolbar displays which domain has been selected. Once you filter a Top N portlet this way, it displays results only for equipment authorized for the selected domain.

Top Configuration Backups

This panel lists the most recent configurations back ed up from devices. The pick list in the upper

right corner lets you select not just the top 10 such backups, but the top 5, 10, 15, 20, and 25.

Right clicking a backup offers the same options as the portlet described in Configuration Files on page 446.

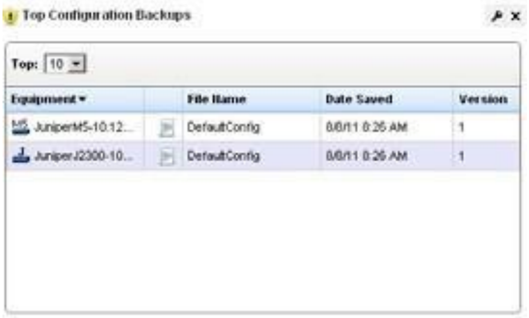

Top: 10

Equipment ▾		File Name	Date Saved	Version
JuniperM5-10.12...		DefaultConfig	8/8/11 8:26 AM	1
JuniperJ2300-10...		DefaultConfig	8/8/11 8:26 AM	1

Dashboard Views

The Dashboard Views portlet lets you assemble several monitors into a single display, or dashboard.

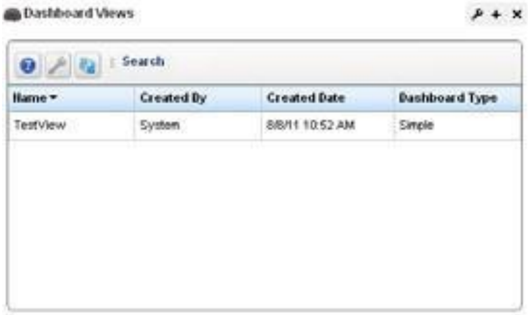

Name ▾	Created By	Created Date	Dashboard Type
TestView	System	8/8/11 10:52 AM	Simple

Youcan create and displaydashboards byright clicking itemsin the Managed Resources portlet, selecting *Show Performance*,or by selecting *New* in the *Dashboard Views* portlet.

Right-click the listed dashboards, and a menu appears that lets you *Rename*, *Delete*, *Copy*, *Edit*, create a *New* simple or custom dashboard, or *Launch* a Dashboard View(either *Maximize*—a larger view—or as a *Popup*). You can also import/export views and share views with other users on your system. See Dashboard Editor on page 420 for information about creating or modifying dashboards. Foran explanation of *Convert*, see *Convert Simple Dashboards to Custom Dashboards* on page 425.

The Performance Dashboard on page 419 and Dashboard Editor on page 420 describe configuring simple dashboards. See the How to: Create a Custom Dashboard View on page 421 section for a description of custom dashboard view creation.

Youcan also Convert Simple Dashboards to Custom Dashboards, as described below. When you *Edit* a view, Dashboard Editor appears. It lets you select which monitors appear in the dashboard, the monitored entities, and attributes.

The expanded portlet offers similar capabilities. To make a monitor appear on a page, use the portlet described in Performance Dashboard on page 419.

When you create dashboards, data rollup is part of what the display shows. If, for example, the monitor displays the results from a boolean (0 or 1 output), rollup may average values for a duration, and values less than one will appear in the graph.

△ CAUTION:
RevisionslikedeletingHierarchical Viewportletsfromapagerequireapagerefreshbeforethe dashboardworkscorrectly. Somepackagescontaina *System* dashboardthatmaynotletyouselect monitors.

Launch a Dashboard View

Launching a view lets you view the monitors active for a Dashboard view.

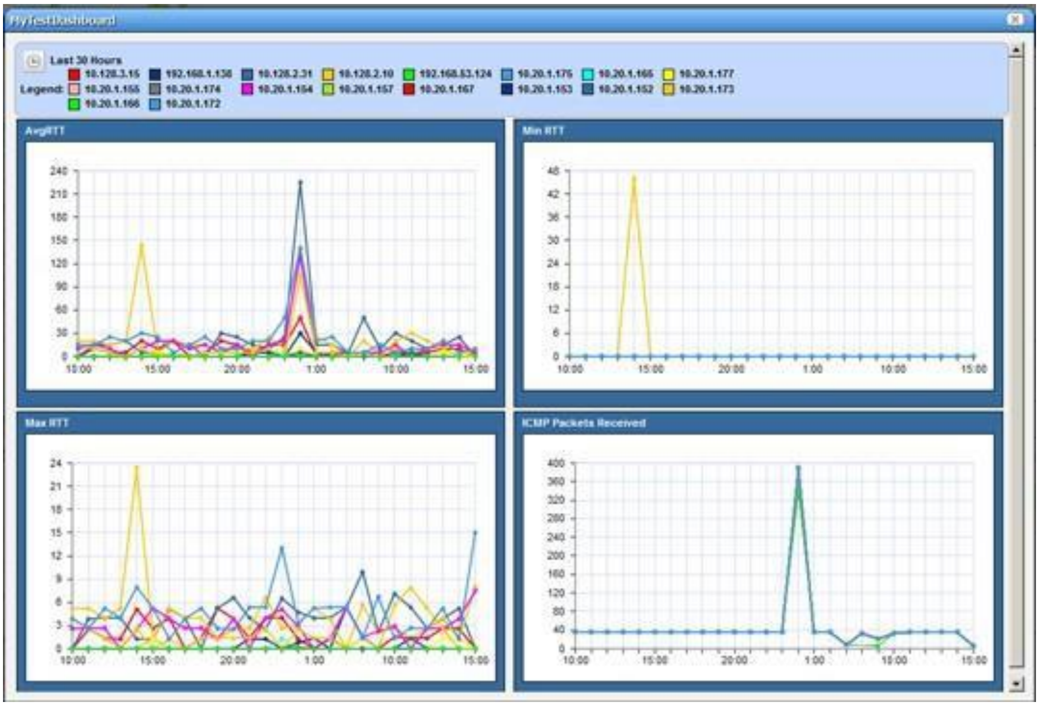

Some packagesdisplaya *Network Dashboard* bydefault. If the Network Dashboard portlet isblank, you can create a new one. Click the *select new* text in the upper right corner of the portlet to select an alternative, already configured view from those in Dashboard Views portlet. Click the *edit* button in that same corner to alter the configuration of any existing dashboard. See Dashboard Editor on page 420 for more about altering views.

Youcan configure Dashboards appear by configuring them in the Dashboard Views portlet, or by selectingadeviceordevicesfrom the Managed Resourcesportlet, right clicking and choosing *Show Performance*. To select more than one device, use the expanded Managed Resources portlet.

The first time you create a default template dashboard for a single device, Cruz saves it in the Dashboard Views manager. Invoking *Show Performance* for thatdevicesubsequently displaysits default view.

The icons in the dashboard's upper right corner let you edit *Dashboard Properties* with the Dashboard Editor, or *Save* the dashboard with the other icon.

NOTE:

No need to reload the browser to update a dashboard; it reloads data every 30 seconds by default, with less overhead.

Displaying Values

Hovering the cursor over the individual points displays the charted attribute value(s) as popup tooltips. If a graph has multiple lines, the data points for different lines are charted at different times (Cruz distributes polling to balance the load on its mediation service). Hover the cursor over the time when a line's data point appears, and that line's value appears as a tooltip. It may seem a device reporting the same value as others is not graphed properly, but mousing over the graph displays the value.

The legend of devices and/or attributes that appear in each graph also provides interactive features. Hover your cursor over a device or attribute color in the legend and only that device or attribute appearsonscreen. Bydefault, allsuch legend colorsquarescontainchecks. Uncheck the ones you do not want to see. The legend can appear consolidated or for each chart, as is appropriate to the distribution of charted devices and attributes.

If no data is available for an attribute in a dashboard, no panel appears for that data.

Changing Dashboard Time/Date Format

Control panel's Redcell > Application Settings screen has a *Performance Chart Settings* panel where you can set the *Day Format* and *Minute Format* so dashboards display time (the x axis) in a meaningful way. If you want European date formats (day/month/year rather than month/day/year), this is available if the language/location settings of the operating system on the computer running Cruz makes it available.

In Control Panel's *Performance Chart Settings* panel, you can also enable Threshold display in dashboards and elect to *Restrict Y-Axis Range to data range* withcheckboxes.

How To:

Create a Simple Dashboard View

Follow these steps to create a simple dashboard view. See How to: Create a Custom Dashboard View on page 421 for more complex monitor creation.

1. In the Dashboard Views portlet, right click to select *New > Simple Dashboard.*
2. Select a name (for example SNMP Interface, to display the monitor configured in How to:Create an SNMP Interface Monitor on page 378).
3. Click *Add Entity* in the Entities panel.
4. In the filter that appears, select the type: Interface.
5. Filter for the IP address of the entity monitored in the previous SNMP interface monitor creation, select it and click *Add Selection* and *Done.*
6. Select the ifInErrors attribute and click the right arrow in the Dashboard View Attributes panel.
7. Click *Save*. The dashboard view you have configured should appear in the portlet.
8. To launch it, right-click and either *Launch (Popup)* or *Launch (Maximize)*

9 If you want to convert this simple dashboard to a custom dashboard so you can alter it further, right-click and click *Convert*.

10 Notice that you can also change the time/date format as described in Changing Dashboard Time/Date Format above.

NOTE:

To improve performance, SNMP Interface Monitor does not retain polled data by default. You cannot use historical attribute data in a dashboard without this retention.

Performance Dashboard

This portlet lets you install and configure Dashboard Views as permanent displays rather than portlets. When youinitiallyinstall thisportlet, it appearsempty. The message “No Dashboard View has been set:” appears with a *Select* button. Click that button to open the Dashboard View Selection screen.

Dashboard View Selection

This screen displays any existing dashboards so you can select one for the Performance Dashboard you want to appear on a page in Cruz.

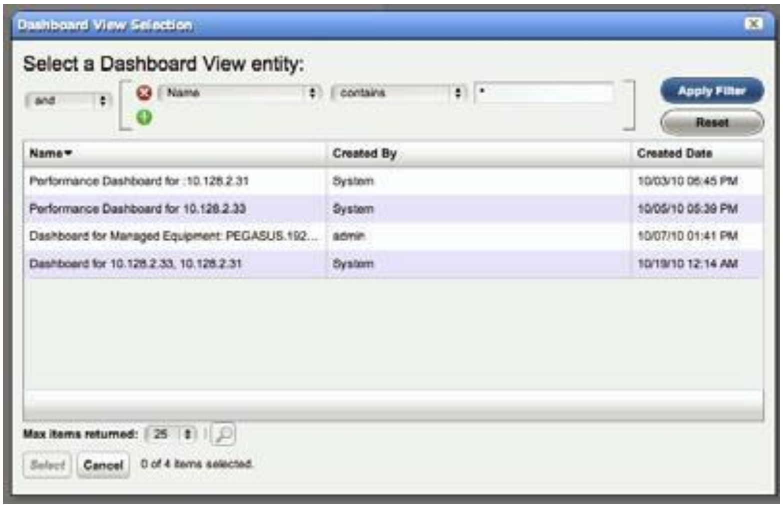

Name filter: contains *

Name	Created By	Created Date
Performance Dashboard for :10.128.2.31	System	10/03/10 06:45 PM
Performance Dashboard for 10.128.2.33	System	10/05/10 05:39 PM
Dashboard for Managed Equipment: PEGASUS.192...	admin	10/07/10 01:41 PM
Dashboard for 10.128.2.33, 10.128.2.31	System	10/19/10 12:14 AM

Max items returned: 25

0 of 4 items selected.

Use the filter at the top of this selector to limit the listed dashboards from which you can select. See Dashboard Views on page 416 for more about creating and configuring the views from which you select.

NOTE:

If you delete the Network Status Dashboard can put it back by adding the Performance Dashboard portlet to the desired page, then select the desired Dashboard View you would like to display as your Network Dashboard.

Dashboard Editor

When you *Edit* a dashboard by right clicking a resource from the Managed Resources portlet and selecting *Show Performance*, or create (select *New*) a dashboard from the Dashboard Views portlet, an editor appears that lets you select and rearrange the monitor components of the dashboard.

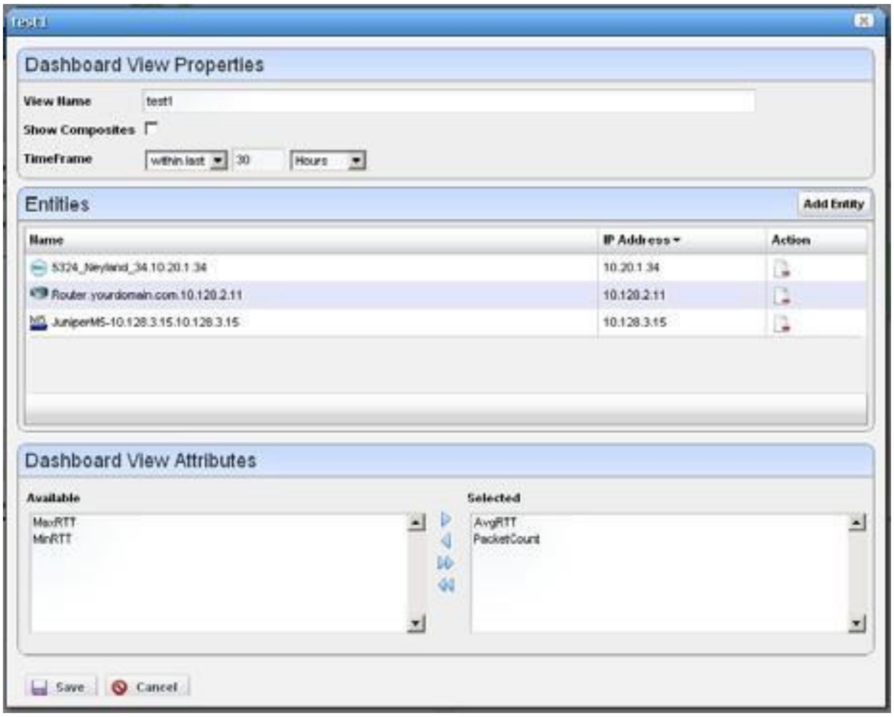

View Name: test1

Show Composites: □

TimeFrame: within last 30 Hours

Name	IP Address	Action
5324_Neyland_34.10.20.1.34	10.20.1.34	
Router.yourdomain.com.10.120.2.11	10.120.2.11	
JuniperM5-10.128.3.15.10.128.3.15	10.128.3.15	

Available: MaxRTT, MinRTT

Selected: AvgRTT, PacketCount

This screen has the following fields:

View Name—The identifier for the dashboard. The de fault is “Performance dashboard for [IP address],” but you can edit this. This is what appears in the Dashboard Views list.

Show Composites — Show attributes that are constructed from other attributes. Composites attributes are special attributes that consist of the attribute name and the instance name. For example: CPU Utilization:cpu1. Some KPI metrics are composite. If you use SNMP Table monitor, then pretty much all values retrieved are composite.

TimeFrame—Use the selectors to configure the time frame for the performance measurement displayed.

Entities —Select the equipment you want to monitor. When you right-click to *Show Performance* with resource(s) selected, those resources appear in this list.

Dashboard View Attributes—Click the arrows between *Available* and *Selected* panels to select monitors for the dashboard. The Available Attributes list shows all the available attributes for that device based on its monitor affiliations. If you select none, a chart appears for each attribute that has data. This is the default. If the user moves some attributes to the *Selected* list then only charts for those attributes appear.

How To:

Create a Custom Dashboard View

The following steps create a custom dashboard view:

1 In the Dashboard Views portlet, select the *New Custom Dashboard* command. An empty default view with twelve components appears.

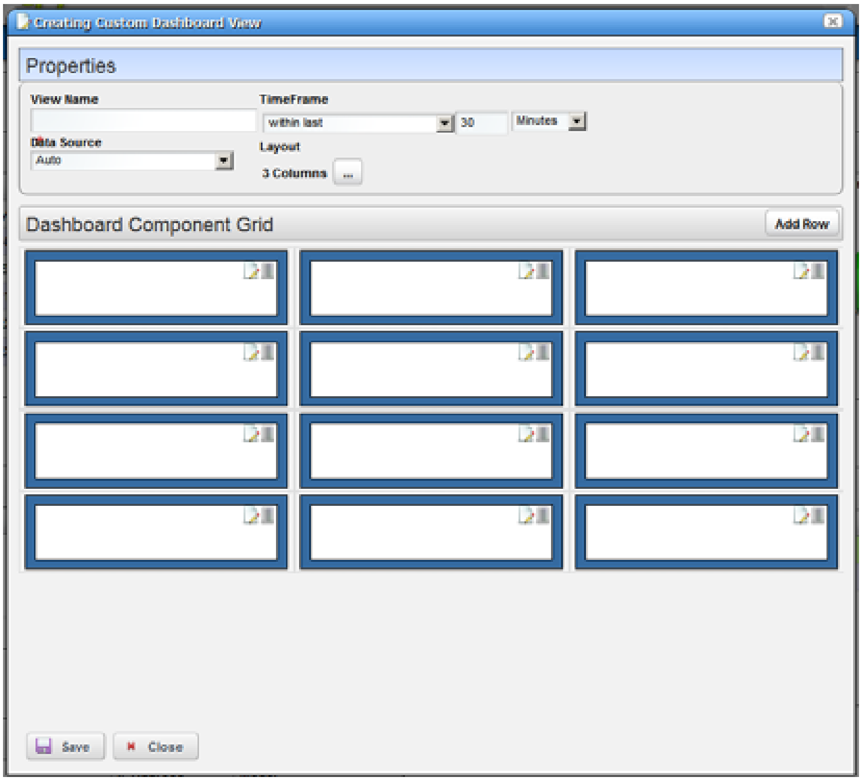

View Name

TimeFrame: within last 30 Minutes

Data Source: Auto

Layout: 3 Columns ...

The Properties panel contains the following controls:

View Name — The name of the dashboard view (Required)

Time Frame — The period over which to display the data. May be either relative (like *last 30 minutes*) or absolute (between specific dates and times). The specified frame applies to all charts in the dashboard.

Data Source —Source for the data. *Current* displays current (raw) data. *Hourly* displays rolled up hourly data. *Daily* displays rolled up daily data. *Auto* (default) determines which data source to use based on the selected time frame.

Layout —Select the desired layout style used to display the dashboard components.

2 To select a layout style, click on the ... button next to the current layout. The layout chooser appears.

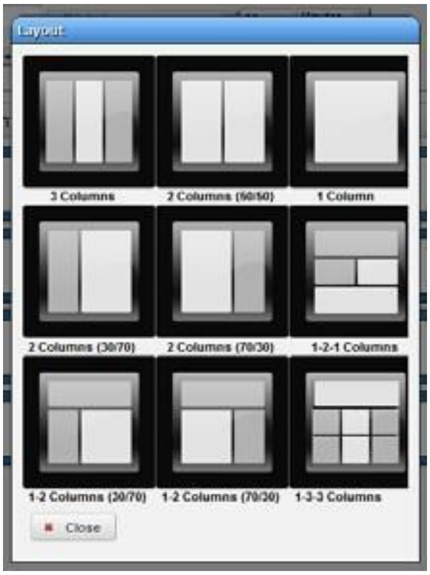

3 Click on the desired layout or click *Close* to keep the current layout. The components displayed to reflect the selected new layout.

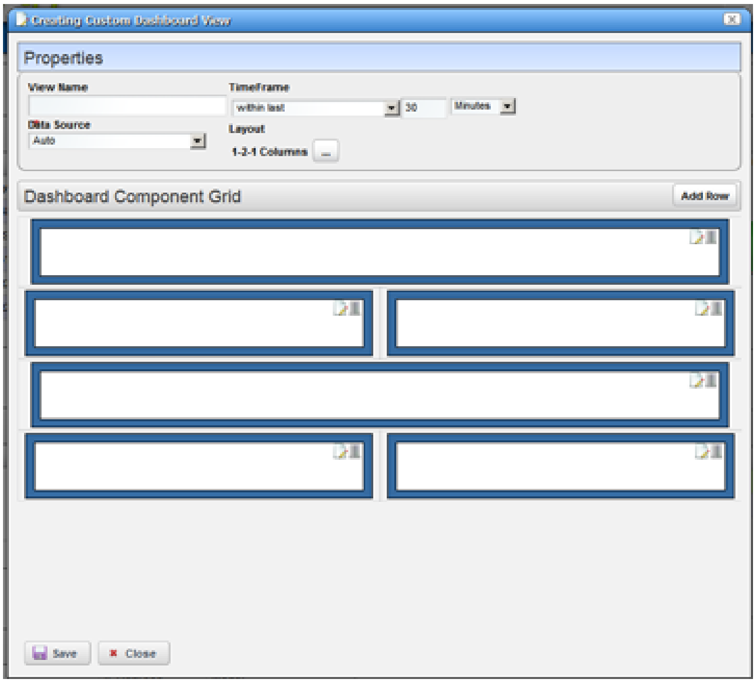

View Name

TimeFrame: within last 30 Minutes

Data Source: Auto

Layout: 1-2-1 Columns ...

If no dashboard components have been configured yet a default configuration appears with three or four rows depending on the dashboard style. If the dashboard components have been configured it will create at least enough rows to display all the configured dashboard components. Add more rows by clicking on the *Add Row* button. An individual dashboard component can be deleted by clicking on the delete button on the component.

Moving Dashboard Components

4 To move a dashboard component to another location, click and drag it over another component. When you release the mouse, the components exchange places.

Configuring Dashboard Components

5 To configure a dashboard component, click the *Edit* button in the upper right corner of the component. The component editor appears.

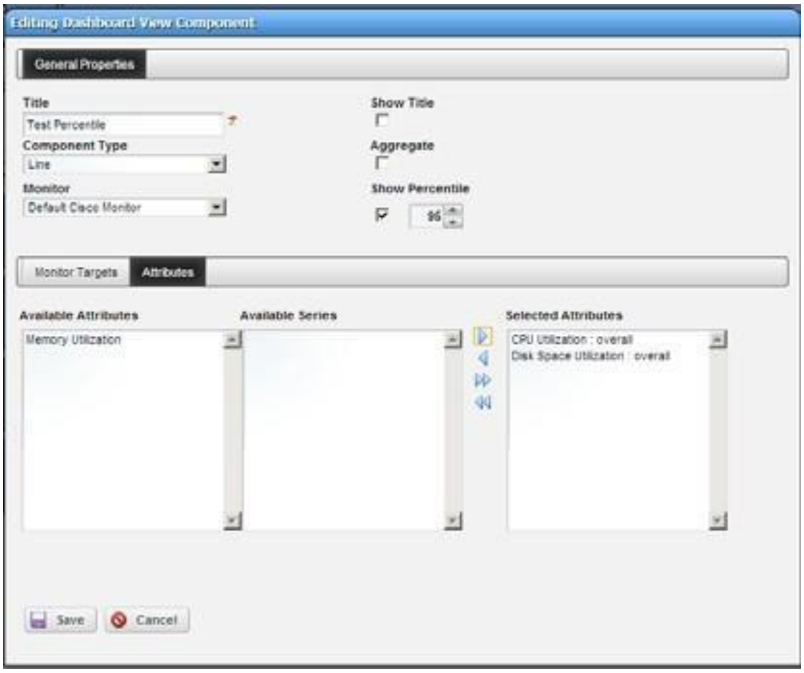

Title: Test Percentile

Component Type: Line

Monitor: Default Cisco Monitor

Show Title: □

Aggregate: □

Show Percentile: ☑ 95

Available Attributes: Memory Utilization

Available Series:

Selected Attributes: CPU Utilization : overall
Disk Space Utilization : overall

The following properties appear in the General Properties section:

Title — Title of this component (required).

NOTE:

Dashboards with a single attribute do not display units, so including them in the title may be informative.

Show Title — Check to display this title above the chart for this component. This overrides the default title that is shown for some charts.

Aggregate — Aggregate the monitored attributes. The aggregate is a sum of all the values for all the entities in the current context.

Show Percentile — Check to enable showing a percentile line for displayed data. Configure the percentile in the field and spinner combination that appears after you check to enable it.

A blue line appears on the dashboard at the selected percentile of the displayed data. If you selected 95%, then of 100 data points, five would be above the line, and 95 below it (so the calculation is more like median than mean). A tooltip for this line displays the calculated value and the selected percentile.

This only supports one line per chart for percentile. If more than one line is on the chart, this computes the percentile based on the first line.

Component Type — Combo Box which specifies what type of component to create. These include the following chart types, *Line, Dial, Bar, Top Talkers* (a line chart showing the top [or bottom] n components for a specific attribute on a specific monitor) *Top Sub-components* (a line chart showing the top [or bottom] n subcomponents belonging to a specific device for a specific attribute. See

Other controls appear depending on the component type selected. These components also have a *Monitor* control, a pick list where you can select from which monitor the charted data originates. See Dial Chart Properties, Top Talkers Properties and Top Subcomponents Properties below for specifics about those.

The line and bar components have two tabs under the general properties section: *Monitor Targets* and *Attributes* . The Monitor Targets section lets you select the devices that are sources of data. You can select any device or attribute in the previously selected monitor. Click the *Add* button displays the monitor target selector.

6 The Attributes tab selects the attribute(s) that appear in the chart. If an attribute is a composite, then its series appears in the Available Series listbox.

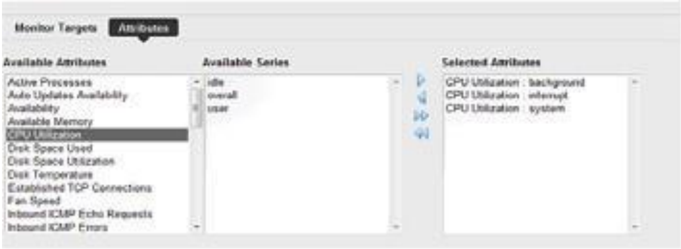

Select the desired series and click the right arrow to move them to the Selected Attributes listbox.

If the attribute is not a composite, then nothing appears in the Available Series listbox. Here, click the right arrow to move the attribute to the *SelectedAttributes* listbox.

You can also elect to *Show Title, Aggregate* (show composite attributes), or *Show Percentile* (display a line at the 95th percentile on the graph) by checking the checkboxes.

Dial Chart Properties

Dial charts have the following additional properties

Monitor — Select which monitor the charted data comes from in the pick list.

Attribute — The attribute to get data for. Note that for aggregate data, each attribute has a minimum (min), maximum (max), and base attribute. If the period is set to Detail, then the aggregate values will always be zero.

Min/Max Value — The minimum/maximum value on the dial.

Entity — The monitor target to get the data for. Clicking on the + button brings up the entity selector.

Top Talkers Properties

Top Talkers components have the following properties.

Monitor — Select which monitor the charted data comes from in the pick list.

Attribute — The attribute to get data for.

NOTE:

Foraggregate data, eachattribute hasaminimum (min), maximum(max), and base attribute.If the period is set to Detail, then the aggregate values will always be zero.

Max # of Entities— The number of entities to display

Order — Select either *Ascending* (Bottom n) or *Descending* (Top n).

Top Subcomponents Properties

Top Subcomponents components have the following properties.

Entity — The parent entity for the found subcomponents. Clicking on the + button brings up the entity selector.

Attribute — The attribute to get data for.

Max # of Entities— The number of entities to display

Order — Select either *Ascending* (Bottom n) or *Descending* (Top n).

Convert Simple Dashboards to Custom Dashboards

To convert a simple dashboard to a custom dashboard use the *Convert* command on the *Dashboard Views* menu. You cannot convert custom dashboards to simple dashboards.

How To:

View Historical Dashboard Data:

1. Click on the clock icon in the dashboard header.
2. Change the timeframe to the period on the day for which you want to view data.
3. Click on the green checkmark.
4. The data appears for the selected time frame.

On simple dashboards the type of data (raw, hourly, daily) depends on the timeframe selected. For less than three hours, raw data appears. For three hours to less than three days, hourly data appears. For three days or more daily data appears.

Do not forget you may need to change the data retention policy. You may not see raw data if you are viewing a timeframe from two days ago because the default retention policy only keeps rawdata for one day.

Show Performance Templates

By default, the Show Performance command displays data for the first twelve attributes it finds. You can control which attributes appear when you select Show Performance by creating a performance template. Aperformance template lets you set dashboard parameters and associate them to one or more device models. Then, when you execute Show Performance on a device of that type, those dashboard parameters display the dashboard for that device.

How To:

Create A Performance Template

To create a performance template, follow these steps:

1 Right click in the Dashboard Views portlet and click on the *Performance Templates* menu item.

2 The Performance Templates manager appears.

3 To create a new performance template, click on the Add button. The Performance Template Editor appears.

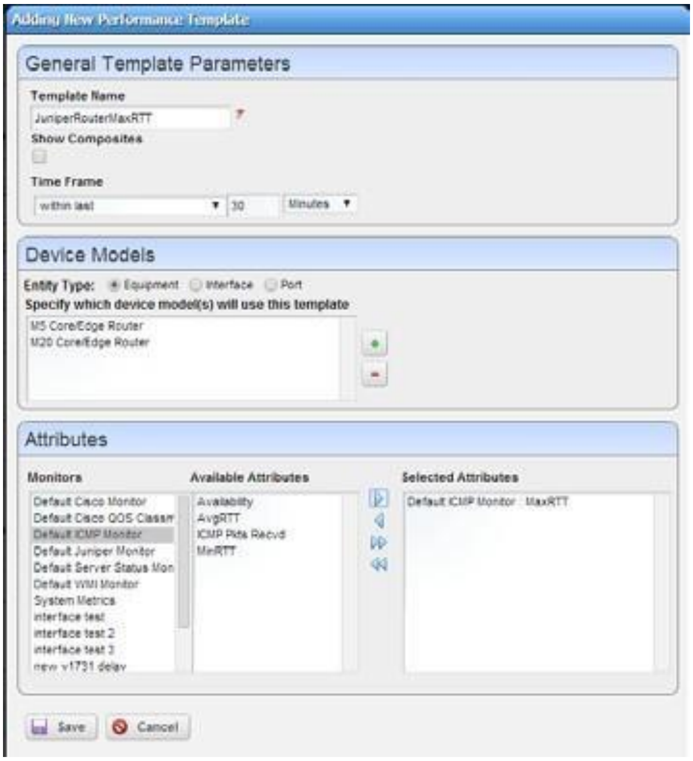

Adding New Performance Template

General Template Parameters

Template Name: JuniperRouterMaxRTT

Show Composites: □

Time Frame: within last ▼ 30 Minutes ▼

Device Models

Entity Type: ● Equipment ○ Interface ○ Port

Specify which device model(s) will use this template

M5 Core/Edge Router
M20 Core/Edge Router

Attributes

Monitors:
Default Cisco Monitor
Default Cisco QOS Classm
Default ICMP Monitor
Default Juniper Monitor
Default Server Status Mon
Default WMI Monitor
System Metrics
interface test
interface test 2
interface test 3
new v1731 delay

Available Attributes:
Availability
AvgRTT
ICMP Pkts Recvd
MinRTT

Selected Attributes:
Default ICMP Monitor : MaxRTT

For this example, we have selected *Entity Type: Equipment* with the radio buttons below *Device Models*. See *Dashboard Templates for Interface and Port Equipment* on page 428

4 Name your template. The Show Composites and Time Frame fields are the same as in the dashboard (see Dashboard Editor on page 420).

5 To specify which device model(s) this template will apply to, click on the + button in the Device Models panel. The model selector appears.

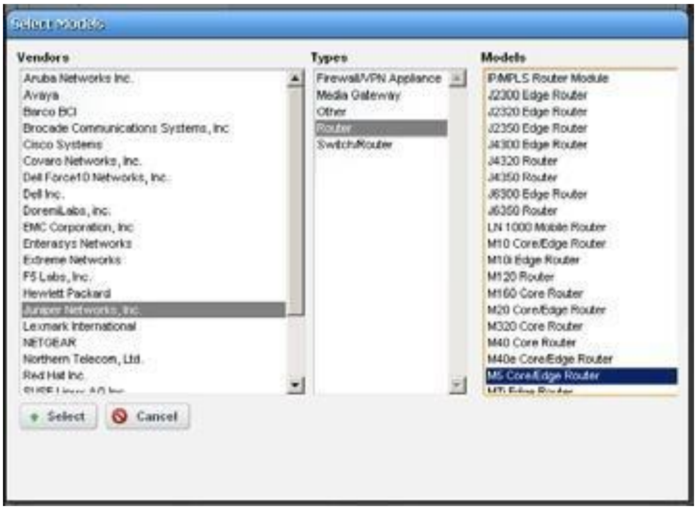

Select multiple devices by clicking + repeatedly, selecting a single device each time. Youcan also make several templates for each device. See Multiple Performance Templates on page 428 for the way that works.

6 Click on a vendor to see the device types for that vendor. Then click on a device type to see the models available for that vendor and device type. Select the model you want and click on the select button.

7 To select the attributes that you want to appear by default in a performance dashboard for the selected device, click on a monitor to see the attributes available for that monitor. Click on the right arrow button to move the selected attributes from *Available* to *Selected*. Those are the attributes that will appear by default in dashboards for the selected device.

8 When you have selected all the parameters you want, click *Save*. It then appears in the template list.

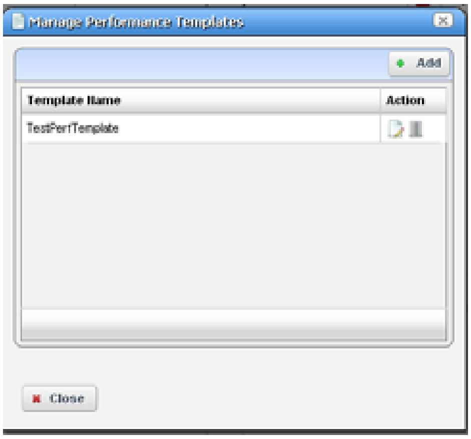

Template Name	Action
TestPerfTemplate	

To edit or delete your template, use the buttons in the action column of the table.

Now when you click on show performance, Cruz checks whether a template for that device type exists. If one exists, then that template guides what appears in the performance view for the device.

Multiple Performance Templates

The template name appears in the upper right corner of dashboards that appear when you select Show Performance.

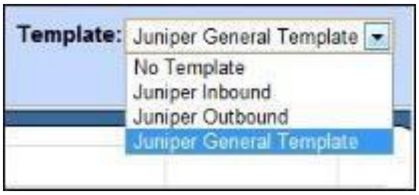

Template: Juniper General Template

If other templates for that device type exist, they also appear in a template pick list in the upper right corner. You can pick another template to display its attribute selection. The *No Template* selection displays the default dozen attributes that would appear if you selected Show Performance without a template defined for the device.

Dashboard Templates for Interface and Port Equipment

Dashboard template types include Interface and Port templates. When configuring a dashboard template in the editor (see How to: Create A Performance Template on page 426), select the *Interface* or *Port* radio button in the middle panel of the editor, then add the appropriate port or interface in that same middle panel.

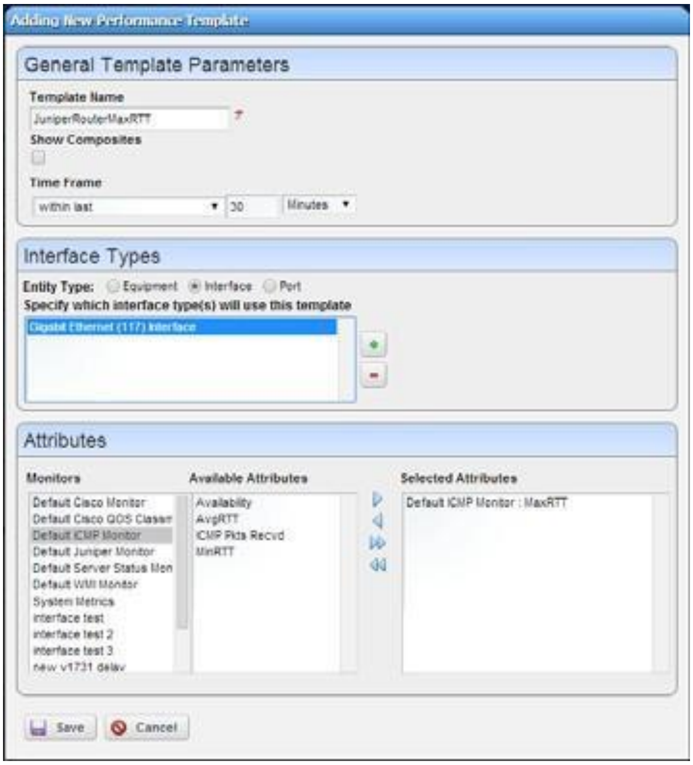

Adding New Performance Template

General Template Parameters

Template Name: JuniperRouterMaxRTT *

Show Composites: □

Time Frame: within last 30 Minutes

Interface Types

Entity Type: ○ Equipment ● Interface ○ Port

Specify which interface type(s) will use this template: Gigabit Ethernet (117) Interface

Attributes

Selected Attributes: Default ICMP Monitor : MaxRTT

After selecting the port or interface types for the template, you can select the Attributes to monitor as you would for other equipment.

Key Metric Editor

When you select *Performance > Show Key Metrics*,this editor appears for devices that have such metrics. It displays the available Metrics, and a Chart panel where you can configure their display.

Metrics

This panel’s display depends on the selected device.

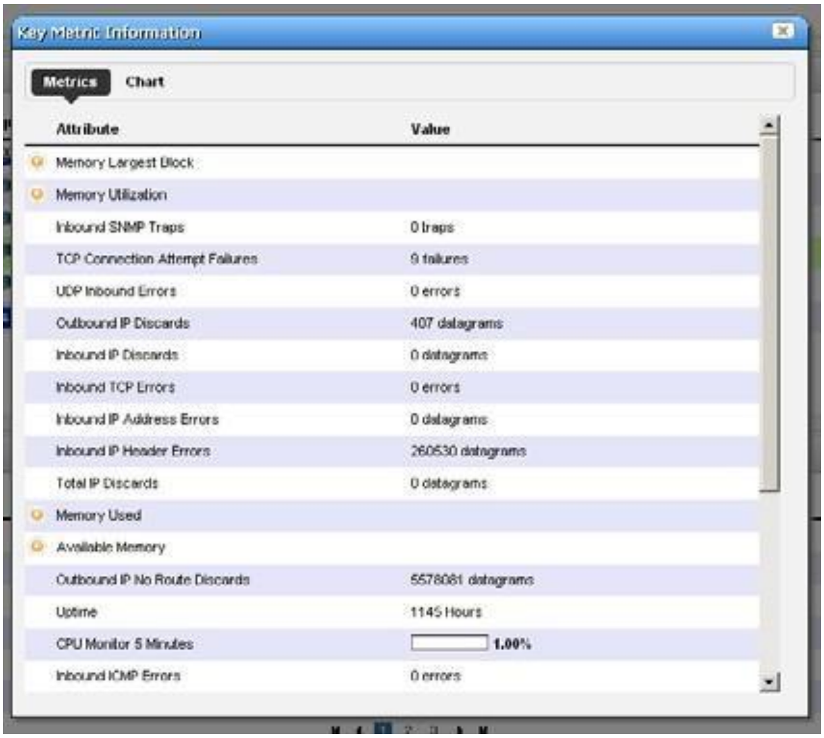

Attribute	Value
Memory Largest Block	
Memory Utilization	
Inbound SNMP Traps	0 traps
TCP Connection Attempt Failures	9 failures
UDP Inbound Errors	0 errors
Outbound IP Discards	407 datagrams
Inbound IP Discards	0 datagrams
Inbound TCP Errors	0 errors
Inbound IP Address Errors	0 datagrams
Inbound IP Header Errors	260530 datagrams
Total IP Discards	0 datagrams
Memory Used	
Available Memory	
Outbound IP No Route Discards	5578081 datagrams
Uptime	1145 Hours
CPU Monitor 5 Minutes	1.00%
Inbound ICMP Errors	0 errors

Chart

Click *Chart* to first select up to three metrics you want to graph, and the polling interval for the graph.

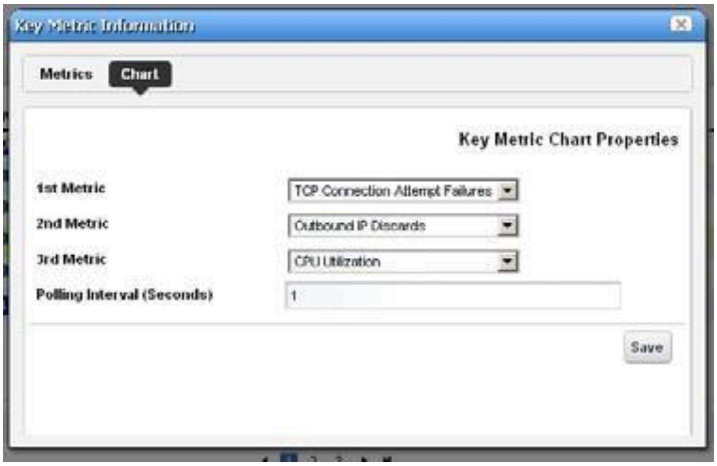

Key Metric Chart Properties

1st Metric	TCP Connection Attempt Failures
2nd Metric	Outbound IP Discards
3rd Metric	CPU Utilization
Polling Interval (Seconds)	1

Then click *Save*, and the graph appears.

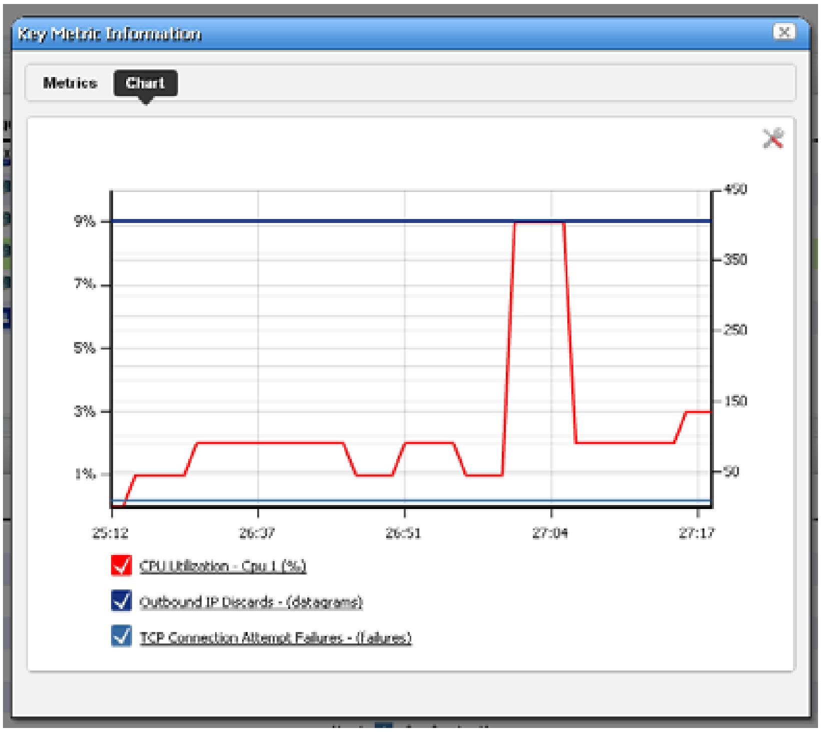
Key Metric Information
Metrics
Chart
9%
7%
5%
3%
1%
450
350
250
150
50
25:12
26:37
26:51
27:04
27:17
☑ CPU Utilization - Cpu 1 (%)
☑ Outbound IP Discards - (datagrams)
☑ TCP Connection Attempt Failures - (failures)

Click the screwdriver/wrench icon in the upper right corner to return to the chart configuration screen.

Best Practices: Performance and Monitors

Monitoring can impact system performance. Monitors with many targets, many attributes per target and frequent polling intervals are most likely to slow system performance.

△ CAUTION:

Limit monitortargetsto 10,000orlesspermonitor(distributethemifyouhave morethan 10,000). The following suggests other ways to improve monitor performance, too.

The following are the primary considerations when configuring monitors to get the desired performance from your system:

Database Insertion Rate — How many rows can your hardware realistically insert per second?

Everysystem has a maximum data insertion rate. Thisrate depends on the system's hardware configuration. Astandard 7200 RPM disk can typically manage 300 insertions per second per disk. 10000-15000 RPM disk can have as many as 600 insertions per second per disk. Your experience may vary depending on your drive's controller and configuration.

The sum of all monitors' insertion rate should not exceed the system's maximum data insertion rate. To calculate the insertion rate of a monitor, apply the following formula:

<# of monitor targets > x <# of retained attributes > / < polling interval in seconds > = inserts /second.

So a monitor with 100 targets, retaining 10 attributes once a minute would have an insertion rate of 17 rows per second (100 * 10 / 60 = 16.67 inserts per second).

Example:

Monitor A: 1000 targets * 10 retained attributes each / 120 = 83 insert per second

Monitor B: 100 targets * 25 retained attributes each / 600 = 4 insert per second

Monitor C: 10000 targets * 10 retained attributes each / 60 = 1667 insert per second

System total insertion rate (A + B + C) : 1754 insertions per second

Thisconfiguration would be too aggressive for a system with a 7200 RPM disk since it dramatically exceeds the 300 insertions per second that the disk can support.

The following alternatives could resolve this:

- Upgrade to disk hardware that can keep up with the insertion rate. If the target insertion rate is 1754 inserts per second, add a disk to the array. If 1754 inserts / 300 insertions per second on a 7200 rpm disk amounts to 5.84 disks, use 6 disks (or more). If using 15000 RPM disks at 600 inserts per second, 1754/600 means you need 2.92 (3 discs) minimum.
- Modify the monitors to achieve a lower insertion rate. If you only have one 7200 RPM disk, it can only support 300 insertions per second. Youhave the option of lowering the target count, reducing what is retained or lengthening the polling interval.

  So, from the example above if we changed the polling interval from once a minute to once every 10 minutes, Monitor C's insertion rate would drop to 167 inserts per second. The overall system would then only have an insert rate of 254 per second well below the hardware's limitation.

NOTE:

Traffic flow analysis can process and retain even larger amounts of information. Flows that correlate 50%, polled everyminute foradayrequire roughly 109G of database, and require 4500 insertions per second.

△ CAUTION:

Thesenumbers and sample calculations represent bestcasescenarios. Any disk or disk arraytypically serves other applications and processes besides monitors. Make sure to take account of that when calculating how to accommodate your monitors. The system admin or system user should assist in making that assessment.

Storage Requirements (Database Size) — How much disk space do you need, based on your retention policy? See Retention Policies on page 366 for more about configuring those.

Cruz stores performance data in three different forms Detail, Hourly and Daily data. It collects Detail metrics directly from the device or calculates these from the collected data with each poll. Hourly data summarizes the detail data collected during the hour. Daily data summarizes the hourly data collected during that day. The retention policy associated to the monitor describes how long Cruz retains each of these data types within the system.

Cruz stores a performance metric as a single row within a database table. Each row in that database consumes roughly 150 bytes of disk space. The sum of each monitor's disk space required determines the amount of disk space. For each monitor add the disk space required for the Detail, Hourly and Daily da ta using the followingformula:

< Detail disk space > + <Hourly disk space > + <Daily disk space > = Monitor disk space in bytes

where...

<# of metrics retained per poll > = <# of monitor targets > x < # of retained attributes >.

<# of metrics retained per day> = <# of metrics retained per poll> x <# of polls per day>.

< Detail disk space> = <# of metrics retained per day> x <# of days to retain Detail data > x 150.

< Hourly disk space > = <# of metrics retained per poll > x <# of days to retain Hourly data > x 24 x 150.

< Daily disk space > = <# of metrics retained pe r poll> x <# of days to retain Hourly data> x 150.

If the system does not have sufficient disk space consider the following options:

- Add more hardware to increase the available disk space.
- Reduce the retention period of one or more monitors to lower the overall disk space requirements. Of the three data forms Detail data will consume the most disk space per day of retention.

Table size — Based on your monitor configuration how large will database tables get? Each monitor has a series of dedicated performance tables that store the Detail, Hourly and Daily performance metrics. The number of tables depends on the retention policy associated with the monitor.

A single table stores the monitor's detail data for a 24-hour period. Detail data are individual performance metrics collected and/or calculated during each poll. After that initial 24 hours, Cruz creates a new table to store the next 24 hours' of detail data and so on.

Because of the resulting table size, the number of performance metrics generated by a single monitor in a 24-hour period impacts performance. Best practice is to configure each monitor to produce less than 10 million rows per day. When monitors exceed that number, noticeable delaysresult when retrieving performance data. Todetermine the number of metrics retained by monitors per day please refer to < # of metrics retained per day > calculation from the previous section.

△ CAUTION:
These numbers depend entirely upon the system hardware, available memory and processor speed.

If a monitor does exceed the target maximum rows per day, consider the following options singly or in combination to change that:

- Reduce the number of retained attributes per poll.
- Reduce the polling frequency.
- Reduce the number of monitor targets per monitor. Notice that you can still have the same number of targets if you split the targets among multiple monitors.

Finally, tune your database for the expected load. Refer to the *Cruz Installation Guide* for MySQL sizing recommendations.

Dashboard Performance Limits

Creating dashboards makes performance demands on your system. If you make too many, or monitor too many attributes within your dashboards, system response times can suffer. Performance can also suffer because you have too many dashboard portlets on a single page.

To work around these limitations, add another page (see Portal Overview on page 122 for details), then move some dashboard portlets from the over-populated page to the new one. You can also split monitored attributes between different dashboards.

Monitoring from a Cloud Server

The following outlines hardware sizing for performance monitoring from a cloud server:

RAM	Max Targets (5-minute poll intervals)	Heap Memory Settings	Recommended CPU Cores and Disk Space
16 GB RAM	10000	4 GB Synergy Web Server heap, 6 GB Application Server heap, 2GB Database buffer	4 + core, 100 GB + disk space
32 GB RAM	25000	6 GB Synergy Web Server heap, 12 GB Application Server heap, 8 GB Database buffer	4 + core, 200 GB + disk space
64 GB RAM	50000	8 GB Synergy Web Server heap, 16 GB Application Server heap, 24 GB Database buffer	8 + core, 400 GB + disk space
128 GB RAM (recommended)	50000	12 GB Synergy Web Server heap, 24 GB Application Server heap, 48 GB Database buffer	16 + core, 400 GB + disk space

Note:

- The VPN tunnel is assumed to be available to cloud services.
- The suggested target numbers are for latency under 50ms. You may have to scale down numbers of targets when latency is greater than 50ms.

monitorTargetDown Event Interval

By default, Cruz only creates the monitorTargetDown event every 30 minutes, even if the polling interval is shorter. This behavior is configurable. Thereare properties through which you can either tell the system to always create these so-called availability events every time it finds the target unreachable.

Youcan alsochange the time threshold forwhich it willcreate another monitorTargetDownevent if you want to have these created more or less often than once every 30 minutes in circumstances where a polling target remains unreachable for an extended period of time.

To configure this interval, modify (or bett er, override) the following properties (from `.../owareapps/performance/lib/pm.properties`)

```
#Monitor alarm settings
#
#pm.monitor.AlwaysReemitAlarm=true
#pm.monitor.AlarmReemitTimeout=15
```

Notice that these are commented out by default, so monitorTargetDown appears by 30-minute intervals, and thedefault for `AlwaysReemitAlarm` is `false`. Restartapplicationserverforedits to take effect. In a cluster, all edits must be consistent.

8

Configuration Management

This section provides tasks and information related to perform configuration management tasks, such as managing your file servers and files. If you already have a good understanding of the portlets and editors, go directly to the tasks you want to perform.

Configuration Management Portlets and Editors

This section describes the following configuration management portlets and editors:

- Configuration Alarms
- Configuration Management Schedule
- Configuration Files
- Image Repository
- File Servers

Configuration Alarms

Use the Configuration Alarms portlet to viewalarms generated by a configuration action, such as backup, restore, or deploy and filter out those alarms not related to configuration actions. You can expand or change to event history to show all events.

This portlet is intended for users who are interested in managing the configuration of their discovered devices.

Access this portlet by selecting the Configuration Management page. This portlet has both a summary view and an expanded view. Each viewdisplays different columns and have the same pop-up menu options available.

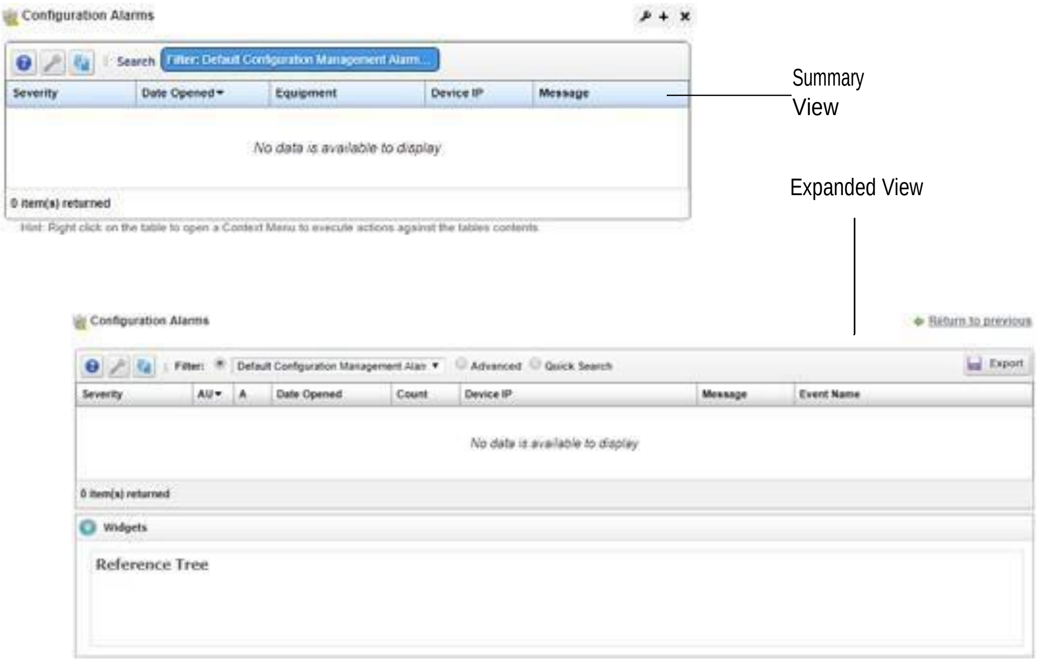

No data is available to display

0 item(s) returned

Hint: Right click on the table to open a Context Menu to execute actions against the tables contents

No data is available to display

0 item(s) returned

See Alarms in Topologies on page 246 to understand how alarms appear in the topology portlet.

Columns

Other than the general navigation and configuration options, the Configuration Alarms portlets (summaryand expanded views) include the followingcolumns. The columns displayed bydefault are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Severity	The alarm severity indicated by the color and text. The severity only has meaning for alarms and security alarms. Informational alarms have a severity level of Indeterminate. Closed alarms show without color. This field displays on the summary and expanded views by default.
Assigned User (AU)	The user currently assigned to this alarm. You can filter your assigned alarms using the Advanced filter option and selecting Assigned by User, contains, and entering your user ID. This field displays on the expanded view by default.
Acknowledged (A)	An indicator whether the alarm is acknowledged (checkmark) or not (X).
Date Opened	The date the alarm was created. This field displays on the summary and expanded views by default.
Count	The number of instances for this alarm. Multiples of the same alarm show as a single row but this value increments as more instances occur. This field displays on the expanded view by default.
Equipment	The entity emitting the alarm. This field displays on the summary view by default.
Device IP	The equipment IP address where the alarm appeared. This field displays on the summary and expanded views by default.
Message	The message associated with the alarm. This field displays on the summary and expanded views by default.
Event Name	The event associated with the alarm. This field displays on the expanded view by default.
Widgets	Additional information about the selected alarm, such as: • Reference Tree shows the connection between the alarm and its resource. • Alarm Details shows the severity, message, date opened, and so on. • MIB Details shows notification OID and MIB text. • Total Occurrences by Date shows a graph of the total occurrences by date. The Widgets field is available **only** from the expanded view.
The remaining field descriptions are available for use but do not appear on the default Configuration Alarms portlets.	
Date Cleared	The date and time the alarm was closed.
Update Date/Time	The time stamp for when the alarm was updated.
Date Assigned	The date and time that the alarm was assigned.
Assigned By	The user ID for the person that assigned this alert.
Ack Time	The time that the alarm was acknowledged.
Alarm State	The alarm’s open or closed state.
Cleared By	The user ID for the person that cleared the alarm.
MIB Text	The alarm’s MIB text.
Ack By	The user ID for the person that acknowledged the alarm.
Correlated By	The user ID for the person that cleared the alarm.
Correlated Time	The date and time when the alarm correlation to a parent alarm (caused by or blocked by).

Column	Description
Correlation State	The role the alarm plays in any correlation, such as top-level alarm, caused by parent, blocked by parent.
Domain	The domain emitting the alarm.
Entity Name	The entity emitting this alarm.
Entity Type	The type of monitored entity.
Has Children	An indicator that shows whether the alarm has children (checkmark) or not (X).
Highest Severity	The highest severity assigned to the selected alarm.
Include child alarms	An indicator that shows whether the alarm includes child alarms (checkmark) or not (X).
Location	The alarm's location.
Notification Type OID	The identifier for the notification displayed as an alarm.
Original Severity	The original severity specified for the selected alarm.
Parent Alarm	The name of the parent alarm.
Region	The alarm's region.
Resource Propagation	The propagation for this alarm, if any.
Service Affecting	Indicates whether the alarm is on equipment in a provisioned service (checkmark) or not (X). A service affecting alarm propagate to show as a component of service- and customer-related alarms. Service-Affecting alarms are of indeterminate or greater severity.
Service Propagation	The alarm propagation if Service Affecting is true in the event definition. This is similar to Resource Propagation, but controls **only** whether alarms affect services associated with entities hierarchically related to the alarmed entity. For example, if a port alarm is created and its event definition specifies that the Service Propagation is Impacts Subcomponents, then the alarm propagates only to the services associated with this port's interfaces. This does not affect the services associated with the top-level device.
Suppressed	An indicator whether the alarm is suppressed (checkmark) or not (X).
Suppression Date	The date and time of the alarm's suppression, if applicable.
Suppressor	The alarm that suppresses the selected alarm.

Pop-Up Menu

The Configuration Alarmspop-up menu providesaccessto the followingoptions. Right-clickarow to access these options.

Menu Option	Description
Edit	Opens the Edit Alarm window (see Alarm Editor on page 284) or the Editing Event Definition window (see Event Definition Editor on page 326).
Details (Click +Shi ft)	Opens a Details screen for the alarm itself, not the entity emitting it. (see Connected Devices on page 189 for an example of this type of screen). This contains information like the MIB text, any Event Processing Rules invoked, and a Reference Tree for the alarm. It also lets you configure alarm correlation. See Parent/Child Alarm Correlation: Alarm Details Panel on page 282. Select Alarm Details or Equipment Details

Menu Option	Description
Topology	Displays a topology map that includes the selected alarms. See Presentation Capabilities for more about these maps.
Acknowledge Alarm	Acknowledges the selected alarms. The current date and time appear in the Ack Time field. The checkmark appears in the expanded portlet for acknowledged alarms. This option is available when the selected alarms **are not** acknowledged.
Unacknowledge Alarm	Unacknowledges previously acknowledged alarms and clears the entries in the Ack By and Ack Time fields. When an alarm is not acknowledged, the red icon appears in the expanded portlet. This option is available when the selected alarms **are** acknowledged.
Assign User	Assigns selected alarms to one of the users displayed in the sub-menu by selecting that user. An icon also appears in the expanded portlet indicating the alarm is assigned to someone.
Clear Alarm	Removes the selected alarms from the default alarm view and marks them as a candidates for the database archiving process (DAP). Clearing alarms is an indication to the system that the alarms are resolved/addressed. If your system has propagation policies enabled, clearing recalculates dependent alarms.
Clear Group of Alarms	Removes a group of alarms from the default alarm view and marks them as candidates for DAP. This option is useful when you have many unimportant alarms that you want to clear at the same time instead of individually. For example, perhaps you want to clear all alarms that are informational and are more than a week old. Before selecting this menu item, create a filter for the group of alarms that you want to clear. When you select Clear Group of Alarms, a list of defined filters appears. Select the appropriate filter and then click Execute. All open alarms that meet the filter criteria are cleared. **Caution:** The Clear Group of Alarms action is irreversible.
Direct Access	Open an SNMP Mib Browser to the alarmed device, a CLI Terminal (Telnet window) to the alarmed device, or ICMP Ping the device alarmed. Only those available appear in the subsequent menu.
Email Alarm	Opens the Email Alarm window, where you specify a subject and e-mail address to which you want to mail the alarm’s content, and then click Send Email. Click Cancel to end this operation without sending e-mail. SMTP setup is required to e-mail an alarm. See SMTP Configuration on page 63 for instructions about setting up e-mail from Cruz.
Show Performance	Shows performance data for the selected interfaces in the Performance Dashboard window. Click the edit tool to modify the dashboard view properties, entities, or attributes. See Dashboard Editor on page 420 for more about the editor and its options.
Aging Policy	This lets you select a policy that determines how long the selected alarms remain in the database. See Implementing DAP on page 80 for information about configuring such policies.
Edit Custom Attributes	Opens the Custom Attribute Editor where you define field characteristics, such as whether it is enabled, the label name, and the tooltip.
View as PDF	Creates an Acrobat PDF document containing the selected alarms’ content displayed in the portlet.
Share with User	Opens the Share with User window where you select the colleague you want to share the selected alarms with and then type your message.

Configuration Management Schedule

Use the Configuration Management Schedule portlet to quickly locate schedules pertaining to managed device configurations and to create and maintain schedules for the followingaction types:

- Config File Backup
- Config File Restore
- OS Image Deploy

The seeded Default Change Determination process backups a device if a configuration file change is detected. This runs against all devices when turned on.

This portlet is intended for users who are interested in managing the configuration of their discovered devices.

Access this portlet by selecting the Configuration Management page. This portlet has both a summary view and an expanded view. Each view displays different columns and could have the same pop-up menu options available.

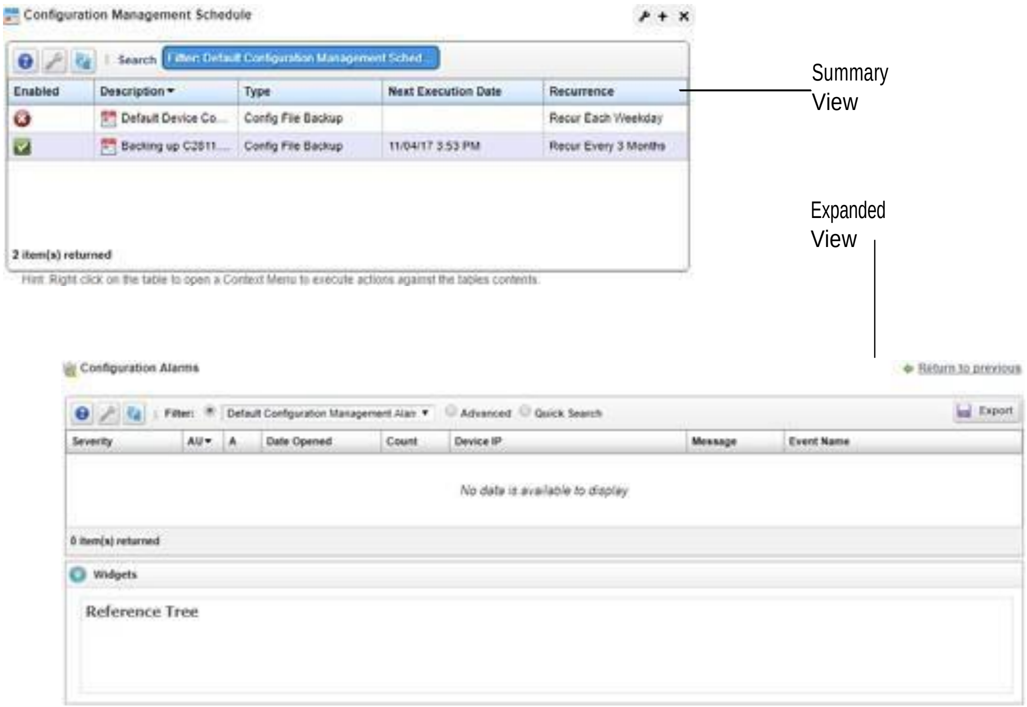

Configuration Management Schedule

Enabled	Description	Type	Next Execution Date	Recurrence
☒	Default Device Co...	Config File Backup		Recur Each Weekday
☑	Backing up C2811....	Config File Backup	11/04/17 3:53 PM	Recur Every 3 Months

2 item(s) returned

Hint: Right click on the table to open a Context Menu to execute actions against the tables contents.

Summary View

Expanded View

Configuration Alarms

Return to previous
Filter: ● Default Configuration Management Alan ▾ ○ Advanced ○ Quick Search
Export

Severity	AU	A	Date Opened	Count	Device IP	Message	Event Name

No data is available to display

0 item(s) returned

Reference Tree

Columns

Other than the general navigation and configuration options, the Configuration Management Schedule portlets (summaryandexpanded views) includesthefollowingcolumns. Thecolumns displayed by default are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Enabled	An indicator that shows whether the schedule is enabled (checkmark) or not (X). This field displays on the summary and expanded views by default.
Description	The schedule description. This field displays on the summary and expanded views by default.
Type	The type of action scheduled, such as Config File Backup, Config File Restore, or OS Image Deploy. This field displays on the summary and expanded views by default.
Next Execution Date Recurrence	The frequency in which to execute the scheduled action, such as each weekday, every three months, only at startup, only once, and so on. This field displays on the summary and expanded views by default.
Submission Date	The date and time that this schedule was submitted. This field displays on the expanded view bydefault.
Start Date	The date and time to start the scheduled action execution. This field displays on the expanded view by default.
End Date	The date and time to end the scheduled action execution or the number of occurrences before ending this scheduled action. This field displays on the expanded view by default.
Finished	An indicator that shows whether the configuration execution completed (checkmark) or not (X). This field displays on the expanded view by default.
Execution Count	The number of times this scheduled action executed. Useful if you specified and end date based on the number of occurrences. This field displays on the expanded view by default.
Last Execution Date	The date and time that the action was executed. This field displays on the expanded view by default.
Domain ID	The Multitenant domain ID emitting the alarm.
Run Status	Indicated the schedules current state, such as waiting. Available to add to the summary/expanded view

Pop-Up Menu

The Configuration Management Schedule pop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
New	Provides access to the following schedule configuration types: • Config File Backup • Config File Restore • OS Image Deploy Once you select the configuration type, the New Schedule (*configType*) window opens, where you configure actions, their targets, and the order in which they execute. You can also create a schedule from a portlet that ordinarily executes them using the Schedule option. For example, Discovering Resources on page 211.
Edit	Opens the Editing Schedule (*configType*) window, where you modify the selected schedule's settings. This option appears for activity-based configurations. You can also edit a schedule from a portlet that ordinarily executes them using the Schedule option. For example, Discovering Resources on page 211.
Delete	Deletes the selected schedule configuration. A confirmation message is displayed. Click Delete to continue with the deletion. Click Cancel to disregard the request.
Enable Schedule	Activates the selected schedule. This option appears if the selected schedule is disabled. You can also enable a schedule using the Edit option and selecting the Enabled option.
Disable Schedule	Deactivates the selected schedule. This option appears if the selected schedule is enabled.
Execute	Runs the scheduled action. If the scheduled action is activity-based or discovery-profile-based, an audit viewer appears progress of the selected item. For other scheduled action types, a message is displayed sating that the scheduled items were sent to the application server for immediate execution. Use the Audit Trail portlet to monitor progress (see Audit Trail Portlet on page 150) for details.
Share with User	Shares selected alarms with other colleagues on your Cruz and consults with them using the conferencing feature.

Config File Backup

Use the New Schedule (Config File Backup) window to specify the backup options, targets for backup, device options, and then schedule the backup.

Access this window by right clicking a configuration and then selecting New > Config File Backup.

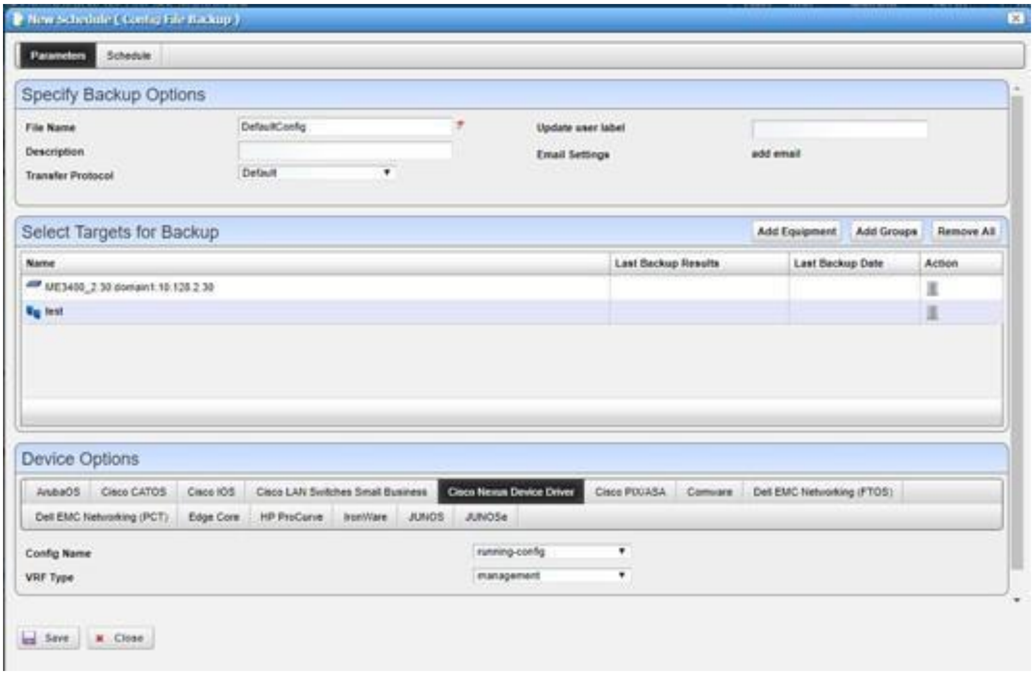

File Name: DefaultConfig

Description:

Transfer Protocol: Default

Update user label:

Email Settings: add email

Name	Last Backup Results	Last Backup Date	Action
ME3400_2.30 domain1:10.128.2.30			
test			

Config Name: running-config

VRF Type: management

The Config File Backup Parameters window has the following fields and options. See Schedules Portlet on page 155 for scheduling fields and options descriptions.

Field/Option	Description
File Name	This name is used as the backed-up configuration file during server transfers and saved in source control. The backed-up configuration file will be name this during the server transfers and saved in the source control. Due to limitations of some devices, do not exceed 15 characters for the filename.
Description	Provides a description of th e file as it’s checked into storage.
Transfer Protocol	Specifies which protocol is used (Default, FTP, TFTP, or SFTP) to transfer files.
Update user label	Creates or points the selected user label to the backed up configuration file.
Email Settings	Opens a window where you specify email options and recipients.
Add Equipment	Adds the selected equipment to the Targets for Backup list.
Add Groups	Adds groups of devices to the Targets for Backup list.
Remove All	Deletes all devices and groups from the Targets for Backup list.
Targets for Backup	Provides the following information for each target: • Name • Last Backup Results • Last Backup Date • Action deletes the entry

Field/Option	Description
Device Options	Lists the tabs for each device based on your target backup list. Click a device tab to view its options, such as Config Name, VRF Type, Config File, Configuration Type, on so on.

Config File Restore

Use the Config File Restore window to specify the targets for restore, device options for the selected targets, and schedule the restore.

Access this window by right clicking a configuration and then selecting New > Config File Restore.

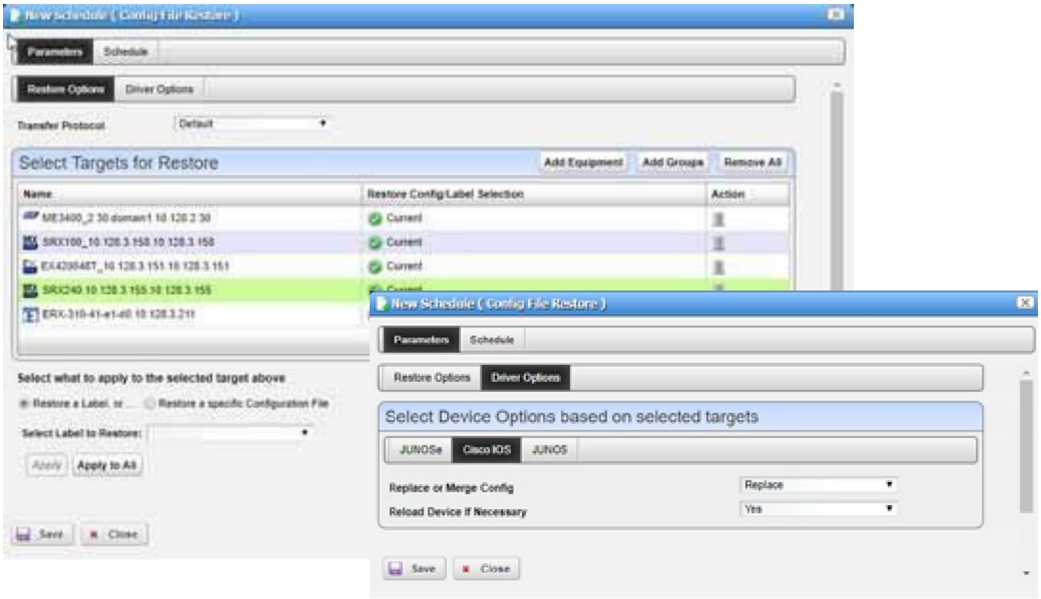

New Schedule (Config File Restore)

Parameters | Schedule

Restore Options | Driver Options

Transfer Protocol: Default

Select Targets for Restore

Name	Restore Config/Label Selection	Action
ME3400_2 30 domain1 10.128.2.30	Current	
SRX100_10.128.3.158 10.128.3.158	Current	
EX4200-48T_10.128.3.151 10.128.3.151	Current	
SRX240 10.128.3.155 10.128.3.155	Current	
ERX-310-41-e1-d0 10.128.3.211		

Select what to apply to the selected target above

● Restore a Label, or ... ○ Restore a specific Configuration File

Select Label to Restore:

New Schedule (Config File Restore)

Parameters | Schedule

Restore Options | Driver Options

Select Device Options based on selected targets

JUNOSe | Cisco IOS | JUNOS

Replace or Merge Config: Replace

Reload Device If Necessary: Yes

The Config File Backup Parameters panes has the following restore and driver fields and options. See Schedules Portlet on page 155 for scheduling fields and options descriptions.

Field/Option	Description
Transfer Protocol	Specifies which protocol is used (Default, FTP, TFTP, or SFTP) to transfer files.
Add Equipment	Adds the selected equipment to the Targets for Restore list.
Add Groups	Adds groups of devices to the Targets for Restore list.
Remove All	Deletes all devices and groups from the Targets for Restore list.
Targets for Restore	Provides the following information for each target: • Name • Restore Config/Label Selection • Action deletes the entry When you select to target to restore, the following options appear: • Restore a Label • Restore a specific Configuration File. If you select Restore a Label, select a label to restore. If you select Restore a specific Configuration File, select the configuration file from the list.

Field/Option	Description
Device Options	Lists the tabs for each device based on the restore target selected. Click a device tab to view its options, such as Replace or Merge Config, Reload Device If Necessary, Load Action, on so on. Select the appropriate action for each option.

OS Image Deploy

Use the OS Image Deploy window to specify firmware image, targets for deployment, device options, and then schedule the OS image deployment.

Access this window by right clicking a configuration and then selecting New > OS Image Deploy.

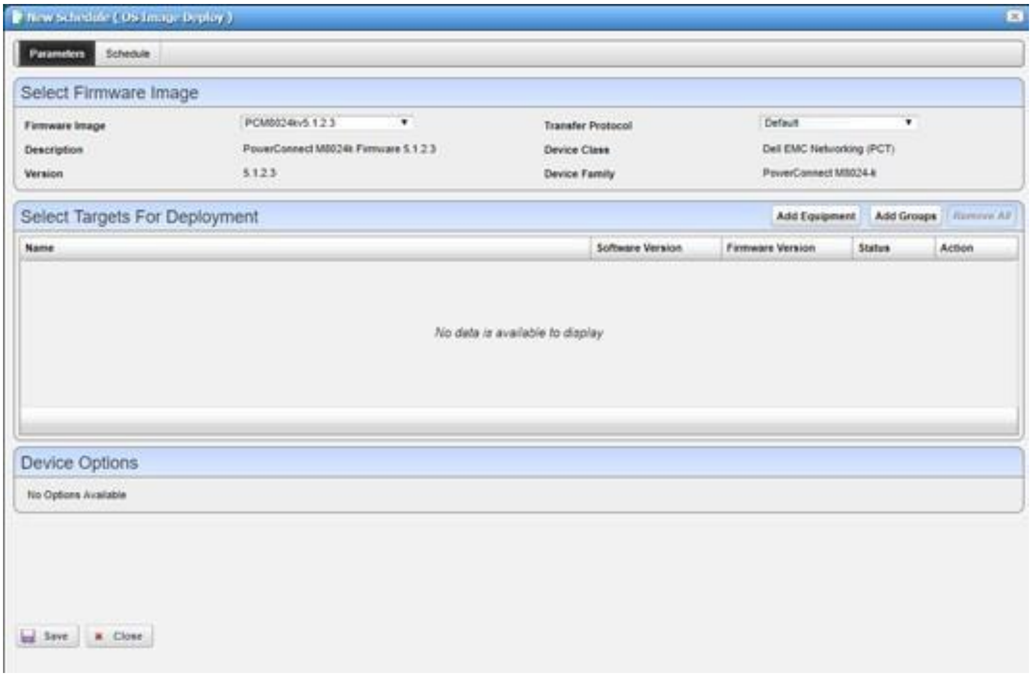

Select Firmware Image

Firmware Image: PCM8024kv5.1.2.3
Description: PowerConnect M8024k Firmware 5.1.2.3
Version: 5.1.2.3
Transfer Protocol: Default
Device Class: Dell EMC Networking (PCT)
Device Family: PowerConnect M8024-k

Select Targets For Deployment

Name	Software Version	Firmware Version	Status	Action

No data is available to display

Device Options

No Options Available

The OS Image Deploy Parameters window has the following fields and options. See Schedules Portlet on page 155 for scheduling fields and options descriptions.

Field/Option	Description
Firmware Image	Specifies which image you want to deploy.
Description	Displays the selected firmware image description.
Version	Displays the selected firmware image version.
Transfer Protocol	Specifies which protocol is used (Default, FTP, TFTP, or SFTP) to transfer files.
Device Class	Displays the selected firmware’s device class.
Device Family	Displays the selected firmware’s device family.
Add Equipment	Adds the selected equipment to the Targets for Deployment list.
Add Groups	Adds groups of devices to the Targets for Deployment list.
Remove All	Deletes all devices and groups from the Targets for Deployment list.

Field/Option	Description
Targets for Deployment	Provides the following information for each target: • Name • Software Version • Firmware Version • Status • Action deletes the entry
Device Options	Lists the tabs for each device based on the deploy target selected. Click a device tab to view its options. Select the appropriate action for each option.

Configuration Files

Use the Configuration Files portlet to view backed up configuration files. To see the most recent configuration files, see Top Configuration Backups on page 416.

Access this portlet from the Configuration Management page. This portlet has both a summary view and an expanded view. Each view could display different columns and has the same pop-up menu options available.

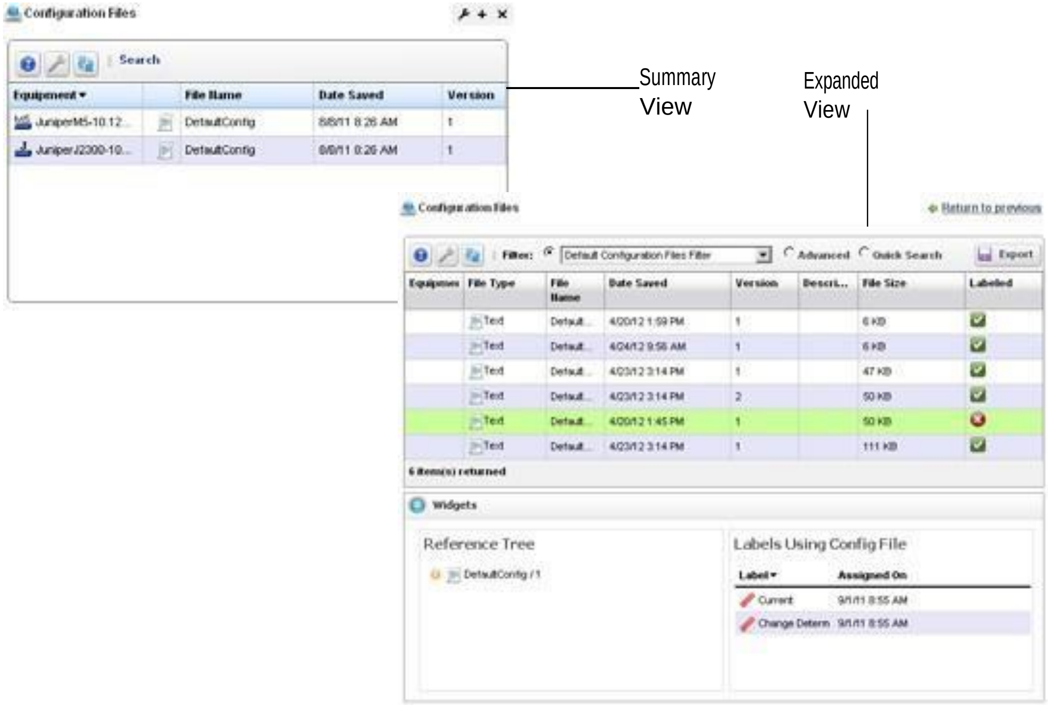

Equipment ▾		File Name	Date Saved	Version
JuniperM5-10.12...		DefaultConfig	8/8/11 8:26 AM	1
JuniperJ2300-10...		DefaultConfig	8/8/11 8:26 AM	1

Equipmen	File Type	File Name	Date Saved	Version	Descri...	File Size	Labeled
	Text	Default...	4/20/12 1:59 PM	1		6 KB	✓
	Text	Default...	4/24/12 9:56 AM	1		6 KB	✓
	Text	Default...	4/23/12 3:14 PM	1		47 KB	✓
	Text	Default...	4/23/12 3:14 PM	2		50 KB	✓
	Text	Default...	4/20/12 1:45 PM	1		50 KB	☒
	Text	Default...	4/23/12 3:14 PM	1		111 KB	✓

6 item(s) returned

Reference Tree

DefaultConfig / 1

Labels Using Config File

Label ▾	Assigned On
Current	9/1/11 8:55 AM
Change Determ	9/1/11 8:55 AM

Columns

Other than the general navigation and configuration options, the Configuration Files portlet includes the following columns. The columns displayed by default are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Equipment	The equipment impacted by the configuration. This field displays on the summary and expanded views by default.
File Name	A descriptive name given to the configuration. This field displays on the summary and expanded views by default.
Date Saved	The date and time the configuration file was saved. This field displays on the summary and expanded views by default.
Version	The version for the selected configuration. This field displays on the summary and expanded views by default.
File Type	The type of file in which the configuration is stored, such as text. This field displays on the expanded view by default.
Description	A detailed description for the configuration. This field displays on the expanded view by default.
File Size	The configuration file size (KB/MB). The Cruz converts from bytes to KB/MB and presents the file size in terms of KB/MB after some rounding. For example, 1484 bytes/1024 = 1.44921875 KB, which is rounded to 1.45 KB by the Cruz. In an advanced search by file size, search for a range to accommodate the rounding conversion. This field displays on the expanded view by default.
Labeled	An indicator that shows whether the configuration file has a label (checkmark) or not (X). When a configuration has a label applied, you cannot delete or archive it. This field displays on the expanded view by default.
Widgets	Additional information about the selected image, such as: • Reference Tree shows the configuration file name and provides access to the operations it supports by right clicking the file name. • Labels connected to the selected configuration file and the date on which the connection was made. The Widgets field is available **only** from the expanded view.

Pop-Up Menu

The Configuration File pop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
View	Opens the Configuration File editor, where you can see the configuration file content if it is a non-binary file.
Edit	Opens the Configuration File editor, where you can edit the backed up, non-binary configuration file. Configuration File Editor on page 451 for a description of these capabilities.
Assign Labels	Labels a single selected configuration file. A label selector appears that lets you select an existing label and create a new one. If you assign one file the Current label, others from the same device cannot have it. Cruz automates moving Current from one file to the other, if another has it. You can delete non-system labels from devices in the selector this menu item produces.

Menu Option	Description
Backup	Backs up the device (again) related to the selected file. **Note:** The Cruz automatically assigns the most recently restored file the Current label.
Compare to Label/ Compare Selected	Compares labeled configuration files to the current selection. See Configuration File Compare Window on page 449 for a description of this capability. You can create labels when you back up a config file, or you can compare to the default labels (Change Determination, Current, Compliant). If you select two configuration files in the expanded portlet, you can also Compare Selected.
Promote	Makes the selected configuration file available for mass deployment. This is a useful way to make a **pattern** configuration file to deploy to several devices. See Configuration on page 456 for additional information about how to do this.
Restore	Restores the selected file to its device. **Note:** The Cruz automatically assigns the most recently restored file the *Current* label.
Aging Policy	Opens the Aging Policy selector., where you specify how long this file remains in the database. See Implementing DAP on page 80 for more about configuring such policies.
Export	Saves the selected configuration file to an XML file. Click Download Export File to specify where to save the file. This option is useful as a backup or to share descriptors with other projects.
Import	Retrieves a file containing XML configuration descriptions. Some imports can come from a URL. This option is useful when sharing descriptors with other projects. You must import data into the correct port let. For example, you cannot import event data into the Actions portlet. If one type of data depends on another, you must import the other data before importing the data that depends on it.
Share with User	Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Configuration File Compare Window

In addition to letting you backup and restore configuration files, and deploy firmware updates to devices, the Configuration File Compare window manages viewing and comparing configuration files backed up from the selected devices.

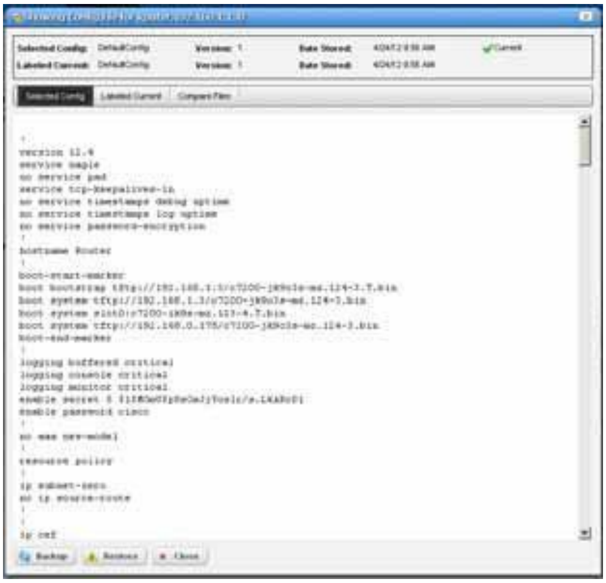

Selected Config: DefaultConfig Version: 1 Date Stored: [illegible]
Labeled Current: DefaultConfig Version: 1 Date Stored: [illegible]

```
!
version 12.4
service nagle
no service pad
service tcp-keepalives-in
no service timestamps debug uptime
no service timestamps log uptime
no service password-encryption
!
hostname Router
!
boot-start-marker
boot bootstrap tftp://192.168.1.1/c7200-jk9o3s-mz.124-3.T.bin
boot system tftp://192.168.1.3/c7200-jk9o3s-mz.124-7.bin
boot system slot0:c7200-ik8s-mz.123-4.T.bin
boot system tftp://192.168.0.175/c7200-jk9s-mz.124-3.bin
boot-end-marker
!
logging buffered critical
logging console critical
logging monitor critical
enable secret 5 [illegible]
enable password cisco
!
no aaa new-model
!
resource policy
!
ip subnet-zero
no ip source-route
!
!
ip cef
```

Compare and View options have the following limitations:

- If you select a config file that is a single file, without any historical precedent, no comparison option appears on the menu since the selected version does not have a prior version.
- If you select a single config file of version two or higher, comparison is an option. When selected, Cruz automatically compares against the most recent prior version for that device and file name.
- If you select two config files of any version, Cruz compares those two versions.
- If you select three or more config files, no comparison option appears.
- The View option appears for a single selection only, and only lets you view files that are not binary.

The File Type column indicates whether a configuration file is binary (not viewable) or text (viewable).

The following configuration file options are available.

Option	Description
View/Edit	Opens a window displaying the configuration file's contents. Use the browser's Find function to locate specific text within the configuration file. You can also select and copy text within this window. The **Selected Config** and **Live Config** (current) version and storage dates are displayed. When you perform a backup that differs from the config that is Labeled Current, that label changes to Live Config if changes are detected. **Selected Config** appears when you open this window from the Configuration Files portlet, but **Live Config/Current Config** appear side-by-side when you open this window from the Managed Resources portlet. You can also compare two different configurations (Selected Config and Labeled Current/Live Config) with the Compare Files tab panel. *Close* the screen with the buttons at its bottom. Notice you can also *Backup* or *Restore* what you are viewing with buttons at the bottom of the screen.
Assign Labels	Use this option to select an existing label or create a new one. You cannot assign System labels (*Current, Compliant* , and so on).
Compare Current v. Previous/to Label /Selected	You can compare configurations by right clicking a device, or two devices then selecting *Compare*. If you right click a single device with a previous backup, then the comparison is between the latest and next-to-latest backup. If it does not have a previous backup, then the menu offers to compare to a designated label. You can compare two different *Selected* devices too. Ctrl + click to select two different devices before you *Compare*. Notice that the *Prev/Next* buttons at the bottom of this screen can cycle through as many as five previous configuration files. The comparison screen appears with the configurations side-by-side (note the file names in the title bar of this screen). **Colors:** Lines that differ between the two configurations appear highlighted green. Lines missing in one, but that appear in another appear highlighted red. Added lines are yellow. Use the right/left arrows to page through the side-by-side comparison. The page numbers and beginning/forward/back/end arrows help you navigate between pages of pairs of files. Notice also that if you have more than two such files, a panel appears at the bottom that lets you navigate between adjacent pairs of such files (1 and 2, 2 and 3, 3 and 4, and so on). Click the Prev/Next links to move between pairs of files. Use the browser's "Find" function (Ctrl +F) to locate text within these views.
Backup/Restore	Select these to backup or restore a configuration file. See Backing Up Configurations on page 462 or Restoring Configurations on page 463 for step-by-step instructions. Some devices merge rather than replace configurations when you select Restore. (Cisco XR, for one)

Option	Description
Deploy	Select this option to deploy an OS Image (firmware). See Deploying Firmware on page 465 for more. Some devices, including the Dell EMC Networking FTOS C-Series and E-Series, first permit then drop telnet connections during deployment or file restoration when you select restart as part of the process. This can take from six to eight minutes, though it can take as long as fifteen minutes for a fully populated chassis. During that time, ping detects the device; however, Cruz cannot log in to the device until the reboot is complete. Restoring configurations to Dell Force 10 devices may produce errors when individual commands already exist in the running config and cannot be overwritten. Cruz ignores such errors and reports success by default since the errors indicate a command was not applied, not that restoration was unsuccessful. Best practice is to restore to startup config to avoid these errors, especially when scheduling backup or backing up a group on such devices. “Console Logging” must be turned off on all devices. The messages from console logging interfere with the communication between Cruz and the device (via CLI) and can disrupt supported functionality in Cruz.
Export/Import	Export lets you save a local copy of the selected config file. Import opens a screen that lets you select a locally-accessible file to store, view, compare and deploy.

Configuration File Editor

Use the Configuration File editor to manually edit configuration files and save them to the Cruz database. You can edit only non-binary configuration files.

When you select a file from the Configuration Files portlet, and then right-click to select Edit, the Configuration File editor opens with the Find/Replace options.

Click the search/replace tool (magnifying glass) to open the text search feature. Select `A/a` to make your search case-sensitive, or `RegEx` to use regular expressions to search.

Click the Find button to locate text in the config file. Click Replace to replace the text found. Select All and then click Replace to globally replace all instances of the text found within the configuration file.

Click Save to preserve your edits, or Close to abandon them. Notice that the edited configuration is listed with the other Configuration Files in the portlet as a different version than the original. The version increments by one every time you edit and save aconfiguration.

Image Repository

Use the Image Repositoryportlet to managefirmware updates to deployto devicesin your network or manage configurations that you want to deploy to several devices. Add these files to your Cruz before deploying them.

This portlet is intended for users who are interested in managing firmware updates deployed to devices or manage configurations deployed to several devices.

Access this portlet from the Configuration Management page. This portlet has both a summary view and an expanded view. Each view could display different columns and has the same pop-up menu options available.

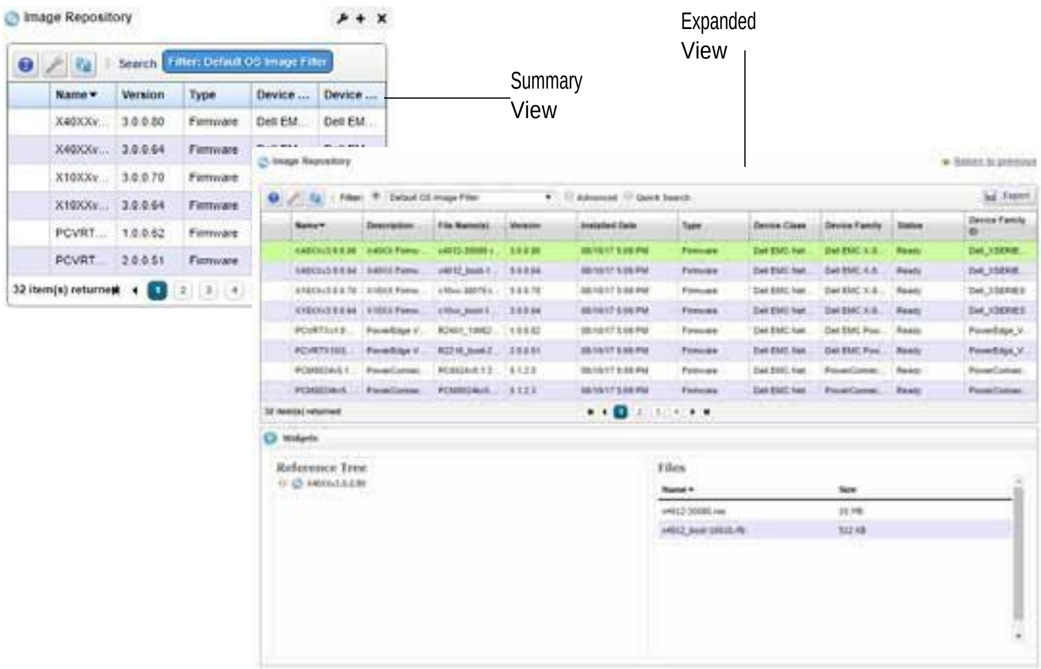

Name	Version	Type	Device ...	Device ...
X40XXv...	3.0.0.80	Firmware	Dell EM...	Dell EM...
X40XXv...	3.0.0.64	Firmware		
X10XXv...	3.0.0.70	Firmware		
X10XXv...	3.0.0.64	Firmware		
PCVRT...	1.0.0.62	Firmware		
PCVRT...	2.0.0.51	Firmware		

32 item(s) returned 1 2 3 4

Name	Description	File Name(s)	Version	Installed Date	Type	Device Class	Device Family	Status	Device Family ID
[illegible]	[illegible]	[illegible]	3.0.0.80	08/10/17 5:08 PM	Firmware	Dell EMC Net...	Dell EMC X-S...	Ready	Dell_XSERIE...
[illegible]	[illegible]	[illegible]	3.0.0.64	08/10/17 5:08 PM	Firmware	Dell EMC Net...	Dell EMC X-S...	Ready	Dell_XSERIE...
[illegible]	[illegible]	[illegible]	3.0.0.70	08/10/17 5:08 PM	Firmware	Dell EMC Net...	Dell EMC X-S...	Ready	Dell_XSERIES
[illegible]	[illegible]	[illegible]	3.0.0.64	08/10/17 5:08 PM	Firmware	Dell EMC Net...	Dell EMC X-S...	Ready	Dell_XSERIES
PCVRT...	PowerEdge V...	[illegible]	1.0.0.62	08/10/17 5:08 PM	Firmware	Dell EMC Net...	Dell EMC Pow...	Ready	PowerEdge_V...
PCVRT...	PowerEdge V...	[illegible]	2.0.0.51	08/10/17 5:08 PM	Firmware	Dell EMC Net...	Dell EMC Pow...	Ready	PowerEdge_V...
[illegible]	PowerConnec...	[illegible]	5.1.2.3	08/10/17 5:08 PM	Firmware	Dell EMC Net...	PowerConnec...	Ready	PowerConnec...
[illegible]	PowerConnec...	[illegible]	5.1.2.3	08/10/17 5:08 PM	Firmware	Dell EMC Net...	PowerConnec...	Ready	PowerConnec...

32 item(s) returned 1 2 3 4

Name	Size
[illegible]	33 MB
[illegible]	522 KB

Columns

Other than the general navigation and configuration options, the Image Repository portlet includes the following columns. The columns displayed by default are noted.

You can view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Name	The image name. This field displays on the summary and expanded views by default.
Version	The version for the selected image. This field displays on the summary and expanded views by default.
Type	The type of image, such as firmware, configuration, and so on. This field displays on the summary and expanded views by default.
Device Class	The device class for the selected image. This field displays on the summary and expanded views by default.

Column	Description
Device Family	The device family for the selected image. This field displays on the summary and expanded views by default.
Description	More details about the firmware or configuration. This field displays on the expanded view by default.
File Names	The files related to the selected image. These files are also listed in the Widgets field. This field displays on the expanded view by default.
Installed Date	The date and time in which the firmware was installed. This field displays on the expanded view by default.
Status	The device's readiness to deploy (failed, importing, ready). This field displays on the expanded view by default.
Device Family ID	The family of models for which the image works. This field displays on the expanded view by default.
Widgets	Additional information about the selected image, such as: • Reference Tree shows the connection between the image and its device. • Related files. Also listed in the File Names field. The Widgets field is available **only** from the expanded view.

Pop-Up Menu

The Image Repository pop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
New	Opens the Firmware Image Editor window or the Configuration Image Editor window, where you define a firmware image or configuration, respectively.
Edit	Opens the Firmware Image Editor or the Configuration Image Editor if the selected line is a configuration image.
Deploy	Opens the Deploy Firmware window, where you define deployment options for the selected image. For this to function, you must enable a server. See Configuration File Compare Window on page 449 for details.
Download Firmware For	Downloads firmware from the internet for the selected device. The devices that support downloading (such as Dell DNOS or Dell FTOS) are listed as an option to the Download Firmware For menu option. Select the device for which you want to download OS images and the Cruz automatically downloads it.
Delete	Removes the selected OS image or configuration from the list.
Share with User	Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Firmware Image Editor

Use the Firmware Image editor to create or maintain OS images. This editor has the following panels:

- General Parameters specifies the general parameters, such as name, version, class, and so on
- Image Files specifies or displays image-related files
- Readme File to view or specify a text file related to the image

Click *Save* to preserve the OS Image you have configured or *Cancel* to exit these screens without saving.

Access this editor from the Image Repository portlet byselecting the New > Firmware Image pop-up menu option or selecting a firmware image and then the Edit pop-up menu option.

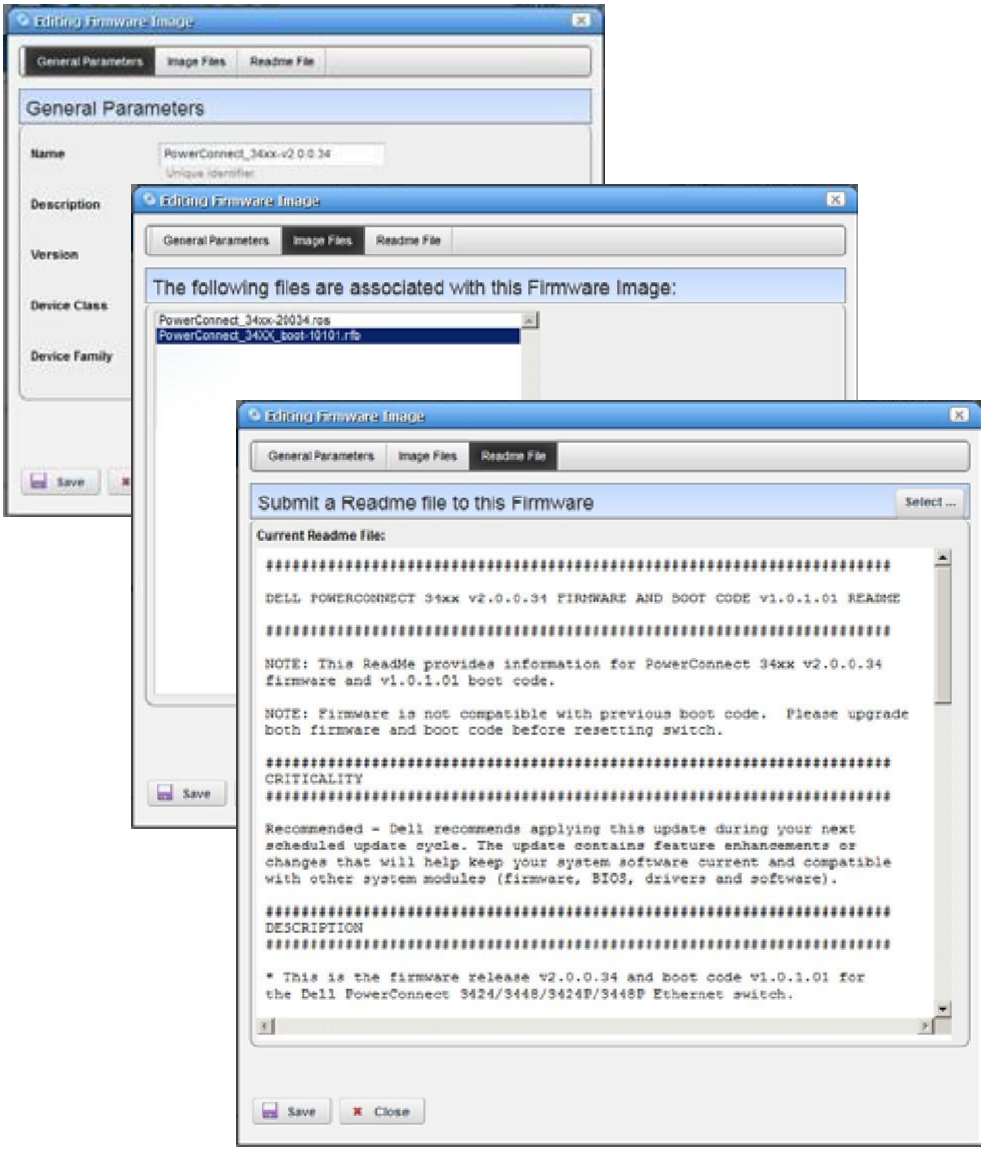

General Parameters

Name: PowerConnect_34xx-v2.0.0.34
Unique identifier
Description
Version
Device Class
Device Family

The following files are associated with this Firmware Image:

PowerConnect_34xx-20034.ros
PowerConnect_34XX_boot-10101.rfb

Submit a Readme file to this Firmware

Current Readme File:

```
########################################################################

DELL POWERCONNECT 34xx v2.0.0.34 FIRMWARE AND BOOT CODE v1.0.1.01 README

########################################################################

NOTE: This ReadMe provides information for PowerConnect 34xx v2.0.0.34
firmware and v1.0.1.01 boot code.

NOTE: Firmware is not compatible with previous boot code.  Please upgrade
both firmware and boot code before resetting switch.

########################################################################
CRITICALITY
########################################################################

Recommended - Dell recommends applying this update during your next
scheduled update cycle. The update contains feature enhancements or
changes that will help keep your system software current and compatible
with other system modules (firmware, BIOS, drivers and software).

########################################################################
DESCRIPTION
########################################################################

* This is the firmware release v2.0.0.34 and boot code v1.0.1.01 for
the Dell PowerConnect 3424/3448/3424P/3448P Ethernet switch.
```

The Firmware Image editor has the following fields or options.

Field/Option	Description
Name	Specifies a user-defined name for the image.
Description	Provides more details about the image.
Version	Provides a descriptive version number.
Device Class	Lists the available device firmware classes from which to choose. Based on the device class selected, the Device Family list is populated.
Device Family	Lists the device families available for the selected Device Class.
Image Files	Imports a file from disk or from UR L when you create an image. Displays imported files when editing and image. Because OS images can consist of multiple files, you can import multiple files here.
Readme File	Imports a read me text file to accompany the image or displays the imported file content.

Configuration Image Editor

Use the Configuration Image editor to create and maintain image configurations. This editor has the following panels:

- General Parameters
- Configuration

Access this editor from the Image Repository portlet by selecting the New > Configuration Image pop-up menu option or by selecting a configuration image and then the Edit pop-up menu.

General Parameters

Use the General Parameters panel to name and describe the configuration file and configure a filter to screen restoration targets.

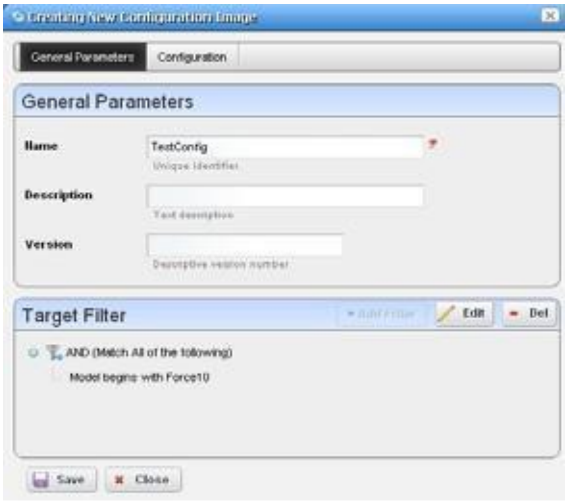

Name: TestConfig
Unique Identifier

Description:
Text description

Version:
Descriptive version number

AND (Match All of the following)
Model begins with Force10

The Configuration Image editor has the following fields or options.

Fields/Options	Description
Name	Specifies a unique identifier for the image configuration.
Description	Provides more details on the image configuration.

Fields/Options	Description
Version	Automatically tracks changes to the original.
Target Filter	Defines which devices to target. When targets fail, restoration skips them. Add filters by creating a filter or copying and existing filter.

Configuration

Use the Configuration panel to configure what is restored, and what is variable in mass deployments. Initially the configuration document is empty if you are creating a configuration image. However, it displaysdata fromanypromoted configuration file if it originated asapromoted configuration file.

Access this information by clicking the Configuration Image edit or Configuration tab.

The Target Params list includes all available discoverable parameters. These parameters are stored in the Cruz database. Some parameters may not appl y to a specific device or configuration file. To insert a parameter into the document, double-click it. If you are adding a parameter to an existing document, make sure that you indicate where you want to add it before double-clicking the parameter.

For example, if a Contact parameter appears in the configuration file, delete the specifics retrieved from a device's configuration and then double-click the Contact target parameter. The following line is inserted into the document where you put the cursor:

```
$_EquipmentManager_RedCell_Config_EquipmentManager_Contact
```

When you deploy this configuration file to the devices that pass the General Parameters target filter, the Cruz updates this parameter with discovered data retrieved from the device before restoring the configuration. This facilitates deploying the same configuration to many devices while retaining individual target parameters like Contacts, DNS Hostname, and so on.

NOTE:

If you want tocompare different promoted configuration file templates for the samedevices, deploy both template #1 and template #2, then compare them in the Configuration Files on page 446 portlet. By default, the Description notes that such configuration files are "Created from template."

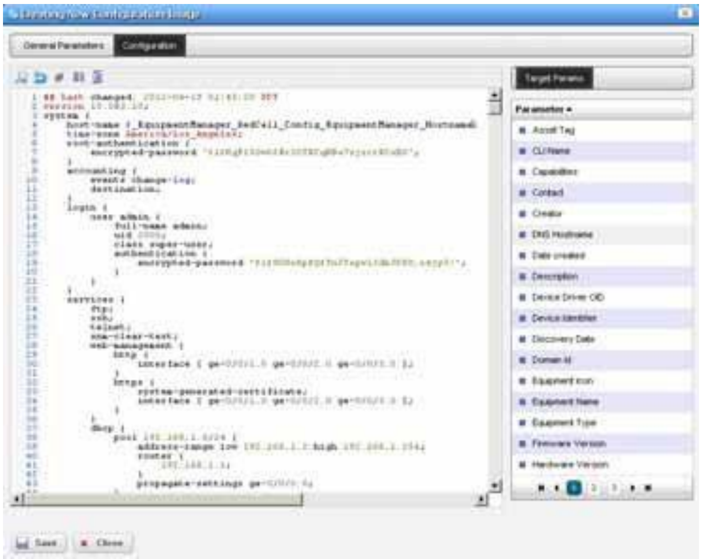

```
## Last changed: [illegible]
version [illegible];
system {
    host-name $_EquipmentManager_RedCell_Config_EquipmentManager_Hostname;
    time-zone America/Los_Angeles;
    root-authentication {
        encrypted-password "[illegible]";
    }
    accounting {
        events change-log;
        destination;
    }
    login {
        user admin {
            full-name admin;
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "[illegible]";
            }
        }
    }
    services {
        ftp;
        ssh;
        telnet;
        xnm-clear-text;
        web-management {
            http {
                interface [ ge-0/0/1.0 ge-0/0/2.0 ge-0/0/3.0 ];
            }
            https {
                system-generated-certificate;
                interface [ ge-0/0/1.0 ge-0/0/2.0 ge-0/0/3.0 ];
            }
        }
        dhcp {
            pool 192.168.1.0/24 {
                address-range low 192.168.1.2 high 192.168.1.254;
                router {
                    192.168.1.1;
                }
                propagate-settings ge-0/0/0.0;
```

Deploy Firmware

Use the Deploy Firmware window to configure a deployment, whether triggered fromresource groups, individual resources, or the Image Repository portlet. Deployment validates that the selected image is appropriate for the selected devices, or appropriate devices within a group.

You have the option to select a transfer protocol, add target equipment or groups, execute immediately, or schedule the deployment for a later date/time. If you create a deployment schedule, save your configuration. Otherwise, close the window to abandon your edits.

NOTE:

When you add firmware to the Image Repository for Dell 35xx and 55xx devices, you must add both the boot image and firmware image together to deploy to these devices.

You may see multiple options for selecting the configuration file to backup for PowerConnect (not Dell EMC Networking FTOS) devices. Layer 2 Powerconnect switches have just running and startup options while the Layer 3 router has running, startup and backup options, so different options appear for the two sets of switches. When you do file backup for a group of devices, all those options are combined. Select only the top entry selection for execution.

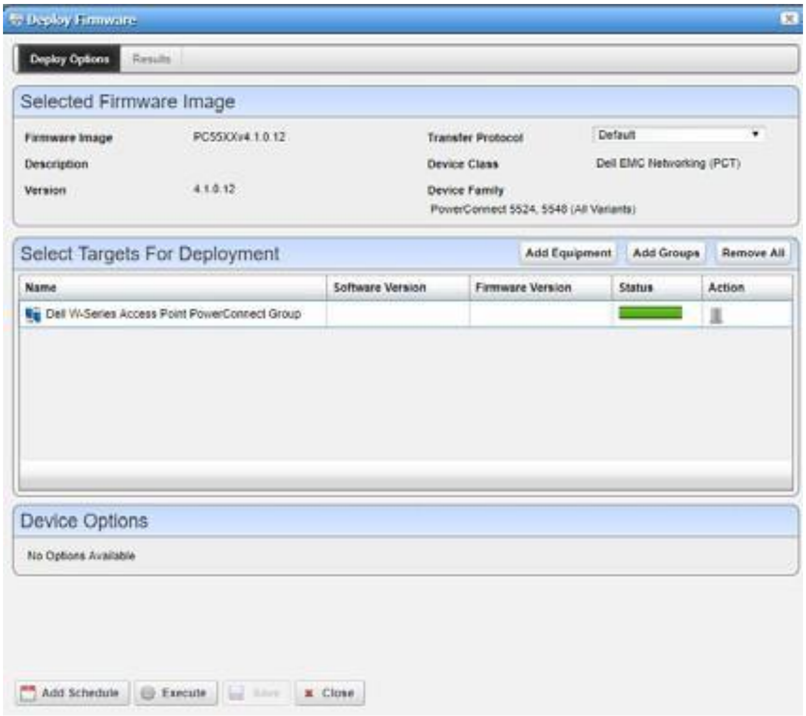

Selected Firmware Image

Firmware Image: PC55XXv4.1.0.12

Description:

Version: 4.1.0.12

Transfer Protocol: Default

Device Class: Dell EMC Networking (PCT)

Device Family: PowerConnect 5524, 5548 (All Variants)

Select Targets For Deployment

Name	Software Version	Firmware Version	Status	Action
Dell W-Series Access Point PowerConnect Group			🟩	

Device Options

No Options Available

File Servers

Use the File Servers portlet to define and maintain a list of internal and external file servers. This portlet has only the summary view.

This portlet is intended for anyone that want to define and maintain file servers and they have the permission to do so.

Access this portlet from the Configuration Management page. This portlet has only a summary view with all columns showing and pop-up menu options available.

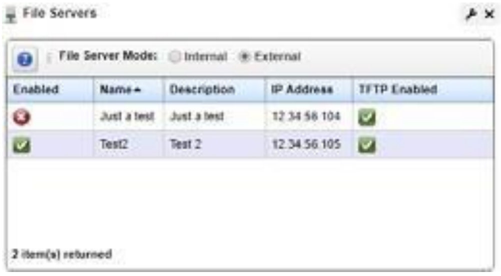

Enabled	Name ▲	Description	IP Address	TFTP Enabled
☒	Just a test	Just a test	12.34.56.104	☑
☑	Test2	Test 2	12.34.56.105	☑

2 item(s) returned

Columns

Other than the general navigation and configuration options, the File Servers portlet includes the following columns. All columns are displayed by default and there is not option to change the columns show.

Column/Option	Description
File Server Mode	An option to specify whether the file server is internal or external. The default is External.
Enabled	An indicator that shows whether the server is enabled (checkmark) or not (X).
Name	The user-defined name for the server.
Description	A detailed description for the server.
IP Address	The IP address used by the application.
TFTP Enabled	An indicator that shows whether the server supports TFTP (checkmark) or not (X).

Pop-Up Menu

The File Servers pop-up menu provides access to th e following options. Right-click a row to access these options.

Menu Option	Description
New	Opens the External File Server editor, where you define an external file server by setting general parameters, specify server type (FTP, FTP/SCP) and whether you want TFTP support, and the authentication settings.
Edit	Opens the External File Server editor, where you modify the selected external file server’s general parameters, server type (FTP, FTP/SCP) and whether you want TFTP support, and the authentication settings.
Disable	Disables the selected file servers from the list. A successful/failed message is displayed. This option shows only if the selected file servers are enabled.

Menu Option	Description
Enable	Disables the selected file servers from the list. A successful/failed message is displayed. This option shows only if the selected file servers are not enabled.
Test	Verifies your settings for the selected file server. Shows the test progress and the results in the Job Status window.
Delete	Removes the selected file servers from the list. A confirmation message is displayed. Click Delete to remove.
Share with User	Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

External File Server

Use the External File Server editor to define and maintain external file servers.

Access this editor by right clicking the File Servers portlet and then selecting New, or by right-clicking and existing file server definition and thenselecting Edit.

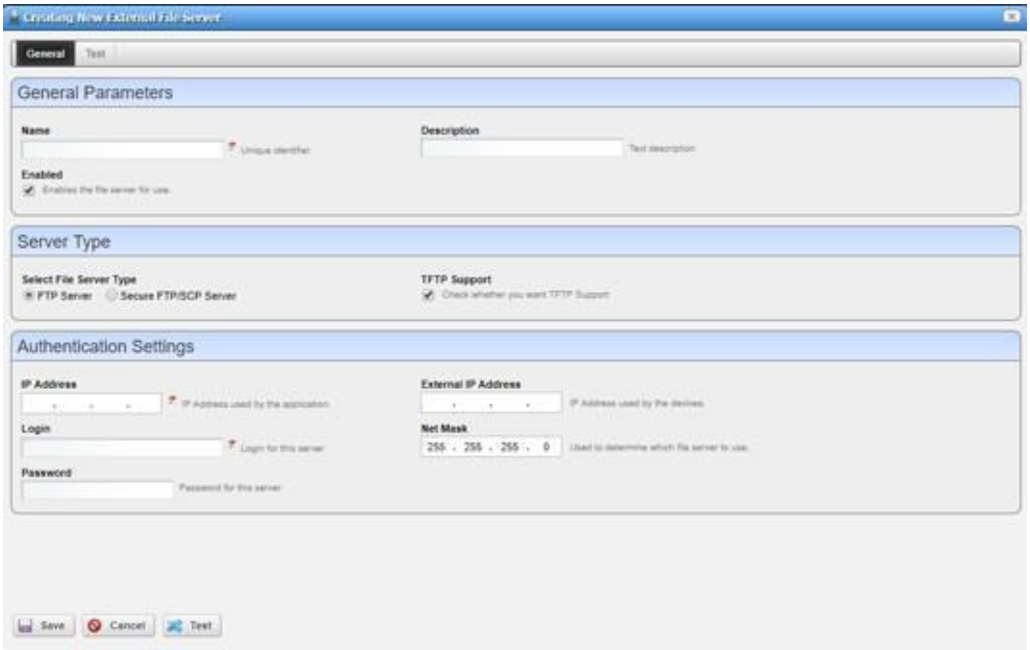

Creating New External File Server

General Parameters

Enabled: ☑ Enables the file server for use.

Server Type

Select File Server Type: ● FTP Server ○ Secure FTP/SCP Server

TFTP Support: ☑ Check whether you want TFTP Support

Authentication Settings

Net Mask: 255 . 255 . 255 . 0 Used to determine which file server to use.

The External File Server editor has the following fields and options.

Field/Option	Description
Name	Is a unique name for the file server.
Description	Provides a detailed description for the file server.
Enabled	Activates the file server for use. Selected by default.
Server Type	Specifies one of the following file server types: • FTP Server (default) • Secure FTP/SCP Server **Note:** Secure FTP connections (scp/sftp) often require SSH services be enabled on the devices addressed. Make sure that your system’s server and sftp/scp file server can also access the devices with SSH.

Field/Option	Description
TFTP Support	Indicates whether you want TFTP support. The default is to have TFTP support (selected).
IP Address	Specifies the IP address used by the application. This address is required.
External IP Address	Specifies the IP address used by the devices.
Login	Specifies the login defined for this server. This is required.
Net Mask	Determines which file server to use. The default is 255.255.255.0. If you have several servers, the specified net mask determines which server communicates with devices in which portion of the network.
Password	Specifies the optional password for the login name.
Test (panel)	Shows the Job Status and whether the test was successful or failed. Click the Test button to verify your file server configuration.

Understanding FTP/TFTPServers

Before you can push and pull configuration files to and from devices or deploy firmware updates, you need to configure FTP, TFTP, or both file servers.

△ CAUTION:

TheinternalFTP/TFTPserverisfortestingonly,notforproductionuse.Servicediscoverymaynot functioncorrectly withtheinternal FTP/TFTPserver. Nointernalserveris availableonLinuxinstallations.

Port conflicts prevent an external file server and internal file server from operating on the same machine.

You need not be concerned that the internal server may provide insecure access to the application. The internal file server was designed to be ultra-secure. It creates a separate authentication and virtual file system for each file retrieved. It also only responds to Cruz internal requests.

NOTE:

The internal FTP server is primarily for testing, not production systems. It may not function in all cases. Configure an external file server and use it instead if and when the internal file server fails.

The following sections provide some more specific information related to FTP/TFTP servers:

- File Permissions
- Recommended Windows File Servers
- External File Server Editor

File Permissions

The application automatically deletes any temporary file created as part of an FTP/TFTP interaction. If the directory you selected for your servers does not have permissions needed to make these deletions, transactions still proceed. The Cruz application does print a warning saying the deletion failed, and the details panel suggests checking to make sure delete permissions exist in the relevant directory. Omitting such permissions causes no loss of functionality, but the server may fill up with the remnants of old transactions.

When you test a TFTP server on a Window s system, you may see the following error:

```
FTP umask/permissions of file on server are incorrect
```

This is an artifact of Windows permission structure and may be safely ignored (test your TFTP server just to makesure).

NOTE:

For Linux TFTP servers, a typical configuration line in the /etc/xientd.d/tftp file is:

```
server_args = -c -p -u ftpuser -U 177 -s /home/ftpuser
```

Recommended Windows File Servers

Theopen source Filezilla server worksasaserviceon Windowsservers. Anylogin/password access to these servers goes in the File Server editor login/password fields. To support TFTP, try the Tftpd32 or Tftpd64 (for 32-bit or 64-bit machines) open source.

These servers must read/write from/to the same directory. Also, make sure that the directory offers openread/write/execute permissions so youcanretrievefilesput there temporarily, and delete them once the process is done with them.

External File ServerEditor

From the File Servers portlet, access the External File Server editor to manage your file servers and files. See External File Server on page 459 for a detailed description of this editor.

NOTE:

Secure FTP connections (scp/sftp) often require SSH services be enabled on the devices addressed. Make sure that your system's server and sftp/scp file server can also access the devices with SSH

FTP servers typically must be on the same side of the firewall as the devices with which they communicate. If you have several such servers, the specified Net Mask also determines which server communicates with devices in which portion of thenetwork.

The file server uses an internal, local LAN address (192.168.100.100 example). However, the routers with which it communicates often cannot communicate to such internal addresses. Therefore, an external/reachable address is necessary. Youcan know an IP address used by the Cruz application, and another External IP Address used by these devices. If you configure multiple file servers, the Cruz application selects the server with the Net Mask whose subnet is closest to the devices with which it communicates.

Backing Up Configurations

The application simplifies backing up devices so you can always have their configuration, even if the device file becomes corrupted or out-of-date.

You can back up several devices at once. Use the standard Ctrl+ click option in the expanded Managed Resources portlet to select several devices. Or, look for the appropriate group in the Managed Resource Groups portlet.

Back up your device configurations a s follows.

1 Make sure that an FTP or TFTP server is configured to handle the backup. See File Servers on page 458.

2 Navigate to the Managed Resources portlet.

3 Right-click the devices and then select File Management > Backup.

The Backup Device window is displayed.

4 Specify options, targets, and device options.

See Config File Backup on page 443 if you need a detailed description of these fields.

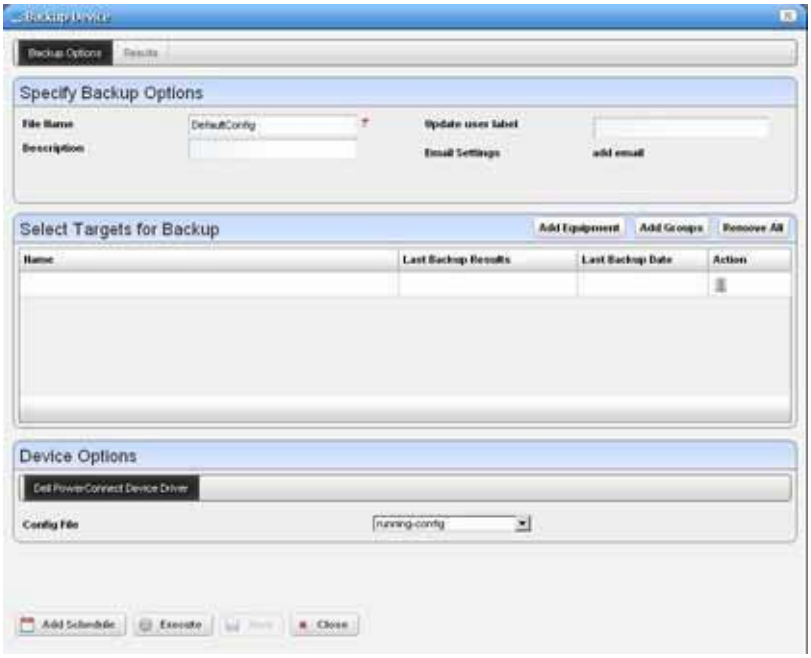

Backup Device

Backup Options | Results

Specify Backup Options

File Name: DefaultConfig *

Description:

Update user label:

Email Settings: add email

Select Targets for Backup

Name	Last Backup Results	Last Backup Date	Action

Device Options

Dell PowerConnect Device Driver

Config File: running-config

5 Execute the backup now. Otherwise, continue with step 6.

The Results panel displays the message traffic between the Cruz application and the devices. See Audit Trail Portlet on page 150.

6 Schedule the backup for a later time.

7 Click Save to preserve this scheduled configuration.

Restoring Configurations

Restore a configuration file to a device as follows.

1 Make sure that an FTP or TFTP server is configured to handle the restore. See File Servers on page 458.

2 Navigate to the Managed Resources portlet.

3 Right-click the devices and then select File Management > Restore.

The Restore Device window is displayed. If you selected multiple devices, select a target from the list before continuing with the next step.

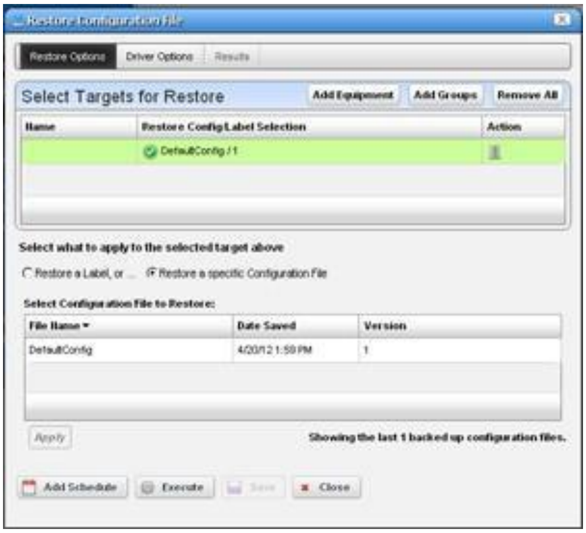

Select Targets for Restore

Name	Restore Config/Label Selection	Action
	✓ DefaultConfig / 1	

Select what to apply to the selected target above

○ Restore a Label, or ... ● Restore a specific Configuration File

Select Configuration File to Restore:

File Name ▾	Date Saved	Version
DefaultConfig	4/20/12 1:59 PM	1

Showing the last 1 backed up configuration files.

4 Modify the Targetsfor Restore list as necessary by:

This screen’s tabs let you configure the following:

- Deleting a target
- Adding equipment
- Adding groups

5 Select a target for restore.

6 Set and apply the restore options for the selected target.

This panel lets you select either a label (like *Current*, *Compliant* and so on—a selector listing available labels appears onscreen once you click this option), or *Restore a specific Configuration File*. The latter lists available files and lets you click to select.

7 Select the driver options for the selected target.

The Driver Options tab lets you select device-specific restoration options. If you are restoring to a group or multi-selected devices, as many tabs appear as are necessary to configure different restorations for different devices

8 Repeat steps 5 through 7 for each target listed.

9 Execute the backup now. Otherwise, continue with step 10.

The Results panel displays the message traffic between the Cruz application and the devices. See Audit Trail Portlet on page 150.

10 Schedule the restore for a later time.

11 Click Save to preserve this scheduled configuration.

Deploying Firmware

Deploy a device firmware image as follows.

1. Make sure that an FTP or TFTP server is configured.

   See File Servers on page 458.

2. Navigate to the Managed Resources portlet.
3. Right-click the devices and then select File Management > Deploy.

   The Deploy Firmware window information based on the selected device.

4. Select the firmware image you want to deploy.

   The fields are populated based on your selection.

5. Select a transfer protocol.
6. Modify the targets as needed, such as deleting a target, adding equipment, or adding groups.

   What you can do is restricted to devices that match the deploy file's type.

△ CAUTION:

Youcanalsoselectdevices, then change the OS selection so apotential mismatch mayoccur. This *may* trigger deployment rejection by the device but is not a recommended experiment.

7. Specifyanydeviceoptions available to fine-tune the deployment.

   The options available are vendor specific.

8. Execute the backup now. Otherwise, continue with step 9.

   The Results panel displays the message traffic between the Cruz application and the devices. See Audit Trail Portlet on page 150.

9. Schedule the deployment for a later time.
10. Click Save to preserve this scheduled configuration.

Restoring a Configuration to Many Devices

You can restore a single configuration to many discovered devices without overwriting those devices' essential information. Restoreasingleconfiguration to manydiscovered devicesasfollows.

1. Back up a single device's configuration that is closest to the configuration you want to see.
2. Navigate to the File Management Menu portlet.
3. Right-click this backed up file in the File Management Menu portlet and Promote it (so it eventually appears in the portlet).

   The Editor appears for the promoted configuration.

4. Name the file.
5. Configure a Target Filter if necessary, from the editor's General Parameters panel to confine it to certain devices by default.
6. Select the Configuration tab.
7. Locate the parameters you want to preserve in discovered devices when you restore this file.

   This can include items like the device's DNS Hostname, IP Address, and so on. Delete the file's specifics and double-click to insert the Target Params in place of these variables.

8 Save the configuration.

9 Right-click and then select Deploy.

This deploys the configuration to the targets you select.

Youcanselect Generate and saveforconfiguration onlyifyou simply want to deploylater and save for now. Also, optionally name a Configuration File Label for the deployed files.

10 Select the devices, or groups of devices to which you want to deploy.

11 Click Save, Execute or Add Schedule depending on your desired outcome.

If you click Execute, confirm the action.

When Cruz deploys the configuration, it reads the Target Params from those discovered for each device, inserts them in the deployed config file, then restores the configuration, device by device, skipping any that do not pass the filter set up in step 4.

Creating/Comparing Promoted Configuration Templates

If you want to store “template” configurations you have promoted, and compare them to previous templates, here are the steps to do that.

1 Select a configuration file and promote it to be a template.

Suggestions: Settle on a naming convention for these, perhaps one that includes a date so you can easily find and compare templates from different dates. You can also create a label for such configurations by simply typing in the *Label for Configuration* field. This should make such configurations easy to retrieve, particularly in the expanded Configuration Files portlet.

2 Enter the needed variables.

3 Save the template.

4 Right-click to *Deploy*it.

5 Select a single target and make sure to check the *Generate and Save Configuration Only checkbox*.

6 *Execute*. Rather than deploying the file, this saves a copy of the file as it would be generated for the single target in the Configuration Files.

7 When you have more than one of these configurations, find the files and Ctrl + click to select a pair of them, then right-click to *Compare* them.

Troubleshooting Backup, Restore or Deploy Issues

Here are some steps to troubleshoot issues you ma n encounter during a backup, restore, or deploy action. The following example steps are for troubleshooting a backup, but the steps apply to a restore or deploy action too.

1 Make sure the FTP/TFTP server you are using is correctly set up, and still active.

Use the External File Server test button to confirm the servers work.

2 Look in the Audit Trail portlet for the failed job.

3 Copythe "Executingreadcommands against the device" informational message content.

For example:

```
copy running-config tftp://192.168.0.138/010128030139_DefaultConfig
```

4 Use Direct Access to get to a Telnet/SSH comma nd line on the device having backup issues.

If you cannot get to a command line, then see Incomplete Discovery on page 104 for the way to remedy that.

5 Paste the command you have copied in step 3 after the prompt.

6 Press [Enter].

7 Observe whether the device executes this command.

If the device does not successfully execute the command, either the authentication used does not have permission to do such commands, or the device is configured to prohibit their execution.

Consult with your network administrator to get the correct authentication, and either revise the Discovery Profile that discovered this device, delete the device from the Cruz, and the discover it again. Or, Edit the device and enter the revised authentication/management interface combination.

If the device is configured to prohibit this command's execution, then consult the device's documentation and revise that.

9

Change Management and Compliance

The Cruz Change Management and Compliance (also known as ProScan) utility lets you scan stored configurations or Adaptive CLI show command output to verify that managed devices comply with company, department, or industry standards. The ProScan utility also automatically tracks all changes to managed devices. You can repo rt on user-specified values found in persisted backup configuration files for a group of devices. This lets network managers, security officers, and external auditors to generate detailed audit trail documents to validate compliance with both internal standards (ISO 17799, NSA Guidelines) and industry regulations (Sarbanes-Oxley, GLBA, HIPAA).

Compliance reporting lets you specifya text string, regular expression, oroptionally the generated configlet from File Management (NetConfig) for matching. Group results must be separated by device like Adaptive CLI Manager. When Compliance policies run, the application emits notifications whose contents depend on whet her compliance was or was not maintained.

NOTE:

Your system may have several Compliance Policy examples. You can use the examples as is or copy and alter them to suit your network.

This section describes the portlets and editors us ed to perform change management tasks and then provides those tasks. If you already have a good understanding of the portlets and editors, go directly to the tasks you want to perform.

Change Management Portlets and Editors

This section describes the portlets and editors related to change management and compliance tasks:

- Compliance Policies/ProScan
- Compliance Remediation Actions
- Compliance Schedules
- Compliance Alarms
- Compliance Policy Summary

By default, access these portlets by selecting Compliance from the navigation bar.

Compliance Policies/ProScan

Use the Compliance Policies/Proscanportlet to configurecompliance requirements. Limit the visible policies using the filtering in the expanded view.

This portlet is intended for users who are interested in configuring and maintaining compliance policies or policy groups.

By default, access this portlet by selecting Compliance from the navigation bar. This portlet has both a summary viewand an expanded view. Each view could displaydifferent columns and has the same pop-up menu options available.

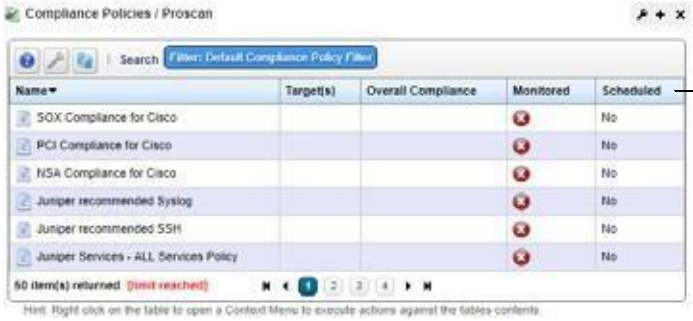
Compliance Policies / Proscan
Search
Filter: Default Compliance Policy Filter

Name▾	Target(s)	Overall Compliance	Monitored	Scheduled
SOX Compliance for Cisco			✗	No
PCI Compliance for Cisco			✗	No
NSA Compliance for Cisco			✗	No
Juniper recommended Syslog			✗	No
Juniper recommended SSH			✗	No
Juniper Services - ALL Services Policy			✗	No

50 item(s) returned (limit reached)

Hint: Right click on the table to open a Context Menu to execute actions against the tables contents

Summary View

Expanded View

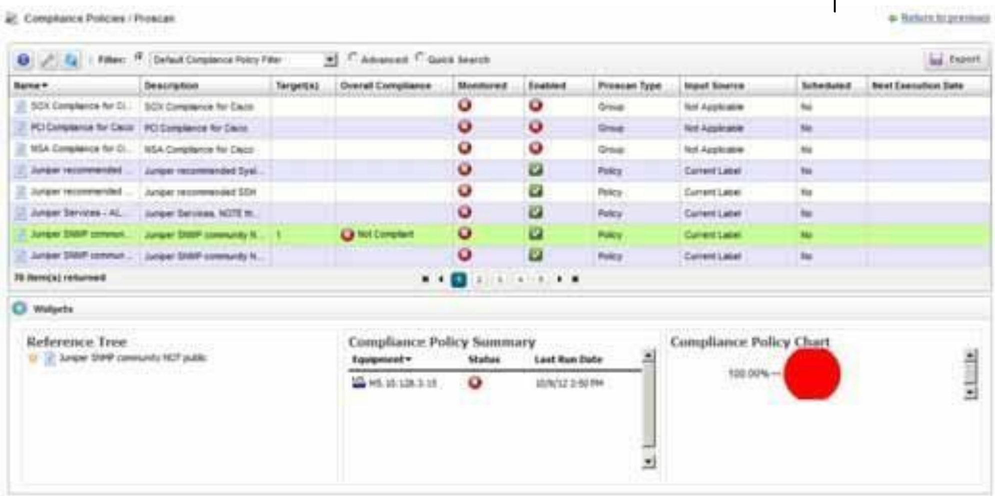
Compliance Policies / Proscan
Return to previous
Filter:
Default Compliance Policy Filter
○ Advanced ○ Quick Search
Export

Name▾	Description	Target(s)	Overall Compliance	Monitored	Enabled	Proscan Type	Input Source	Scheduled	Next Execution Date
SOX Compliance for Ci...	SOX Compliance for Cisco			✗	✗	Group	Not Applicable	No	
PCI Compliance for Cisco	PCI Compliance for Cisco			✗	✗	Group	Not Applicable	No	
NSA Compliance for Ci...	NSA Compliance for Cisco			✗	✗	Group	Not Applicable	No	
Juniper recommended ...	Juniper recommended Sysl...			✗	✓	Policy	Current Label	No	
Juniper recommended ...	Juniper recommended SSH			✗	✓	Policy	Current Label	No	
Juniper Services - AL...	Juniper Services. NOTE m...			✗	✓	Policy	Current Label	No	
Juniper SNMP commun...	Juniper SNMP community N...	1	✗ Not Compliant	✗	✓	Policy	Current Label	No	
Juniper SNMP commun...	Juniper SNMP community N...			✗	✓	Policy	Current Label	No	

78 item(s) returned

Equipment▾	Status	Last Run Date
HS 10.128.3.15	✗	10/9/12 2:50 PM

Columns

Other than the general navigation and configuration options, the Compliance Policies/ProScan portlets (summary and expanded views) include the followingcolumns. The columns displayed by default are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Name	The user-defined name for the policy or policy group. This field displays on the summary and expanded views by default.
Targets	The number of devices to scan. This field displays on the summary and expanded views by default.
Overall Compliance	The overall policy compliance using the following values and flags: • **All Compliant** —Icon: Green. indicates that all selected equipment is in compliance with the policy. • **None Compliant** —Icon: Red. indicates that none of the selected equipment is in compliance with the policy. • **None Determined** —Icon: blank. indicates that none of the equipment was tested for compliance. • **Partial Compliance** —Icon: Yellow. indicates that not all equipment complies with the policy but all equipment was tested. • **Compliance Varies**—Icon: Yellow indicates that not all equipment was tested for compliance. The tested equipment might be compliant or not compliant. This field displays on the summary and expanded views by default.
Monitored	An indicator that shows whether the policy is monitored (checkmark) or not (X). See ProScan on page 404 in Performance Monitoring for details. This field displays on the summary and expanded views by default.
Scheduled	An indicator that the policy's execution is scheduled (Yes) or not (No) and whether the schedule has occurred. This field displays on the summary and expanded views by default. **Note:** You can execute a policy manually from the Managed Resources portlet by right clicking the targeted device, selecting Change Management, and then selecting one of the following options: • **Execute ProScan** to execute policies that target the device. • **Execute ProScan Policy** to execute a ProScan policy that is not already associated with the device or group. A selection screen appears where you can select a policy and either execute or schedule it.
Description	A detailed description of the policy/policy group. This field displays on the expanded view by default.
Enabled	An indicator that shows whether the policy is enabled (checkmark) or not (X). This field displays on the expanded view by default.
ProScan Type	An indicator that the policy is a single policy or a group policy. This field displays on the expanded view by default.
Input Source	The inputs source, such as Not Applicable, Current Label, and so on. This field displays on the expanded view by default.

Column	Description
Next Execution Date	The frequency in which to execute the scheduled action, such as each weekday, every three months, only at startup, only once, and so on. This field displays on the expandedview by default.
Icon	A graphical representation of the selected policy/rule.
Widgets	Additional information about the selected alarm, such as: • Reference Tree shows the connection between the alarm and its resource. • Compliance Policy Summary catalogs the compliance policy's history and lists the equipment scanned and an indicator that the run discovered equipment in (checkmark) or out (X) of compliance. • Compliance Policy Chart The Widgets field is available **only** from the expanded view.
Adaptive CLI	The Adaptive CLI script.
Config File Label	The configuration file label, such as Current.
Configured Date	The configuration date and time.

Pop-Up Menu

The Compliance Policies/ProScan pop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
New	Select either a new policy or group. Creating a new policy opens the Compliance Policy Editor, through which you can define one. See Creating or Modifying a Compliance Policy on page 486 for more information about the Editor. See Creating or Modifying Compliance Policy Groups on page 501 for the group editor. Provides access to the following options to create policies/policy groups: • Group opens the Creating New Compliance Policy Group window, where you define general properties, grouped policies, and grouped targets. See Creating or Modifying Compliance Policy Groups for editor details. • Policy opens the Creating New Compliance Policy, where you define a policy's general properties, targets, and compliance criteria.
Edit	Opens the selected policy or group for modification. See Creating or Modifying a Compliance Policy on page 486 for more information. See Creating or Modifying Compliance Policy Groups on page 501 for the group editor.
Execute Compliance	Initiates a compliance check for the selected policy. The Job Status window opens displaying each check as it occurs and its outcome as well as a Success/Failed when complete. This option is available only for compliance alarms.
Refresh Targets	Queries to check targets, particularly those in dynamic groups, are up-to-date. Best practice is to Refresh ProScan Targets before running a scan particularly if your network has changed since the last scan. You can also schedule this. See Schedules Portlet on page 155.
Modify Targets	Lets you modify and/or select target equipment for the policy.
Schedule	Configures a policy to run on a schedule.
Audit	Opens an Audit Viewer with the results of a selected policy's runs. This is one way to see the historical results of the Compliance Policy runs. Another is to consult the Compliance Policy Summary widget in the Compliance Schedules.
Delete	Removes the selected compliance policy/group. **Caution:** This can impact anything that refers to what you are deleting.

Menu Option	Description
Import/Export	Provides the following actions when available for the selected image: • Import retrieves a file containing XML compliance descriptions. Some imports can come from a URL. • Export Selection exports the selected description to an XML file. • Export All exports all descriptions to an XML file. Click Download Export File to specify where to save the file. The Import/Export option is useful as a backup or to share descriptors with other projects. You must import data into the correct portlet. For example, you cannot import event data into the Actions portlet. If one type of data depends on another, you must import the other data before importing the data that depends on it.
Share with User	Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Compliance Remediation Actions

Use the Compliance Remediation Actions portlet to manage remediation actions. Compliance policies are used to scan a target for configuration that it expects to be present or check for configurations that it expects to not exist.

Youcan create a remediation action script that fix es a failed policy. A remediation action can apply Commands to a target or remove configuration from a target. Youcan manually execute the action or automate the corrective action so that when a policy failure event occurs, the event triggers the remediation action and brings the device into compliance.

The Compliance Remediation Actions portlet allows you to create a custom filter so that you only see actions that apply to compliance policies. Th e current portlet configuration will display all available actions. The easiest way to filter is by Name. Youshould have a naming convention for yourusercreatedcompliance-relatedactions. Forexample, allnewcompliancerelatedactionsstart with “Remediation-”. Then create a portlet filter like:

```
Name begins with Remediation-
```

Once you apply the filter, only the relevant compliance actions are listed.

By default, access this portlet by selecting Compliance from the navigation bar. This portlet has both a summary viewand an expanded view. Each viewcould displaydifferent columns and has the same pop-up menu options available.

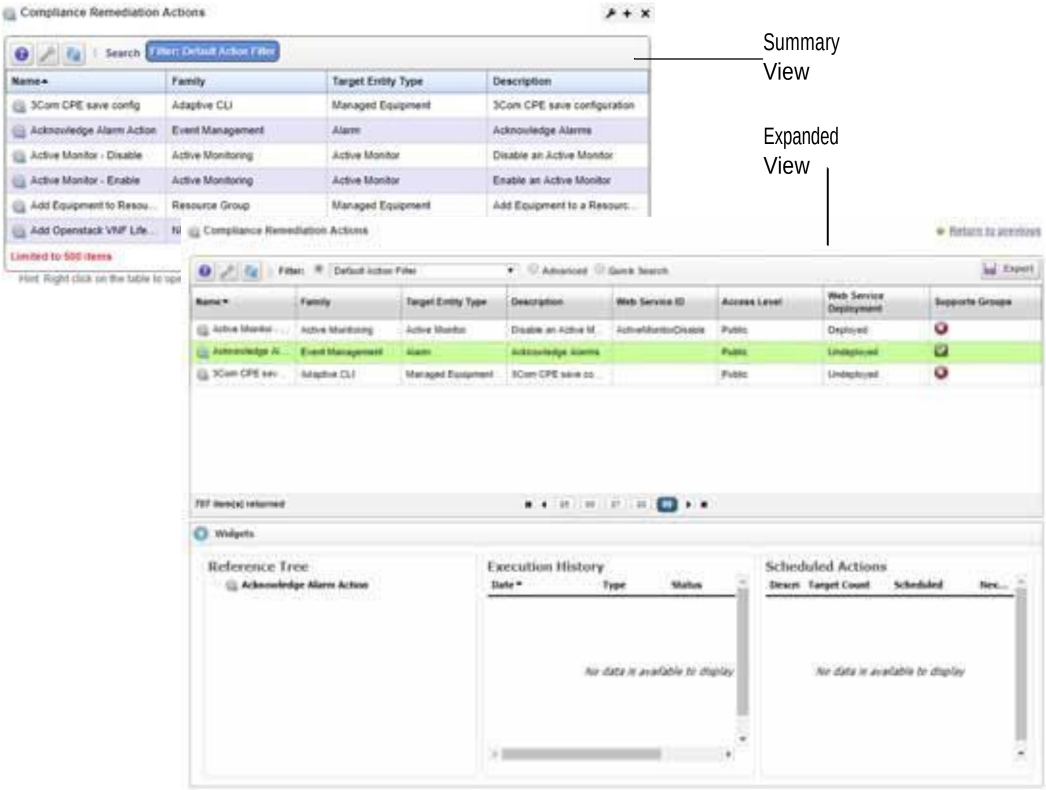

Name	Family	Target Entity Type	Description
3Com CPE save config	Adaptive CLI	Managed Equipment	3Com CPE save configuration
Acknowledge Alarm Action	Event Management	Alarm	Acknowledge Alarms
Active Monitor - Disable	Active Monitoring	Active Monitor	Disable an Active Monitor
Active Monitor - Enable	Active Monitoring	Active Monitor	Enable an Active Monitor
Add Equipment to Resou...	Resource Group	Managed Equipment	Add Equipment to a Resourc...
Add Openstack VNF Life...	Ni		

Name	Family	Target Entity Type	Description	Web Service ID	Access Level	Web Service Deployment	Supports Groups
Active Monitor ...	Active Monitoring	Active Monitor	Disable an Active M...	ActiveMonitorDisable	Public	Deployed	✗
Acknowledge Al...	Event Management	Alarm	Acknowledge Alarms		Public	Undeployed	✓
3Com CPE sav...	Adaptive CLI	Managed Equipment	3Com CPE save co...		Public	Undeployed	✗

787 item(s) returned

Columns

Other than the general navigation and configuration options, the Compliance Remediation Actions portlets (summary and expanded views) include the following columns. The columns displayed by default are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Name	The user-defined name for the policy or policy group. This field displays on the summary and expanded views by default.
Family	Pre-seeded grouping/categorization of actions based on the feature it supports.
Target Entity Type	The target type for the action, such as alarm, port, managed equipment, and so on. This field displays on the summary and expanded views by default.
Description	A detailed description of the remediation action. This field displays on the summary and expanded views by default.
Web Service ID	A unique identifier for the Web service. This field displays on the expanded view by default.
Access Level	The level of access for the action, such as public or private. This field displays on the expanded view by default.

Column	Description
Web Service Deployment	The Web service state (Deployed/Undeployed). This field displays on the expanded view by default.
Supports Groups	An indicator that shows whether the action supports groups (checkmark) or not (X). This field displays on the expanded view by default.
Widgets	Additional information about the selected alarm, such as: • Reference Tree shows the connection between the alarm and its resource. • Execution History. • Scheduled Actions The Widgets field is available **only** from the expanded view.
Component	This identifies the Redcell component that seeded the action. It will specify 'User' if the action was created by the user.
Data Type	The name of the data schema name for the action. The data schema defines the parameters for the action.
Domain Id	The identifier for the resource domain.
Icon	A graphical representation of the selected action.
Implementor Type	List this as internal use
Web Service Deployed	An indicator that shows whether the Web service is deployed (checkmark) or not (X).

Pop-Up Menu

The Compliance Remediation Actions pop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
New	Opens an editor where you specify genera l settings, action associations, attribute settings, and scripts for the following actions: • CLI Show Command • CLI Configure Command • Config File Generation • External Executable • Perl Script • RESTful Web Service See Adaptive CLI Editor on page 546 for a detailed description of the editors fields and options.
Copy	Provides access to the following functions for the selected action: • Copy Full copies the original schema to a new schema. • Copy reuses or references the original ACLI’s schema.
Execute	Initiates a compliance check for the selected policy. The Job Status window opens displaying each check as it occurs and its outcome as well as a Success/Failed when complete. This option is available only for compliance alarms.
Details	Shows more details for the selected action, such as activity details, execution history, scheduled actions, reference tree, and associated scripts.
View Scripts	Displays the scripts for the selected action.
History	Lists the selected action’s execution history results and execution details.

Menu Option	Description
Web Services	Allows you to deploy/undeploy the selected action. You also have the option to export the selected actions to a WSDL (Web Services Description Language) XML-based file.
Audit	Opens the Audit Trail Viewer.
Show Last Results	Shows the results for the last time you executed the selected action.
Schedule	Configures a policy to run on a schedule.
Add to Action Group	Opens a window where you select the actions you want to add to a group.
Import/Export	Retrieves a file containing XML compliance descriptions. Some imports can come from a URL. You must import data into the correct port let. For example, you cannot import event data into the Actions portlet. If one type of data depends on another, you must import the other data before importing the data that depends on it. **Note:** Use the Export button to save the current table to PDF, Excel, or CSV format. Click Export, select the format type, and then click Generate Export.
Share with User	Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Compliance Schedules

Use the Compliance Schedules to view and modify Compliance-specific schedules. You can use Schedules to initiate the Change Determination process. See Change Determination Process on page 506. Change Determination is disabled by default.

This portlet is intended for users who are interested in viewingand modifying Compliance-specific schedules.

By default, access this portlet by selecting Compliance from the navigation bar. This portlet has both a summary viewand an expanded view. Each viewcould display different columns and has the same pop-up menu options available.

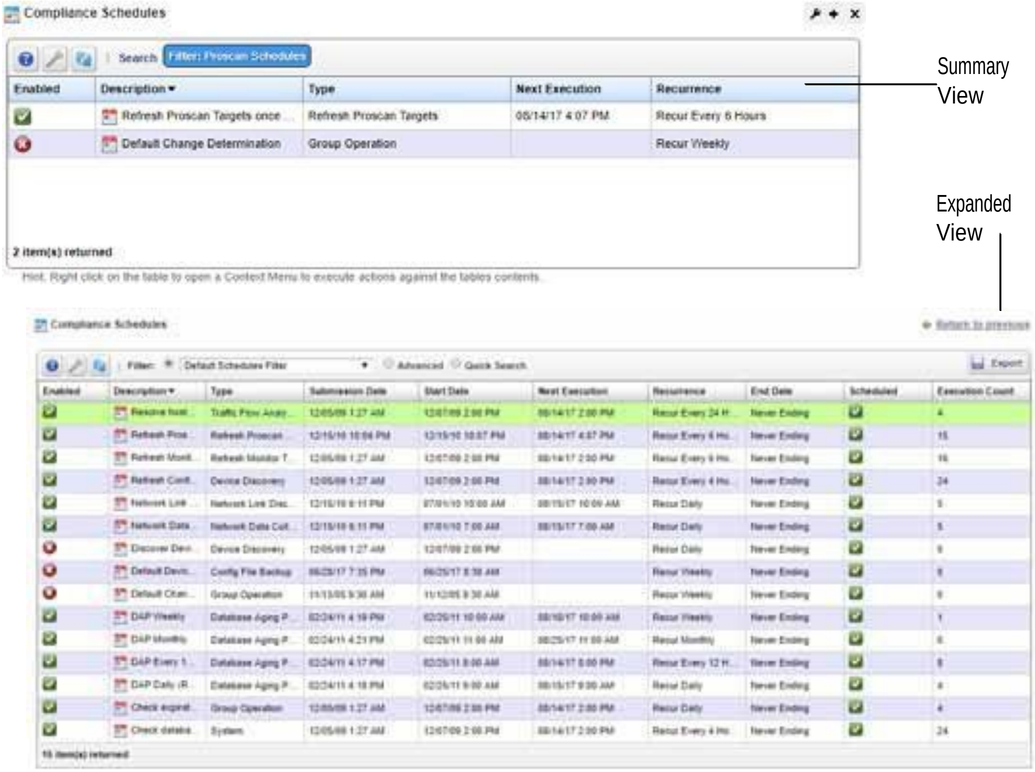

Enabled	Description ▾	Type	Next Execution	Recurrence
✓	Refresh Proscan Targets once ...	Refresh Proscan Targets	08/14/17 4:07 PM	Recur Every 6 Hours
✗	Default Change Determination	Group Operation		Recur Weekly

2 item(s) returned

Hint: Right click on the table to open a Context Menu to execute actions against the tables contents.

Enabled	Description ▾	Type	Submission Date	Start Date	Next Execution	Recurrence	End Date	Scheduled	Execution Count
✓	Resolve host...	Traffic Flow Analy...	12/05/08 1:27 AM	12/07/09 2:00 PM	08/14/17 2:00 PM	Recur Every 24 H...	Never Ending	✓	4
✓	Refresh Pros...	Refresh Proscan ...	12/15/10 10:06 PM	12/15/10 10:07 PM	08/14/17 4:07 PM	Recur Every 6 Ho...	Never Ending	✓	15
✓	Refresh Monit...	Refresh Monitor T...	12/05/08 1:27 AM	12/07/09 2:00 PM	08/14/17 2:00 PM	Recur Every 6 Ho...	Never Ending	✓	16
✓	Refresh Conf...	Device Discovery	12/05/08 1:27 AM	12/07/09 2:00 PM	08/14/17 2:00 PM	Recur Every 4 Ho...	Never Ending	✓	24
✓	Network Link ...	Network Link Disc...	12/15/10 6:11 PM	07/01/10 10:00 AM	08/15/17 10:00 AM	Recur Daily	Never Ending	✓	5
✓	Network Data...	Network Data Coll...	12/15/10 6:11 PM	07/01/10 7:00 AM	08/15/17 7:00 AM	Recur Daily	Never Ending	✓	5
✗	Discover Devi...	Device Discovery	12/05/08 1:27 AM	12/07/09 2:00 PM		Recur Daily	Never Ending	✓	9
✗	Default Devic...	Config File Backup	06/23/17 7:35 PM	06/25/17 8:38 AM		Recur Weekly	Never Ending	✓	8
✗	Default Chan...	Group Operation	11/13/05 9:30 AM	11/12/05 9:30 AM		Recur Weekly	Never Ending	✓	9
✓	DAP Weekly	Database Aging P...	02/24/11 4:19 PM	02/26/11 10:00 AM	08/16/17 10:00 AM	Recur Weekly	Never Ending	✓	1
✓	DAP Monthly	Database Aging P...	02/24/11 4:21 PM	02/25/11 11:00 AM	08/25/17 11:00 AM	Recur Monthly	Never Ending	✓	6
✓	DAP Every 1...	Database Aging P...	02/24/11 4:17 PM	02/25/11 8:00 AM	08/14/17 8:00 PM	Recur Every 12 H...	Never Ending	✓	8
✓	DAP Daily (R...	Database Aging P...	02/24/11 4:18 PM	02/25/11 9:00 AM	08/15/17 9:00 AM	Recur Daily	Never Ending	✓	4
✓	Check expirat...	Group Operation	12/05/08 1:27 AM	12/07/09 2:00 PM	08/14/17 2:00 PM	Recur Daily	Never Ending	✓	4
✓	Check databa...	System	12/05/08 1:27 AM	12/07/09 2:00 PM	08/14/17 2:00 PM	Recur Every 4 Ho...	Never Ending	✓	24

15 item(s) returned

Columns

Other than the general navigation and configuration options, the Compliance Schedule portlets (summaryand expanded views) include the followingcolumns. Thecolumns displayed bydefault are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Enabled	An indicator that shows whether the schedule is enabled (checkmark) or not (X). This field displays on the summary and expanded views by default.
Description	A detailed description of the scheduled action, such as Network Data Collection, Default Device Config Backup, and so on. This field displays on the summary and expanded views by default.
Type	The type of action scheduled, such as Device Discovery, Config File Backup, Database Aging Policy, System, and so on. This field displays on the summary and expanded views by default.
Next Execution	The next date and time that the schedule will execute. This field displays on the summary and expanded views by default.

Column	Description
Recurrence	The frequency in which to execute the scheduled action, such as each weekday, every three months, only at startup, only once, and so on. This field displays on the summary and expanded views by default.
Submission Date	The date and time that this schedule was submitted. This field displays on the expanded view by default.
Start Date	The date and time to start the scheduled action execution. This field displays on the expanded view by default.
End Date	The date and time to end the scheduled action execution or the number of occurrences before ending this scheduled action. This field displays on the expanded view by default.
Scheduled	An indicator that shows whether definition is scheduled (checkmark) or not (X). This field displays on the expanded view by default.
Execution Count	The number of times this scheduled action executed. Useful if you specified an end date based on the number of occurrences. This field displays on the expanded view by default.
Domain ID	The identifier for the resource domain.
Run Status	Indicate the schedules current state, such as waiting.

Pop-Up Menu

The Compliance Schedules pop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
New	Opens the New Schedule window, where you define a schedule for the following option selected: • Action (see Scheduling Actions on page 587 for the details) • Alarm Suppression (See Managed Resources on page 177 for details) • Config File Backup on page443 • Config File Restore on page444 • Database Aging Policy (Implementing DAP on page 80 for more about DAP) • OS Image Deploy on page 445 The subsequent window’s fields and options depend on the action selected.
Edit	Opens the Editing Schedule (*ProScan Target*), where you modify the schedule settings.
Delete	Removes the selected schedule. A successful/unsuccessful confirmation message is displayed.
Disable Schedule	Deactivates a currently enabled schedule. This option displays if the selected schedule is currently enabled. You can also disable a schedule by right clicking it and then selecting the deselecting the Enable Schedule option from the Editing Schedule window.
Enable Schedule	Activates a currently disabled schedule. This option displays if the selected schedule is currently disabled. You can also enable a schedule by right clicking it and then selecting the Enable Schedule option from the Editing Schedule window.

Menu Option	Description
Execute	Runs the actions for the selected schedule. The audit viewer displays execution progress for activity-based or discovery-profile-base schedules. For all other types of scheduled actions, a message is displayed stating that the item was sent to the application server for immediate execution. Monitor it progress from the Audit Trail portlet (see Audit Trail Portlet on page 150 for more details).
Share with User	Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Compliance Alarms

Use the Compliance Alarms portlet to manage Compliance-related alarms. This portlet filters all other non-Compliance-related alarms.

This portlet is intended for users who are interested in managing ProScan-related alarms.

By default, access this portlet by selecting Compliance from the navigation bar. This portlet has both a summary viewand an expanded view. Each viewcould display different columns and has the same pop-up menu options available.

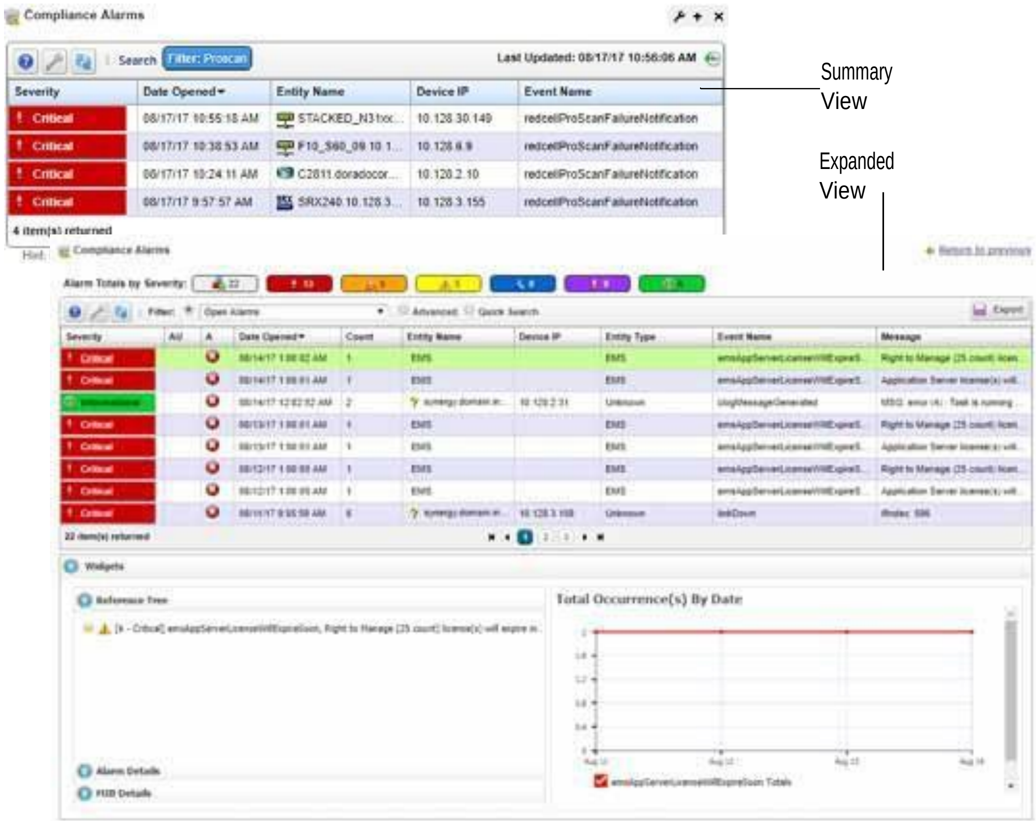

Last Updated: 08/17/17 10:56:06 AM

Severity	Date Opened	Entity Name	Device IP	Event Name
Critical	08/17/17 10:55:18 AM	STACKED_N3.txx...	10.128.30.149	redcellProScanFailureNotification
Critical	08/17/17 10:38:53 AM	F10_S60_09.10.1...	10.128.6.9	redcellProScanFailureNotification
Critical	08/17/17 10:24:11 AM	C2811.doradocor...	10.120.2.10	redcellProScanFailureNotification
Critical	08/17/17 9:57:57 AM	SRX240.10.128.3...	10.128.3.155	redcellProScanFailureNotification

4 item(s) returned

Severity	Ail	A	Date Opened	Count	Entity Name	Device IP	Entity Type	Event Name	Message
Critical			08/14/17 1:00:02 AM	1	EMS		EMS	emsAppServerLicenseWillExpireS...	Right to Manage (25 count) licen...
Critical			08/14/17 1:00:01 AM	1	EMS		EMS	emsAppServerLicenseWillExpireS...	Application Server license(s) will...
Informational			08/14/17 12:02:02 AM	2	synergy.domain.in...	10.128.2.31	Unknown	syslogMessageGenerated	MSG: error (4) : Task is running...
Critical			08/13/17 1:00:01 AM	1	EMS		EMS	emsAppServerLicenseWillExpireS...	Right to Manage (25 count) licen...
Critical			08/13/17 1:00:01 AM	1	EMS		EMS	emsAppServerLicenseWillExpireS...	Application Server license(s) will...
Critical			08/12/17 1:00:00 AM	1	EMS		EMS	emsAppServerLicenseWillExpireS...	Right to Manage (25 count) licen...
Critical			08/12/17 1:00:00 AM	1	EMS		EMS	emsAppServerLicenseWillExpireS...	Application Server license(s) will...
Critical			08/11/17 9:55:58 AM	6	synergy.domain.in...	10.128.3.100	Unknown	linkDown	ifIndex: 506

22 item(s) returned

[6 - Critical] emsAppServerLicenseWillExpireSoon, Right to Manage (25 count) license(s) will expire in...

Columns

Other than the general navigation and configuration options, the Compliance Alarms (summary and expanded views) include the followingcolumns. The columns displayed bydefault are noted.

Youcan view the value for most of the hidden columns by clicking the Settings tool, selecting the columns tab, clicking Show for the appropriate column, and then applying the change.

Column	Description
Severity	The alarm severity indicated by the color and text. The severity only has meaning for alarms and security alarms. Informational alarms have a severity level of Indeterminate. Closed alarms show without color. This field displays on the summary and expanded views by default.
Date Opened	The date the alarm was created. This field displays on the summary and expanded views by default.
Entity Name	The entity emitting this alarm. This field displays on the summary and expanded views by default.
Device IP	The equipment IP address where the alarm appeared. This field displays on the summary and expanded views by default.
Event Name	The event associated with the alarm. This field displays on the summary and expanded views by default.
Assigned User (AU)	The user currently assigned to this alarm. You can filter your assigned alarms using the Advanced filter option and selecting Assigned by User, contains, and entering your user ID. This field displays on the expanded view by default.
Acknowledged (A)	An indicator whether the alarm is acknowledged (checkmark) or not (X).
Count	The number of instances for this alarm. Multiples of the same alarm show as a single row but this value increments as more instances occur. This field displays on the expanded view by default.
Entity Type	The type of monitored entity. This field displays on the expanded view by default.
Message	The message associated with the alarm. This field displays on the expanded view by default.
Widgets	Additional information about the selected alarm, such as: • Reference Tree shows the connection between the alarm and its resource. • Alarm Details shows the severity, message, date opened, and so on. • MIB Details shows notification OID and MIB text. • Total Occurrences by Date shows a graph of the total occurrences by date. The Widgets field is available **only** from the expandedview.
Alarm State	The alarm’s open or closed state.
Date Cleared	The date and time the alarm was closed.
Updated Date/ Time	The time stamp for when the alarm was updated.
Notification OID	The identifier for the notification displayed as an alarm.
Equipment	The entity emitting the alarm.
Date Assigned	The date and time that the alarm was assigned.
Assigned By	The user ID for the person that assigned this alert.
Ack Time	The time that the alarm was acknowledged.
Cleared By	The user ID for the person that cleared the alarm.
MIB Text	The alarm’s MIB text.
Ack By	The user ID for the person that cleared the alarm.

Column	Description
Correlated By	The date and time when the alarm correlation to a parent alarm (caused by or blocked by).
Correlated Time	The role the alarm plays in any correlation, such as top-level alarm, caused by parent, blocked by parent.
Correlation State	The domain emitting the alarm.
Domain ID	The identifier for the resource domain.
Entity Description	The alarmed entity's description.
Has Children	An indicator that shows whether the alarm has children (checkmark) or not (X).
Highest Severity	The highest severity assigned to the selected alarm.
Location	The alarm's location.
Notes	A text field for information about the resource.
Original Severity	The original severity specified for the selected alarm.
Parent Alarm	The name of the parent alarm.
Region	The alarm's region.
Resource Propagation	The propagation for this alarm, if any.
Service Affecting	Indicates whether the alarm is on equipment in a provisioned service (checkmark) or not (X). A service affecting alarm propagate to show as a component of service- and customer-related alarms. Service-Affecting alarms are of indeterminate or greater severity.
Suppressed	An indicator whether the alarm is suppressed (checkmark) or not (X).
Suppression Date	The date and time of the alarm's suppression, if applicable.
Suppression End	The alarm' s suppression termination.
Suppressor	The alarm that suppresses the selected alarm.

Pop-Up Menu

The Compliance Alarms pop-up menu provides access to the following options. Right-click a row to access these options.

Menu Option	Description
Edit	Opens the Edit Alarm window or Editing Event Definition window, where you modify the alarm or event definition, respectively.
Details (Click + Shift)	Displays either equipment or alarm details, such as general information, MIB text, reference tree, advisory text, event processing rules, correlations, performance, and so on. The details displayed vary depending on the option selected (Alarm Details/ Equipment Details).
Topology	Displays a topology map that includes the selected alarms. See Presentation Capabilities for more about these maps.
Acknowledge Alarm	Acknowledges the selected alarms. The current date and time appear in the Ack Time field. The checkmark appears in the expanded portlet for acknowledged alarms. This option is available when the selected alarms **are not** acknowledged.
Unacknowledge Alarm	Unacknowledges previously acknowledged alarms and clears the entries in the Ack By and Ack Time fields. When an alarm is not acknowledged, the red icon appears in the expanded portlet. This option is available when the selected alarms **are** acknowledged.

Menu Option	Description
Assign User	Assigns this alarm to one of the users displayed in the sub-menu by selecting that user. An icon also appears in the expanded portlet indicating the alarm is assigned to someone.
Clear Alarm	Removes the alarm from the default al arm view and marks it as a candidate for the database archiving process (DAP). Clearing an alarm is an indication to the system that the alarm was resolved/addressed. If your system has propagation policies enabled, clearing recalculates dependent alarms.
Clear Group of Alarms	Removes a group of alarms from the default alarm view and marks them as candidates for DAP. This option is useful when you have many unimportant alarms that you want to clear at the same time instead of individually. For example, perhaps you want to clear all alarms that are informational and are more than a week old. Before selecting this menu item, create a filter for the group of alarms that you want to clear. When you select Clear Group of Alarms, a list of defined filters appears. Select the appropriate filter and then click Execute. All open alarms that meet the filter criteria are cleared. **Caution:** The Clear Group of Alarms action is irreversible.
Direct Access	Open an SNMP Mib Browser to the alarmed device, a CLI Terminal (Telnet window) to the alarmed device, or ICMP Ping the device alarmed. Only those available appear in the subsequent menu.
Email Alarm	Opens the Email Alarm window, where you specify a subject and e-mail address to which you want to mail the alarm’s content, and then click Send Email. Click Cancel to end this operation without sending e-mail. SMTP setup is required to e-mail an alarm. See SMTP Configuration on page 63 for instructions about setting up e-mail from Cruz.
Show Performance	Shows performance data for the selected interfaces in the Performance Dashboard window. Click the edit tool to modify the dashboard view properties, entities, or attributes. See Dashboard Editor on page 420 for more about the editor and its options.
Aging Policy	This lets you select a policy that determines how long this alarm remains in the database. See Implementing DAP on page 80 for information about configuring such policies.
Edit Custom Attributes	Opens the Custom Attribute Editor where you define field characteristics, such as whether it is enabled, the label name, and the tooltip.
View as PDF	Creates an Acrobat PDF document containing this alarm’s contents displayed in the summary portlet.
Share with User	Opens the Share with User window where you select the colleague you want to share this asset with and then type your message.

Compliance Policy Summary

This snap panel appears at the bottom of the expanded portlet described in Compliance Policy Summary on page 482. It catalogs the compliance policy’s history and lists the *Equipment* scanned, a status icon indicating whether the run discovered equipment *in* (green) or *out* (red) of compliance. If you added equipment to a policy before it has run, you may also see a *Not Executed* (blue) status. Each run date for the policy and equipment combination selected in the list at the top of the detail panel screen appears as a row in this panel. You can also see compliance failure messages in Cruz’s audit trails.

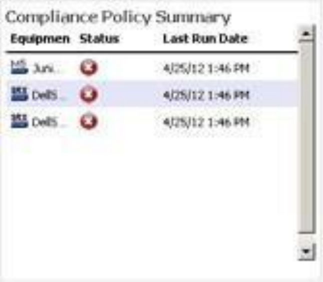

Compliance Policy Summary

Equipmen	Status	Last Run Date
Juni...	☒	4/25/12 1:46 PM
DellS...	☒	4/25/12 1:46 PM
DellS...	☒	4/25/12 1:46 PM

Compliance scans do not stop the first time they fail. They continue so all failures ofcompliance in the entire device configuration appear cataloged in the result.

Each time Cruz executes a compliance policy it st ores a history record in the database. Similarly, edits to these policies update history records. When you edit a compliance policy to add/remove equipment, Cruz creates or deletes the corresponding history record. Every time Cruz executes the compliance policy, it updates the Last Run Date, Status and Details on the history record.

Groups

When you run a Compliance Policygroup, the history for the group appears in this detail panel just as it would for a single policy. History concatenates the results of the component policies, as does reporting. See Compliance and Change Reporting on page 511.

Tosee the Compliance Policy History, print a *Compliance Policy Violation* report from Report Manager.

Using Change Management and Compliance

The following outlines common use cases for this software, and the steps to achieve the goals of each case:

Goal: Regularly verify configurations are compliant

1 Create compliance policy(ies) based on what indicates compliance. Right-click *New > Policy* in the Compliance Policies/ProScan portlet.

2 Specify the Name and Input source (based on Device Backup, Current Config, Configuration Label, By Date and Adaptive CLI Results)

3 Add Targets > Filter Option available for selecting Equipment/Group

NOTE:

Theadvantageofselectingdynamicdevicegroupsisthatnewlydiscovereddevicesoftheselectedtype automatically becomemembersof the group, so they are scanned too. Abenignwarning ("No compliancepolicieshave targetgroup(s)") letsyouknowyouhave notselected groups whenyou execute a compliance policy without them.

4 Specify Compliance Criteria. Add Criteria. For example, that devices' SNMPcommunities *Do not contain* the following:

```
snmp {
    community public {
```

5 Save.

6 Execute or schedule your created compliance policies.

7 Anyout-of-compliance devices throwan alarm, which you can email, or configure to trigger other actions (see the next use case).

Goal: If devices are not compliant restore compliant configuration

In addition to the steps in the previous section:

8 Create an event automation rule that responds to the redcellProScanFailureNotification event byexecuting the Netconfig Restore action. Typically, you would select to restore the *Compliant* label. (28493)

If you have multiple device types you do not need to assign actions for each device, or even each device type. Cruz supports the assigned policies, so it knows which actions to do to that device based on which device sent the trap.

Avoiding Restoring Files for Trivial Differences

Because automated network updates, for example from NTP servers, can change configurations, you may want to tune this to avoid restoring files that differ onlyinsignificantly. Do this byediting, or better overriding, the property file: ... `owareapps/netrestore/lib/nr.properties.` Alter the following property:

```
append.com.dorado.redcell.netrestore.backup.change.omit=
```

For example:

```
append.com.dorado.redcell.netrestore.backup.change.omit=,ntp clock-period
```

Configuring Compliance Policy Groups

If you have different compliance policies for different device type, then you can run a Compliance Policy Group and automatically scan even different types of devices. For more about this, see Creating or Modifying Compliance Policy Groups on page 501.

1 Right-click and select *New* > *Group*.

2 Specify the Compliance Policy Group Parameters.

3 Add Compliance Policies. These policies can be in multiple groups.

4 Add Targets. Notice that group targets appear in the "child" policies, grayed out. Child policies can add more targets.

5 Save.

6 Execute or schedule the group policies to run against the selected targets.

Do Change Management (Example)

The following describes an example use of Change Manager. This backs up a configuration file, modifies it, then scans the file for the modified text and acts according to the result. The following steps describe how to do this:

1 Back up a device configuration. Select a device and click the *File Management* > *Backup* right-click menu in Managed Resources portlet.

2 Right click and Export this backup to a file in the Configuration Files portlet.

3 Edit this config file, adding the word "MyTestContact" somewhere in its text that has no impact. For example, the snmp-server contact, or in comments. Some devices let you create descriptions within their configurations so you can enter a word without impact there.

4 Now import this edited file from the Managed Resources portlet after you have right-clicked on the same device from which you exported it. Renaming it something distinctive is helpful.

5 Right-click this file and *Restore* to the device. Since the name is a comment or description, it should not interfere with the device's operations.

6 Right-click the deviceand select *File Management* > *Backup.*This makes the MyTestContact file label Current.

To confirm MyTestContact is labeled Curren t, you can use an Advanced filter in the expanded Configuration Files portlet to view only Currentlabels.

7 Now,createacompliance policybyright clicking inthe Compliance Policies/ProScanportlet, selecting *New* > *Policy.*

8 In the General tab, name this policy MyTestContactScan, and as an input, select the *Configuration Label* > *Current* label as the Input Source.

9 In the Targets tab, select the equipment from which you exported the config file.

10 In the Criteria tab, click *Add Criteria* enter *contains* MyTestContact as the *Match All ofthe following criteria*.

11 Click *Save.*

12 Right-click the new policy and select *Execute Compliance*.

13 The audit screen that appears should indicate *Success.*

14 Right-click and *Open* the MyTestContactScan policy and change the Criteria to “does not contain” MyTestContact.

15 *Save*

16 Re-execute the policy.

17 The audit screen that appears should indicate *Failure.*

Alarms/Events

Once you have a compliance policy that has failed, the redcellProScan FailureNotification alarm appearsin the Alarms portlet. Success produces an event, not an alarm (visible in the Event History portlet) called redcellProScanClearNotification.

To create a response, create processing rules for the event/alarm (see Creating Event Processing Rules on page 292). For example, you could rest ore the Compliant-labeled configuration file if redcellProScanFailureNotification oc curs, or send an e-mail to a technician, among many other responses.

Some Limitations in this Example

Note that this example does not change authentication, either for telnet or SNMP. If it did alter the SNMP authentication, you would have to create an SNMP authentication alternative before scanning could occur.

Creating or Modifying a Compliance Policy

This series of screens lets you configure ProScan policies.

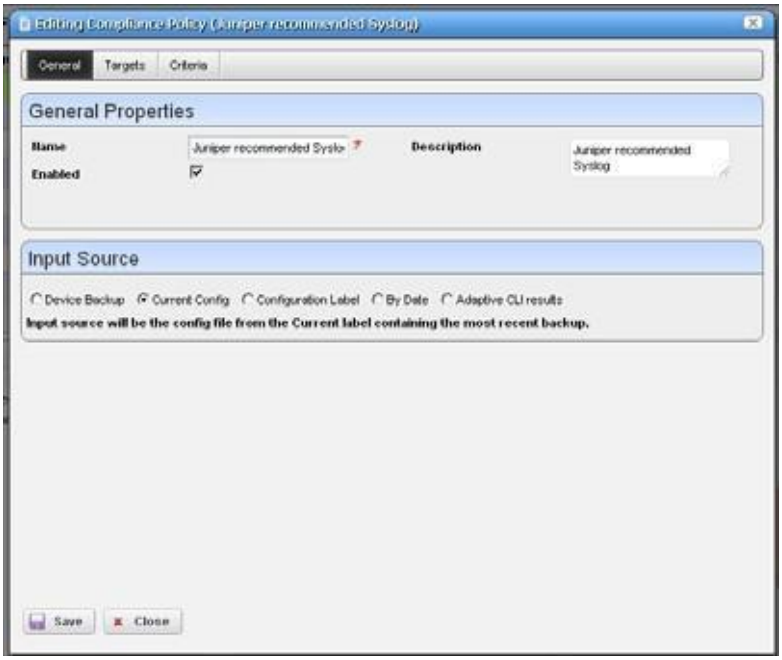

Name: Juniper recommended Syslo

Enabled: ☑

Description: Juniper recommended Syslog

○ Device Backup ● Current Config ○ Configuration Label ○ By Date ○ Adaptive CLI results

Input source will be the config file from the Current label containing the most recent backup.

This screen has the following tabs:

- General
- Targets
- Criteria

The Compliance Policy Job Status screen displays progress of a compliance policy as it executes.

△ CAUTION:
Compliance policies work only with text files; it does not work with binary configuration files.

If you have more than one type of device, you must typically have more than one compliance policy to address each device type. To run more than one compliance policy, so you can address multiple types of devices, create a compliance policy group. See Creating or Modifying Compliance Policy Groups on page 501.

General

This tab has the following fields:

General Properties

Name — An identifier for the policy (editable only when you click *New*, not on existing policies).

Enabled — Check to enable this policy.

Description—A text description of the policy. This also appears when the policy is listed in the manager.

Input Source

Use the radio buttons to select a source. Select from among the following options:

Device Backup— Retrieve the configuration from the device and scan it for compliance.

Current Config — The scan the current configuration backed up from the device.

Configuration Label—Select the configuration to run against based on a label. This software automatically updates the *Current* label, so it points to the most recently backed up configuration files.

By date—When you click this radio button, you can then select a configuration file backed up that precedes a specified date most closely in a selector that appears below the radio button. You can scan even historic configurations for compliance, with the *Based-on Date* field. No validation ensures this date is the current one.

Adaptive CLI —Select a desired *Show* Adaptive CLI to scan the target device below the radio button. The policy configured scans the show results, and that show appears in the Audit screen.

Targets

The top of this screen (*Current Inherited Targets)* displays any targets inherited from already-configured Compliance Policy Groups. Click *Add Targets* in the *Current Explicit Targets* panel at the bottom to select equipment that are targets to scan with this policy. You can also select listed equipment click the *Remove* icon to delete it from the list.

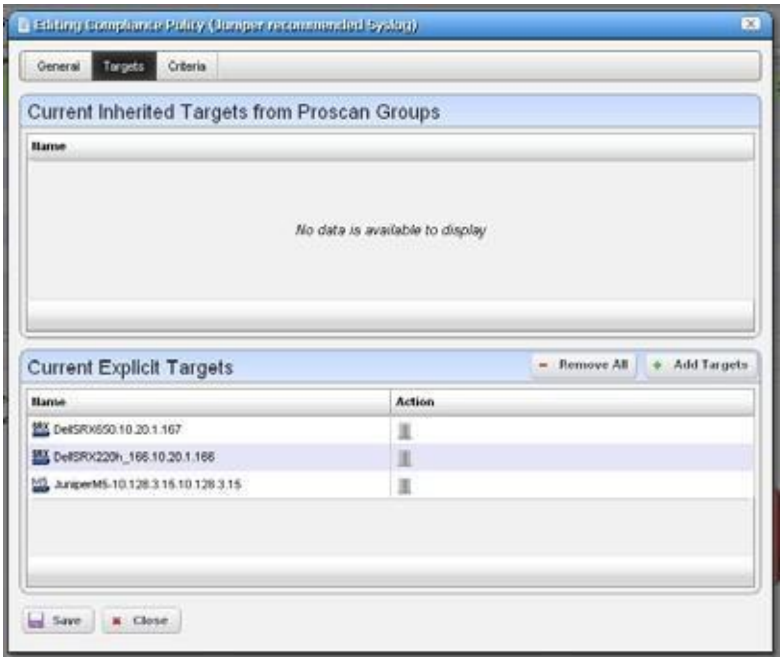

Current Inherited Targets from Proscan Groups

Name
No data is available to display

Current Explicit Targets

Name	Action
DellSRX650.10.20.1.167	
DellSRX220h_166.10.20.1.166	
JuniperM5-10.128.3.15.10.128.3.15	

NOTE:

Use filtering in the subsequent selector screen to make individual selection easier, but do not forget this is *not* dynamic selection. You must assign policies whenever your managed environment adds new equipment.

To provide information for individual policies that are part of groups, this screen displays inherited group targets grayed out. See Creating or Modifying Compliance Policy Groups on page 501 for more about groups.

Criteria

This screen lets you filter configuration files based on text, or Regular Expressions. Click *Add* to open an editor line.

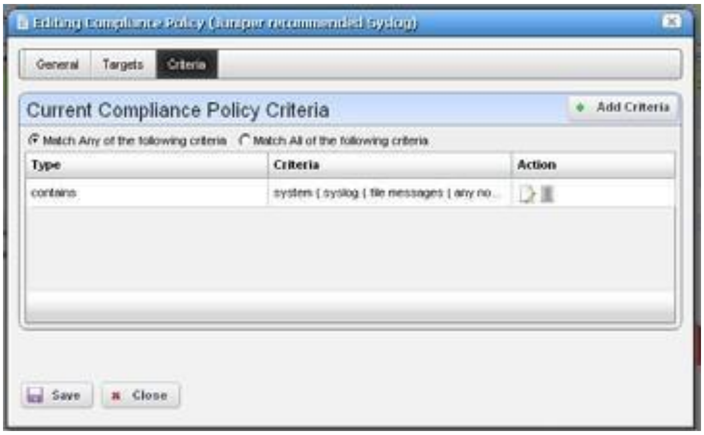
Editing Compliance Policy (Juniper recommended syslog)
General
Targets
Criteria
Current Compliance Policy Criteria
Add Criteria
Save
Close

● Match Any of the following criteria ○ Match All of the following criteria

Type	Criteria	Action
contains	system { syslog { file messages { any no...	

This screen ultimately determines whether the configuration file(s) for the selected equipment complies with the applicable policy. To create a policy, first select whether you want to *Match Any* (logical OR), or *All* (logical AND) of the criteria you configure with the radio buttons at the top of this screen.

See these sections for more about criteria:

- Editing Compliance Policy Criteria
- Match Regex for each line
- Count number of occurrences
- Input Source Grouping
- Properties

For additional criteria information consult these sections:

- Create Source Group Criteria
- Regular Expressions
- Perl/Java (Groovy) Language Policies

Editing Compliance Policy Criteria

After clicking *Add Criteria*, use the pick list on the upper right to select an operation to select a criteria match type (*Contains*, *Doesn't contain*, *[does not] match Regex* (see *Regular Expressions* on page 495*), [does not]* Match Regex for each line, Count number of occurrences, Perl or Java (Groovy)). Specify the match string or regular expression (Regex) in the text editor below the pick list.

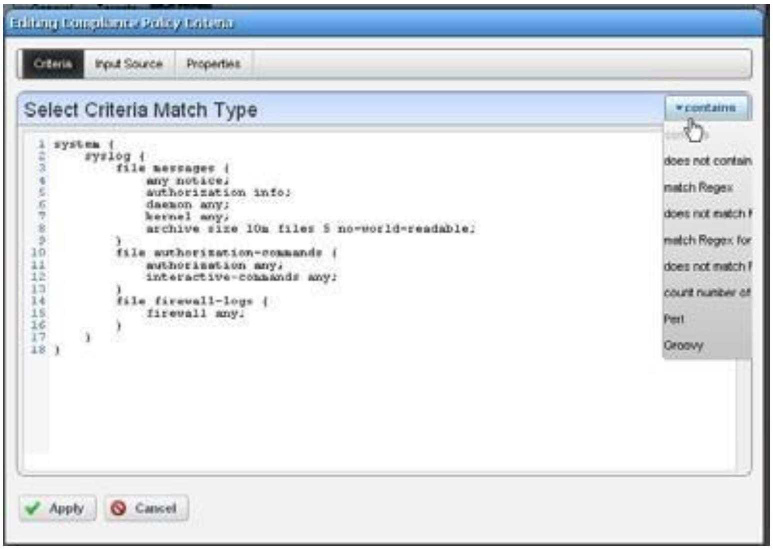

```
system {
    syslog {
        file messages {
            any notice;
            authorization info;
            daemon any;
            kernel any;
            archive size 10m files 5 no-world-readable;
        }
        file authorization-commands {
            authorization any;
            interactive-commands any;
        }
        file firewall-logs {
            firewall any;
        }
    }
}
```

With the *Add Criteria* button, you can configure multi-criteria policies with several lines. For example, configure one saying a maximum of four lines containing `name-server` can appear (<5), in any order (Match Regex for each line), and another that says the configuration must contain `no ip domain lookup [domain]`.

Notice the radio buttons *Match Any of the following* and *Match all of the following*. Selecting *Any* means that if either of the lines matched the policy would succeed. Selecting *All* says that both lines must pass before the policy is successful.

Formore complex scans, you can also enter Perl or Java (Groovy) language policies. See Perl/Java (Groovy) Language Policies on page 497 for details about these. The does not operators are just the negative of the match without does not.

Click the *Apply* green check button to accept your term, or the *Cancel* button to abandon your edits.

You can edit already listed compliance tests by clicking the *Edit* button (pencil and paper) in the list row. You can delete them by clicking the *Delete* button next to the criterion.

Match Regex for each line

In using this type of term, Cruz processes each line separately, comparing the input source to the match criteria. This returns a true value only if the criteria find a match in the source. The order of matching is not important since Cruz processes each lineseparately.

Count number of occurrences

This operator lets you specify a less than, greater than, or equal mathematical operator (< , > , =)

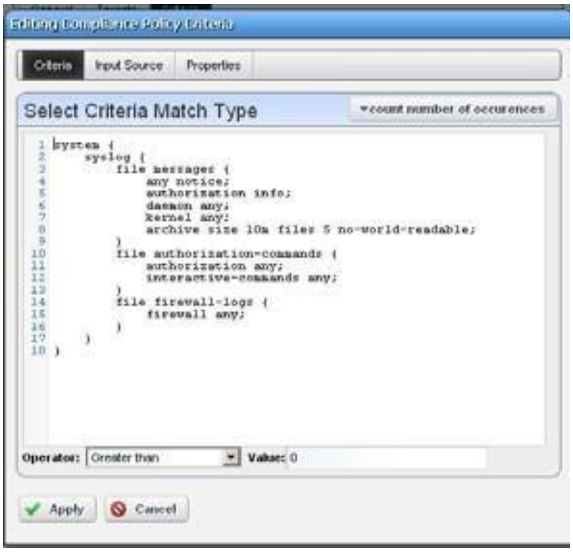
Editing Compliance Policy Criteria
Criteria
Input Source
Properties
Select Criteria Match Type
count number of occurences
Apply
Cancel

```
system {
    syslog {
        file messages {
            any notice;
            authorization info;
            daemon any;
            kernel any;
            archive size 10m files 5 no-world-readable;
        }
        file authorization-commands {
            authorization any;
            interactive-commands any;
        }
        file firewall-logs {
            firewall any;
        }
    }
}
```

Operator: Greater than

Value: 0

and several lines after you provide regex or string criteria with the operator and count value.

This returns true if the criteria (as a whole) match the input source count and operator combination. On the other hand, for example, if you choose a match criterion that includes =9 lines as the operator, and the scanned configuration has ten lines that match, the scan returns *false*.

Input Source Grouping

Adaptive CLI show commands and configuration files often have repeating sections or groups of parameters. Cruz scan configurations by section using *Start Criteria* and *End Criteria* Regex group criteria patterns. A configuration can contain multiple start and stops. This is especially useful when the criteria provided might occur multiple times in the input source, but you want to find only the instances which are preceded by a line in the source.

Click *Add new group* in the *Input Source* panel in the Criteria editor, and the grouping editor appears. (Click the red icon to the source grouping's left to delete it.)

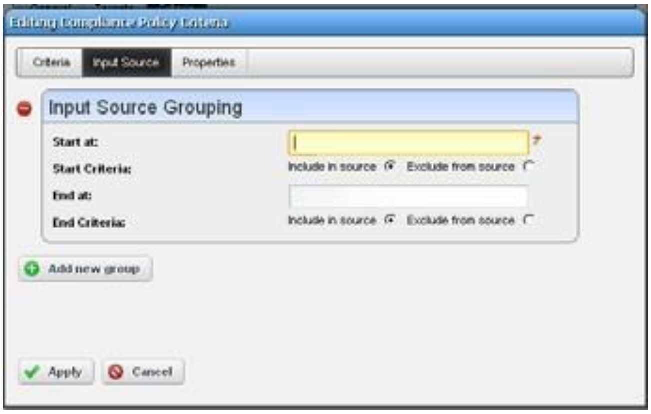

Start at:

Start Criteria: Include in source ● Exclude from source ○

End at:

End Criteria: Include in source ● Exclude from source ○

Enter the starting and ending regular expressions (*Start at/End at*) and elect whether the beginning or end of the source group includes or excludes what that expression matches. Click *Apply* to accept your edits or *Cancel* to abandon them. You can create multiple group criteria. Cruz applies the group criteria in order, from top to bottom.

When you have defined a *Start* and *Stop*, Cruz finds the information between these. Cruz logically extracts the data from the main config (essentially creating sections) and then does the audit.

Forexample, if your configuration has one section of *router bgp* and multiple sections for each bgp neighbor, you can specify matches within each neighbor. Your policy can audit each router bgp section and each neighbor within each router bgp.

See Create Source Group Criteria below for an ex ample of how to use these capabilities. Also, see Regular Expressions below for more about what match criteria are supported.

Properties

Checkboxes on this page configure whether the compliance policy match is *Case Sensitive* or has *Multi-Line Support.* By default, they are disabled. Check to enable them. If (upper/lower) case matters in what you are scanning for, check *Case Sensitive.* If you want to scan for a regular expression that spans more than one line, check *Multi-Line Support.* Lines do not have to be consecutive. For example, you could scan for hostname [line(s) intervene]ipv6.

Create Source Group Criteria

Here is an example of how you can use source group criteria. Suppose you want to scan for the following text:

```
neighbor 2.3.4.5 activate
neighbor 2.3.4.5 route-map allanRM01
```

This is within the fo llowing configuration:

```
router ospf 888
 log-adjacency-changes
 redistribute bgp 88 metric 10010 metric-type 1 subnets tag 334 route-map
 allanRM02
 network 2.3.4.0 0.0.0.255 area 123
 network 2.3.5.0 0.0.0.255 area 124
 network 2.3.6.0 0.0.0.255 area 125
!
router isis
!
router rip
 version 2
 network 175.92.0.0
 no auto-summary
!
address-family ipv4 vrf VPN_PE_A
no auto-summary
 no synchronization
 exit-address-family
```

```
!
router bgp 88
 bgp log-neighbor-changes
 neighbor 2.3.4.5 remote-as 22
 neighbor description "This is Test"
 neighbor test-parameter xxx
 neighbor 4.5.6.7 remote-as 66
 neighbor description "This is Test"
neighbor test-parameter xxx
 !
 address-family ipv4
 redistribute connected route-map map-12
 redistribute static route-map hjlhjhjhjk
 redistribute ospf 888 metric 500 match internal external 2 nssa-external
  1 nssa-external 2 route-map allanRM03
 neighbor 2.3.4.5 activate
 neighbor 2.3.4.5 route-map allanRM01 in
 neighbor 4.5.6.7 activate
 neighbor 4.5.6.7 route-map allanRM02 in
 default-information originate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf VPN_PE_A
 redistribute ospf 10 vrf VPN_PE_A match internal external 1 external 2
 no auto-summary
 no synchronization
 exit-address-family
!
```

In addition, within this configuration, you want to check if the target lines are present under each address-family in the *router bgp* section. To scan for this, follow these steps:

1 Select the *Match All of the following* radio button and enter both above lines as match criteria. Select the *Config Term* as *match Regex for each line*, so the order in which these lines appears does not matter.

2 *Add* a source group criterion to search for a sect ion that begins with “routers bgp”—in regex: `routers\sbgp`. No end match criterion is needed. Click *Apply*.

3 Click *Add* to make another criterion. This time, the start is `address-family\s`, and the end is `exit-address-family`. Click *Apply*.

4 You should see both criteria listed in the editor

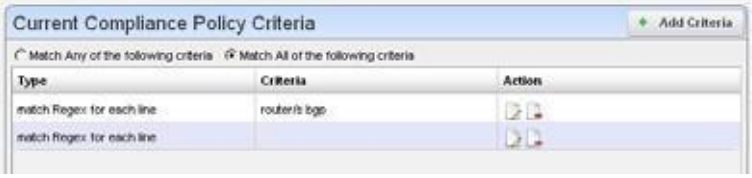

○ Match Any of the following criteria ● Match All of the following criteria

Type	Criteria	Action
match Regex for each line	router\s bgp	
match Regex for each line		

5 Applying the first group criterion finds the match (underlined) in the following:

```
router bgp 88
 bgp log-neighbor-changes
 neighbor 2.3.4.5 remote-as 22
 neighbor description "This is Test"
neighbor test-parameter xxx
neighbor 4.5.6.7 remote-as 66
neighbor description "This is Test"
neighbor test-parameter xxx
 !
 address-family ipv4
 redistribute connected route-map map-12
 redistribute static route-map hjlhjhjhjk
 redistribute ospf 888 metric 500 match internal external 2 nssa-external
  1 nssa-external 2 route-map allanRM03
neighbor 2.3.4.5 activate
neighbor 2.3.4.5 route-map allanRM01 in
 neighbor 4.5.6.7 activate
 neighbor 4.5.6.7 route-map allanRM02 in
 default-information originate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf VPN_PE_A
 redistribute ospf 10 vrf VPN_PE_A match internal external 1 external 2
 no auto-summary
 no synchronization
 exit-address-family
!
```

6 Applying the second group criterion on the above result divides the source:

Source 1:

```
address-family ipv4
 redistribute connected route-map map-12
 redistribute static route-map hjlhjhjhjk
 redistribute ospf 888 metric 500 match internal external 2 nssa-external
  1 nssa-external 2 route-map allanRM03
```

```
neighbor 2.3.4.5 activate
neighbor 2.3.4.5 route-map allanRM01 in
 neighbor 4.5.6.7 activate
 neighbor 4.5.6.7 route-map allanRM02 in
 default-information originate
 no auto-summary
 no synchronization
exit-address-family
```

Source 2:

```
address-family ipv4 vrf VPN_PE_A
 redistribute ospf 10 vrf VPN_PE_A match internal external 1 external 2
 no auto-summary
 no synchronization
 exit-address-family
```

This creates two sources sections.

7 Now Cruz applies the regex in the criteria field to each of the sources. It returns *true* only if both sources pass (we selected the *Match All* radio button). In this case “Source 2" does not have those lines, so Cruz returns a false value.

8 The error details appear in the audit trail panel.

Regular Expressions

Regular expressions include metacharacters to in struct the program how to treat characters it encounters. These include the following: ^, $, . , | , { , } , [,] , (,), *, +, ? , \. If you want to match one of these metacharacters, you must prepend a backslash (\). So to match a literal question mark, rather than instructing regular expression matching to match 0 or 1 of a previous expression, you must enter \?.

The following table outlines standard, supported regular expressions.

Label	Pattern
Single digit	\d
Two digits	\d{2}
Three digits	\d{3}
Four digits	\d{4}
Five digits	\d{5}
Number	[0-9] + One or more [0-9]* Zero or more
Decimal	.[0-9] +
Float	[0-9] + .[0-9] +
IP Address	(\d{1,3}.){3}\d{1,3}
IP Address/Mask	(\d{1,3}.){3}\d{1,3}/\d +
Domestic phone number with extension	1?[\s\-\/\.]*\(?([1-9]\d{2})\)?[\s\-\/\.]*([0-9]{3})[\s\-\/\.]*([09]{4})[\s\-\/\.x]*([0-9]{3,4})?
MAC Address	([0-9a-fA-F]{1,2}:){5}[0-9a-fA-F]{1,2}

Label	Pattern
MAC Address	([0-9a-fA-F]{1,2}.){5}[0-9a-fA-F]{1,2}
MIB2 OID	(1.3.6.1.6.1.2.1.(\d +\.) +\d)
Enterprise OID	(1.3.6.1.4.1.(\d +\.) +\d)
Time	[0-1][0-3]:[0-5][0-9]:[0-5][0-9]
All	.*
Ending Number	\d +$
Character	\w
Word	\w+ One or more. \w* Zero or more.
Whitespace	\s +One or more. \s* Zero or more.
String w/o space	\S + One ormore. \S* Zero or more.
New Line	\n
FormFeed	\f
Tab	\t
Carriage Return	\r
Backspace	\b
Escape	\e
Backslash	\B
URL	(?: ^ | ")(http |ftp |mailto):(?://)?(\w +(?:[\.:@]\w +)*?)(?:/ |@)([^"\?]*?)(?:\?([^\?"]*?))?(?:$ | ")
HTML Tag	<(\w +)[^ >]*? >(.*?) </\1 >

Here are some examples of such expressions:

Label	Pattern
Email address (U.S.)	^[A-Za-z0-9._%+-] +@[A-Za-z0-9.-] +\.[A-Za-z]{2,4}$
MAC Address	([0-9a-fA-F]{1,2}:){5}[0-9a-fA-F]{1,2}
Time hh:mm:ss	(0[0-9] | 1[0-2]):[0-5][0-9]:[0-5][0-9]
IP Address	(\d{1,3}.){3}\d{1,3}
Validated IP Address (restricts what matches better than the previous example)	(25[0-5] |2[0-4][0-9] | [01]?[0-9][0-9?])\.(25[0-5] |2[0-4][0-9] | [01]?[0-9][0-9?])\.(25[0-5] |2[0-4][0-9] |[01]?[0-9][0-9?])\.(25[0-5] |2[0-4][0-9] | [01]?[0-9][0-9]?)
MIB2 OID	(1.2.6.1.6.1.2.1.(\d +\.) +\d

The following are examples of the kinds of matching possible:

△ CAUTION:

Cutting and pasting from notepad into Cruz may cause carriage return or line-feed issues. Best practice is to compose these within Cruz.

Simple (Cisco ACL)

To match the following rows in a Cisco ACL:

```
access-list 159 permit icmp any any
access-list 159 permit tcp any any eq smtp
access-list 159 permit tcp any any eq www
```

To match these lines, simply create a compliance policy for *Config Term contains* (line contents) for each line.

Complex (Juniper)

When you have a multi-line statement to match, with varying elements, regular expressionsare necessary. For example:

```
lab@MyServer# show protocols
bgp {
  group internal {
       type internal
       export nhs
       neighbor 10.1.1.1
  }
}
```

In the above statement, the goal is to ensure an export policy in the BGP group internal called *nhs.* A suggested regex expression to match with the goal:

```
bgp/s+{/n/s+group/s+internal/s+{/n/s+type/s+internal;/n/s+export/s+nhs
```

NOTE:

Make sure you check Multi-line Support.

Another example:

```
lab@MyServer# show policy-options
policy-statement nhs {
  term set-nhs{
        then {
             next-hop self;
             }
  }
```

}

The following regex statement matches this example:

```
policy-statement\s+ns\s+{\n\s+term\s+set-nhs\s+{\n\s+then\s+{\n\s+next-
  hop\s+self
```

Perl/Java (Groovy) Language Policies

In addition to regular expressions, you can enter Config Terms that use either Perlor Java (Groovy) language capabilities for scans. The following sections describe these.

- Perl
- Java (Groovy)

These scans are compiled at runtime, and the Java scan uses the Groovy libraries, included with Cruz. You may need to install Perl on Windows application servers if you want to use that type of Config Term (it often comes with other supported operating systems). Refer to the *Cruz Installation Guide* for upgrading Perl instructions.

Perl

When you select Perl as the type of Config term, an editor appears that lets you enter Perl scans.

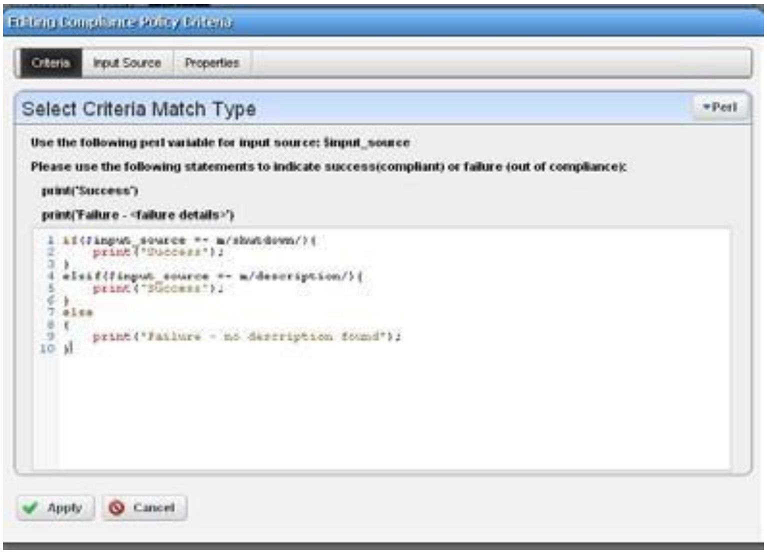
Editing Compliance Policy Criteria
Criteria
Input Source
Properties
Select Criteria Match Type
▾Perl
Apply
Cancel

Use the following perl variable for input source: $input_source

Please use the following statements to indicate success(compliant) or failure (out of compliance):

print('Success')

print('Failure - <failure details>')

```
if($input_source =~ m/shutdown/){
    print("Success");
}
elsif($input_source =~ m/description/){
    print("Success");
}
else
{
    print("Failure - no description found");
}
```

As the screen says `$input_source` is what the code scans. The following is example of the type of Perlyoucan enter that scans forcontents like `description`in shut down interfaces, and prints output “Success” visible in the Audit viewer when it finds a matching term like `description` in whatever source you select:

```
if($input_source =~ m/shutdown/){
    print("Success");
}
elsif($input_source =~ m/description/){
    print("Success");
}
else
{
    print("Failure - no description found");
}
```

Notice that you can also combine these scans with the *Edit Source Group Criteria* regular expressions to streamline them.

Java (Groovy)

When you select Groovy as the type of Config term, an editor appears that lets you enter that type of scans.

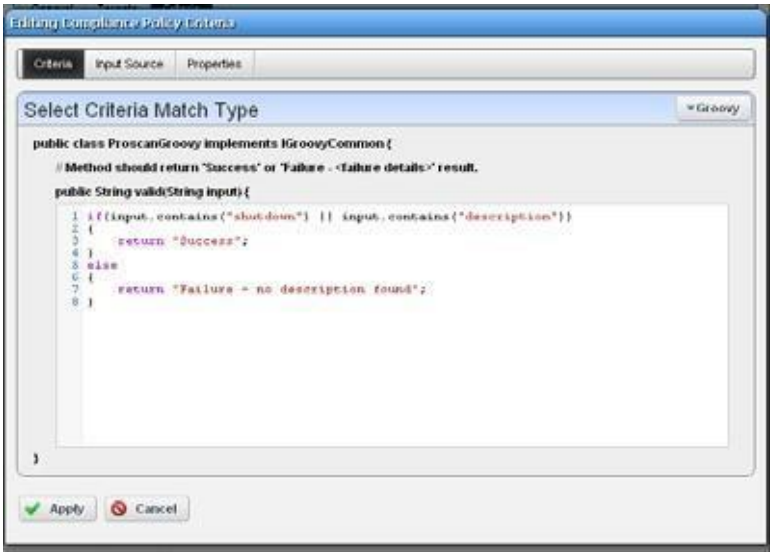
Editing Compliance Policy Criteria
Criteria
Input Source
Properties
Select Criteria Match Type
▾Groovy
✓ Apply
Cancel

```
public class ProscanGroovy implements IGroovyCommon {
    // Method should return 'Success' or 'Failure - <failure details>' result.
    public String valid(String input) {
if(input.contains("shutdown") || input.contains("description"))
{
    return "Success";
}
else
{
    return "Failure - no description found";
}
}
```

As the screen says this implements ProScanGroovy or Groovy Java classes. The method should return ‘Success or ‘Failure -’ results, and assumes `public String validate (String input) {` precedes what you enter in the text editor. The following is example of the type of Java code you can enter that scans for contents like `description` in shut down interfaces, and prints output “Success” visible in the Audit viewer when it finds a matching term like `description` in whatever source you select:

```
if(input.contains("shutdown") || input.contains("description"))
{
    return "Success";
}
else
{
    return "Failure - no description found";
}
```

NOTE:

Notice that you can also combine these scans with the *Edit Source Group Criteria* regular expressions to streamline them.

Click *Save* to preserve the policy you have configured in these screens or click *Close* (in the tool bar) to abandon your edits.

Compliance Policy Job Status

This screen displays the progress of compliance scanning you have configured.

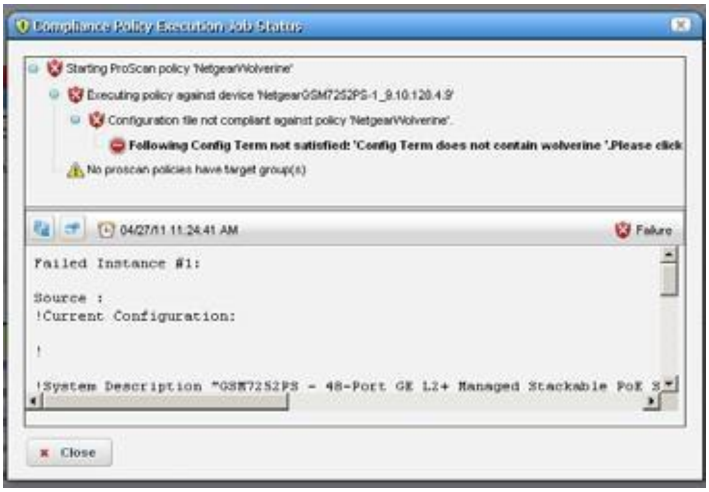

Starting ProScan policy 'NetgearWolverine'

Executing policy against device 'NetgearGSM7252PS-1_9.10.120.4.9'

Configuration file not compliant against policy 'NetgearWolverine'.

Following Config Term not satisfied: 'Config Term does not contain wolverine '.Please click

No proscan policies have target group(s)

04/27/11 11:24:41 AM

Failure

```
Failed Instance #1:

Source :
!Current Configuration:

!

!System Description "GSM7252PS - 48-Port GE L2+ Managed Stackable PoE S
```

You can revisit history of this policy's use in the Audit portlet (see Audit Trail Portlet on page 150). Select an audit trail in this portlet to review details.

When you see the *Success* indicator, then the scanned item is compliant. If you also see a warning message that no policies have target groups, this does not have an impact on compliance.

When you see the *Failure* indicator, then the scanned item is *Not* compliant. Select the "Following Config Term not satisfied" message to see the contents of the failed file at the bottom of this screen.

Executing compliance policies may trigger a benign warning that "No compliance policies have target group(s)." You can safely ignore this warning message.

The advantage of selecting dynamic device groups is that newly discovered devices of the selected type automatically become members of the group, so ProScan scans them too.

Creating or Modifying Compliance Policy Groups

When you create or modify a Compliance Policy Group after right clicking *New > Group* or *Open* when you have selected a group, the Compliance Policy Group editor appears.

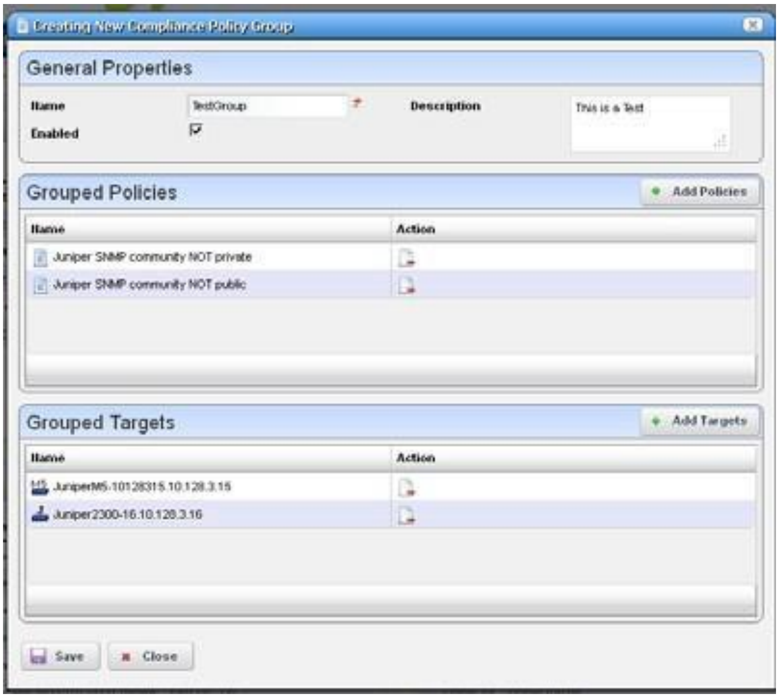

Name: TestGroup

Enabled: ☑

Description: This is a Test

Name	Action
Juniper SNMP community NOT private	
Juniper SNMP community NOT public	

Name	Action
JuniperM5-10128315.10.128.3.15	
Juniper2300-16.10.128.3.16	

This has the following to configure:

Name — A text identifier for the group.

Enabled — Check to enable this grouping.

Grouped Policies — Click *Add Policy* to select compliance policies in a selector screen. Click the *Remove* icon to delete a selected policy. Youcan use individual policies in several groups. Individual policies that are part of groups display inherited group targets grayed out.

Grouped Targets — Click *Add Targets* to select targets for the scans.

Executing a group executes all the member policies and update the historyrecords of the group and member policies. Any policy execution also updates its parent group history records.

Standard Policies

Change Managementcomeswithseveralpoliciesandactions bydefault. These includecompliance policies and policy groups, as well as the corresponding Actions for correcting any violations, and Event Processing Rules that automate remedy actions. The following sections briefly describe a representative set of these (you may have more or less, depending on your package).

- Cisco Compliance Policies
- Cisco Compliance Actions

- Cisco Event Processing Rules

△ CAUTION:
Seeded compliance policies are not necessarily correct by default. You must specify device targets at least. Given the variance in responses, particularly for Cisco devices, best practice is to test any such policy before you use it.

Cisco Compliance Policies

Thefollowingare Cisco Compliance policiesincluded bydefault with your Change Management installation. Policies listed here are part of Compliance Policy Groups scanning for PCI, HIPPA, SOX, NSA, and CISP compliance. These appear at the bottom of this list.

COMPLIANCE Cisco Enable Secret —Use enable secret for enable level access to device; PCI 8.4

COMPLIANCE Cisco Finger Service (12.1 +) — Disable Finger service; PCI 2.2.2

COMPLIANCE Cisco HTTP Server —HTTP server should not be running; PCI 2.2.2

COMPLIANCE Cisco Finger Service (11.3-12.0) —Disables finger service; PCI 2.2.2

COMPLIANCE Cisco Identd Service —Disable Identd service globally

COMPLIANCE Cisco Timestamps Logging — Use the timestamps service to show date and time on all log messages; PCI 10.2

COMPLIANCE Cisco Disable MOP — Disable MOP support on all Ethernet and VLAN interfaces; PCI.

COMPLIANCE Cisco NTP Redundant Servers —Ensures that more than one NTP server is defined; PCI 10.4

COMPLIANCE Cisco Disable NTP — Disable NTP if not in use; PCI 2.2

COMPLIANCE Cisco PAD Service — The packet assembler/disassembler (PAD) service supports X.25 links. This service is on by default, but it is only needed for devices using X.25; PCI 2.2.

COMPLIANCE Cisco Service Config —Disableautoloading ofconfiguration files fromaserver; PCI 2.2.2

COMPLIANCE Cisco Password Encryption — The password-encryption service shows user passwords as encrypted strings within the configuration; PCI 8.4

COMPLIANCE Cisco IP Source Route — Disable handling of source routed packets.

COMPLIANCE Cisco SNMP RW Communities — Do not use SNMP Read-Write strings, and only use Read-Only strings with associated access lists; PCI 2.2.3.

COMPLIANCE Cisco TCP Small-Servers (11.2-) —Disables unneeded TCP services such as echo, discard, chargen, etc; PCI 2.2.2

COMPLIANCE Cisco TCP Small-Servers (11.3+) —Disables unneeded TCP services such as echo, discard, chargen, etc; PCI 2.2.2

COMPLIANCE Cisco UDP Small-Servers (11.2-) —Disables unneeded UDP services such as echo, discard, chargen, etc.; PCI 2.2.2.

COMPLIANCE Cisco UDP Small-Servers (11.3 +) —Disables unneeded UDP services such as echo, discard, chargen, etc; PCI 2.2.2

COMPLIANCE Cisco VTY Exec Timeout — Set Exec Timeout on VTY ports; PCI 8.5.15

COMPLIANCE Cisco VTY Access Class Inbound —Set inbound access class on VTY ports; PCI 2.2.3.

COMPLIANCE Cisco VTY Login — Enable Login on VTY ports; PCI 2.2.3

COMPLIANCE Cisco VTY Transport Input Limit —Limit Input Transport on VTY ports; PCI 2.3

COMPLIANCE Cisco Set Login on Console Port — Enable login on console port; PCI 2.2.3

COMPLIANCE Cisco AAA Login — AAA login should be enabled; PCI 8.3

COMPLIANCE Cisco BOOTP Server — The BOOTSP server should be disabled; PCI 2.2.2

COMPLIANCE Cisco CDP Service — Disable CDP (Cisco Discovery Protocol) globally

COMPLIANCE Cisco Console Exec Timeout — Set an exec timeout console port; PCI 8.5.15

Cisco tacacs + enabled

Cisco monitor logging Enabled

Cisco console logging Enabled

Cisco buffered logging Enabled

Cisco SNMP Community String NOT public

Cisco SNMP Community String NOT private

Cisco RADIUS Enabled

Cisco Interfaces MUST have Description

Cisco Banner Enabled

Cisco ACL RFC 1918 space

Cisco ACL Permit Transit Traffic

Cisco ACL Permit RIP

Cisco ACL Permit OSPF

Cisco ACL Permit IGRP

Cisco ACL Permit EIGRP

Cisco ACL Permit BGP

Cisco ACL Deny access to internal infrastructure

Cisco ACL BGP AS Source

Cisco ACL Anti Spoofing

Cisco ACL - Deny special use address source

Cisco session-timeout' Enabled - ALL LINES

Cisco exec-timeout' enabled ALL LINES

Compliance Policy Groups

The following combine the Compliance Policies described above into groups to scan for compliance.

PCI Compliance for Cisco — This includes the following COMPLIANCE policies: Cisco VTY Transport Input Limit, Cisco VTY Login, Cisc o VTY Exec Timeout, Cisco VTY Access Class Inbound, Cisco SNMP RW Communities, Cisco Password Encryption, Cisco Finger Service (12.1 +), Cisco Finger Service (11.3-12.0), Cisco Disable NTP, Cisco Identd Service, Cisco AAA Login, Cisco UDP Small-Servers (11.2-)

HIPPA Compliance for Cisco —A policy group. This includes the following COMPLIANCE policies: Cisco VTY Transport Input Limit, Cisco VTY Login, Cisco SNMP RW Communities, Cisco Set Login on Console Po rt, Cisco Password Encryption, Cisco PAD Service, Cisco HTTP Server, Cisco Enable Secret, Cisco Timestamps Logging, Cisco NTP Redundant Servers, Cisco Finger Service (11.3- 12.0), Cisco Finger Service (12.1+), Cisco BOOTP Server, Cisco CDP Service.

SOX Compliance for Cisco — A policy group. This includes the following COMPLIANCE policies: Cisco VTY Login, Cisco VTY Transport Input Limit, Cisco SNMP RW Communities, Cisco Set Login on Console Po rt, Cisco Password Encryption, Cisco PAD Service, Cisco Finger Service (11.3-12.0), Cisco Finger Service (12.1+), Cisco HTTP Server, Cisco Identd Service, Cisco UDP Small-Servers (11.3 +).

NSA Compliance for Cisco —A policy group. This includes the following COMPLIANCE policies: Cisco VTY Login, Cisco VTY Transport Input Limit, Cisco SNMP RW Communities, Cisco VTY Exec Timeout, Cisco Service Config, Cisco Password Encryption, Cisco PAD Service, Cisco HTTP Server, Cisco Fi nger Service (12.1 +), Cisco Finger Service (11.3-12.0), Cisco Enable Secret, Cisco Disable MOP, Cisco Disable NTP, Cisco NTP Redundant Servers.

CISP Compliance for Cisco — A policy group. This includes the following COMPLIANCE policies: Cisco UDP Small-Servers (11.3+), Cisco VTY Transport Input Limit, Cisco VTY Login, Cisco VTY Exec Timeout, Cisco VTY Access Class Inbound, Cisco Password Encryption, Cisco Finger Service (12.1+), Cisco Finger Service (11.3-12.0), Cisco Enable Secret.

Cisco Compliance Actions

Remedial actions are often part of the process of change management. These may be triggered by the Cisco Event Processing Rules, and are included as part of the Standard Policies.

Compliance Cisco AAA Login —To avoid being locked out of the router, define username and password on the access server before starting the AAA configuration.

Compliance Cisco Finger Service — Disables the ip finger service.

Compliance Cisco HTTP Server — Disables http.

Compliance Cisco Identd Service — Disables identd

Compliance Cisco IP Source Route —Disables ip source route

Compliance Cisco UDP Small-Servers (11.3 +) — Disables PCI UDP Small-Servers (11.3 +).

Compliance Cisco TCP Small-Servers — Displace PCI Cisco TCP Small-Servers.

Compliance Cisco BOOTP Server — Disables PCI Cisco BOOTP Server.

Compliance Cisco PAD Service — Disables the PAD service.

Compliance Cisco Timestamps Logging — Enables PCI Cisco Timestamps Logging.

Compliance Cisco SNMP RW Communities —Removes RW community string with user input.

Compliance Cisco Password Encryption — Enables PCI Cisco Password Encryption.

Compliance Cisco CDP Service — Disables CDP Cisco Discovery Protocol.

COMPLIANCE Cisco VTY Transport Input Limit

COMPLIANCE Cisco VTY Login

COMPLIANCE Cisco VTY Exec Timeout

COMPLIANCE Cisco VTY Access Class Inbound

COMPLIANCE Cisco Set Login on Console Port

COMPLIANCE Cisco Service Config

COMPLIANCE Cisco SNMP RW Communities

COMPLIANCE Cisco Password Encryption

COMPLIANCE Cisco PAD Service

COMPLIANCE Cisco NTP Redundant Servers

COMPLIANCE Cisco Enable Secret

COMPLIANCE Cisco Disable NTP

COMPLIANCE Cisco Disable MOP

COMPLIANCE Cisco Console Exec Timeout

Cisco Event Processing Rules

Theevent processingrules here typically tie Cisco Compliance Policieswith remedial Cisco Compliance Actions.

Compliance Cisco AAA Login Remediation — Triggers a task to configure an AAA login.

Compliance Cisco BOOTP Server —Corrects PCI Cisco BOOTP Server compliance failures.

Compliance Cisco CDP Service—Corrects PCI Cisco CDP Service compliance failures.

Compliance Cisco Finger Service — Corrects PCI Cisco Finger Service compliance failure.

Compliance Cisco HTTP Server — Corrects http server compliance failures.

Compliance Cisco Identd Service — Corrects PCI Cisco Identd Service compliance failures.

Compliance Cisco IP Source Route —Corrects PCI Cisco IP Source Route compliance failures.

Compliance Cisco PAD Service — Corrects PCI Cisco PAD Service compliance failures.

Compliance Cisco TCP Small-Servers—Corrects PCI Cisco TCP Small-Servers compliance failures.

Compliance Cisco Timestamps Logging—Corrects PCI Cisco Timestamps Logging compliance failures.

Compliance Cisco UDP Small-Servers (11.3 +) —

Juniper Compliance Policies

Packages that support Juniper devices have the following policies:

Juniper FW Filter Private IP — RFC 1918

Juniper Policer DNS — Protect from source address spoofing

Juniper Policer NTP — Protect from source address spoofing

Juniper Policer RADIUS—Protect from source address spoofing

Juniper Policer SNMP — Protect from source address spoofing

Juniper Policer SSH — Protect from source address spoofing

Juniper Policer Small BW—Protect from source address spoofing

Juniper Policer TCP—Protect from source address spoofing

Juniper Recommended Logging — Confirms recommended logging is on.

Juniper SNMP community NOT public — Checks the SNMP community is not “public” closing a potential security hole.

Juniper SNMP community NOT private — Checks the SNMP community is not “private” closing a potential security hole.

Juniper ALLServices Policy—*Note:* this compliance policywill typicallybe modified per deployment.

Juniper Recommended SSH — Confirms recommended SSH is on.

Juniper Recommended Syslog — Confirms recommended syslogging is on.

Change Determination Process

If you run the *Change Determination Process*, it collects all the configuration changes that occurred on the target resources since the last time this process ran. It also associates these changes with the date and time when the Change Determination process runs. After running Change Determination, you can then produce a report (see *Compliance and Change Reporting* on page 511), outlining all such changes by date and time. This report comes seeded with installation.

Cruz stores incremental changes as *RedcellConfigChangeRecords* by device/timestamp. The *ConfigChangeRecordsDAP* Database Aging Policy (DAP) manages howlong the Cruz database retains these records. This DAP’s default setting stores incremental records for 30 days, then archives or purges them. Reporting shows only records in the database; therefore, by default, the *Configuration Change Report* shows only resource changes made in the last 30 days, but no older. Change this default by changing the number of days to retain such records with the DAP.

The next section describes Change Determination Process Workflow.

Change Determination Process Workflow

Change Manager seeds Change Determination Process and compliance group operations. Youcan configure this to run on groups of your choosing if you create a new Change Determination Process group operation.

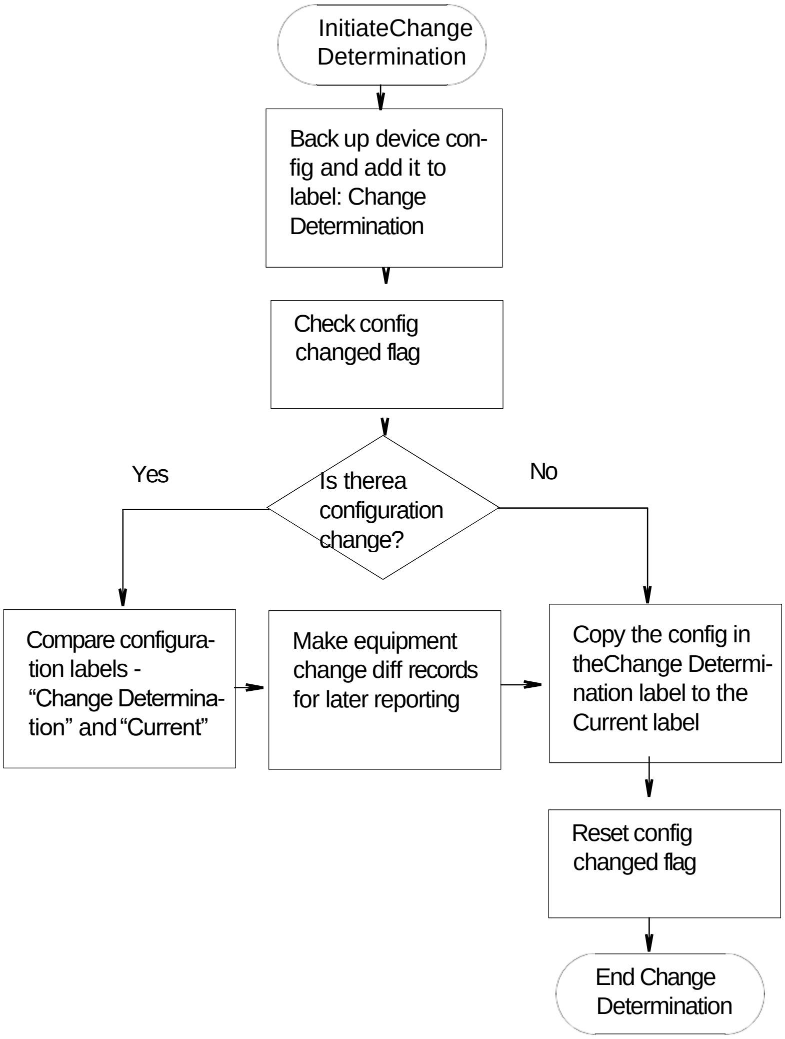
InitiateChange Determination
Back up device config and add it to label: Change Determination
Check config changed flag
Is therea configuration change?
Yes
No
Compare configuration labels - "Change Determination" and "Current"
Make equipment change diff records for later reporting
Copy the config in theChange Determination label to the Current label
Reset config changed flag
End Change Determination

This process records what is removed, updated or added since it last ran on a scanned device's configuration. If you run the Change Determination Process, it first backs up the devices' configuration(s), and stores those with the Change Determination label.

Change Determination Process then looks for Con fig Changed Flags, and if it finds such flags, indicating a change occurred on the device and/or Change Determination has not run on it, the process then compares the device's changed configuration (in the Change Determination label) to the one in the Current label, storing the difference for future reporting.

At its end, the Change Determination Process re-labels the configuration with the Change Determination label to the Current label, and it un-sets the Config Changed Flag on scanned resources so the flag will not signal change occurred when Change Determination runs again.

After running the Change Determination Process, you can run the Configuration Change report to display what changed for a defined period. The contents of that report depend on the report filter, and the specified period. This report lists changed attributes in theconfigurations.

Triggering Change Management and ProScan

To trigger the Change Management for a device, right-click it in the Managed Resources portlet and select *Change Management > Change Determination*. You can also schedule Change Determination to run repeatedly, on regular intervals in the Schedules portlet.

Youcan similarly trigger ProScan by right clicking a device and selecting *Change Management > Execute ProScan* or *Execute ProScan Policy*. The former executes all policies connected with the selected device, while the latter allows you to select policy (or policies) to run. Creating a Compliance Policy Group, lets you run all compliance policies for each device within the selected group, scanning groups even if they consist of de vices from different vendors. In ProScan, you can scan device configurations (of specified labels) or Adaptive CLI command output. (See How to: Using Change Management and Compliance on page 484).

How To:

Run the Change Determination Process

The following describes an exercise for the Change Determination process based on manually running it. To run the process as a response to events devices must transmit traps to Cruz. The next sections describe using Change De termination in the followingways:

- Change Determination Confirmation
- Event/Trap-Based Change Determination

Change Determination Confirmation

The following steps confirm change determination is working.

1. Initialize the Change Determination Report and let it do a configuration backup. The first time this runs, Cruz creates no diffs. It just initializes the Change Determination label.
2. Edit a configuration to make a change. For example, make a change in a device you have discovered. One benign change is to add a contact or a description to an interface.
3. Restore it to the device.
4. Execute the Change Determination process on the device by right clicking it in the Managed Resources Portlet, and selecting *Change Management > Change Determination* .

   This then backs up the device, compares the original and altered configurations, and writes the difference to report later (see How to: Report on Change Determination on page 513 for the steps to run the report to see such changes).

Since we have initialized the report in step 1, the updated report shows the changes made to the config file.

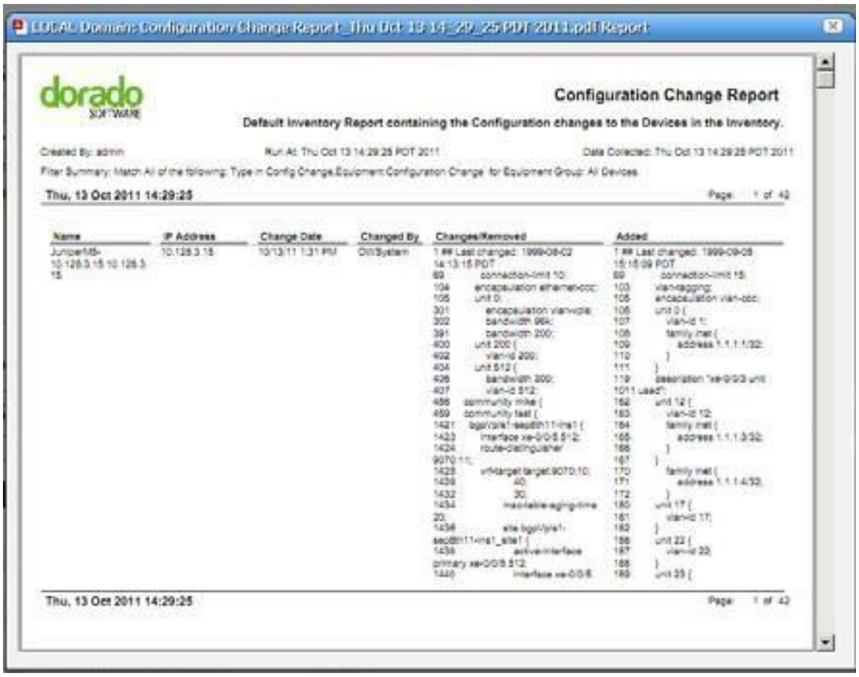
LOCAL Domain: Configuration Change Report_Thu Oct 13 14_29_25 PDT 2011.pdf Report
dorado
SOFTWARE

Configuration Change Report

Default Inventory Report containing the Configuration changes to the Devices in the Inventory.

Created By: admin   Run At: Thu Oct 13 14:29:25 PDT 2011   Data Collected: Thu Oct 13 14:29:25 PDT 2011

Filter Summary: Match All of the following: Type in Config Change.Equipment Configuration Change for Equipment Group: All Devices

Thu, 13 Oct 2011 14:29:25   Page 1 of 42

Name	IP Address	Change Date	Changed By	Changes/Removed	Added

JuniperMX-10.128.3.15 10.128.3.15	10.128.3.15	10/13/11 1:31 PM	OI/System	1 ## Last changed: 1999-08-02
				14:13:15 PDT
				69 connection-limit 10;
				104 encapsulation ethernet-ccc;
				105 unit 0;
				301 encapsulation vlan-vpls;
				302 bandwidth 96k;
				391 bandwidth 200;
				400 unit 200 {
				402 vlan-id 200;
				404 unit 512 {
				406 bandwidth 200;
				407 vlan-id 512;
				458 community mike {
				459 community test {
				1421 bgpVpls1-secdth11-ins1 {
				1423 interface xe-0/0/5.512;
				1424 route-distinguisher
				9070:11;
				1428 vrf-target target:9070:10;
				1430 40;
				1432 30;
				1434 mac-table-aging-time
				20;
				1436 site bgpVpls1-
				secdth11-ins1_site1 {
				1438 active-interface
				primary xe-0/0/5.512;
				1440 interface xe-0/0/5
					1 ## Last changed: 1999-09-08
					15:15:09 PDT
					69 connection-limit 15;
					103 vlan-tagging;
					105 encapsulation vlan-ccc;
					106 unit 0 {
					107 vlan-id 1;
					108 family inet {
					109 address 1.1.1.1/32;
					110 }
					111 }
					119 description "xe-0/0/3 unit
					1011 used";
					162 unit 12 {
					163 vlan-id 12;
					164 family inet {
					165 address 1.1.1.3/32;
					166 }
					167 }
					170 family inet {
					171 address 1.1.1.4/32;
					172 }
					180 unit 17 {
					181 vlan-id 17;
					182 }
					186 unit 22 {
					187 vlan-id 22;
					188 }
					189 unit 23 {

Thu, 13 Oct 2011 14:29:25   Page 1 of 42

5 Repeat step 2 through 5 if you like after you have made further changes.

NOTE:

Best practice in production is to schedule a recurring run for Change Determination in the Schedules portlet. Notice that you can also disseminate the report by e-mail, or view previous reports in the web client, as described in the Reports portion of this guide.

Event/Trap-Based Change Determination

The following steps to trigger Change Determination based on events received by Cruz. Your devices must transmit traps to the Cruz installation, and must emit traps when changes occur, or this does not work.

1 Back up the configuration file for a device you havediscovered.

2 Make a change to that device with the Manage d Resources editor, or from a Direct Access command line.

3 Such changes make the device emit an event that may have further consequences. For example, for Juniper devices, the Juniper JUNOS Configuration Changed event is a correlation event.

4 To provide a response (and to normalize the emitted event), create an automation rule that emits a `redcellEquipmentConfigChangeNotification` event when Cruz receives creates an event in response to events like the `jnxCmCfgChange` event that occurs when Juniper devices change.

5 Create a rule to respond to `redcellEquipmentConfigChangeNotification` by running the Change Determination process. Youdo not have to back up the configuration after the change. See How to: Create Event Processing Rules to Trigger Change Determination Process below.

6. To see the change itself, run the Change Determination Report (see Compliance and Change Reporting on page 511 and How to: Report on Change Determination on page 513). The report displays the changes made.

How To:

CreateEvent Processing Rules to Trigger Change Determination Process

Thisexercise createsan Event Processingrule that has Change Determination respond to anevent. The steps to configure such an event processing rule are as follows:

1. Create a new event processing rule by right clicking in Event Processing Rules > *New* > *Post Processing* rule in the Event Processing Rule portlet.
2. Enter the name in the field labeled *Name*. (Example: Update Config Change Flag)
3. Click *Next* to go to the *Filter* tab.
4. For the *Specify Events* panel, click on the *Add* button to select the event to which this rule responds. A selector listing available events appears.

NOTE:

Notice you can limit the selector's displayed events by entering text in the filter at the top of the selector screen.

5. In the selector, click the event definition (here: `redcellEquipmentConfigChangeNotification`), and confirm your selection.
6. Click *Done* to accept the *Event(s)* you have configured.
7. Notice you can further filter which events this ru le responds to with the lowest panel in this screen's *Filter Conditions* panel by clicking *Add Filter*. For example, you could create a rule that responds only to events from a particular IP address. For now, we will not configure additional filters.
8. Click *Next* to open the *Actions* tab.
9. Click *Add Action*, and click the *Custom* action alternative, then click Keyword Search and select Change Determination. That action appears in the drop-down combo box. Notice you can also select a target in the action selector. By not selecting one, we run change determination against all Managed Equipment.
10. Click *Apply* and view the Change Determination action listed in the Actions screen.

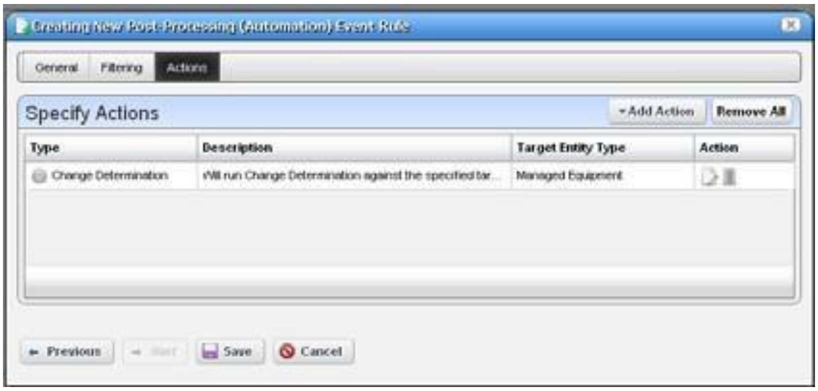

Creating New Post-Processing (Automation) Event Rule

Specify Actions

Type	Description	Target Entity Type	Action
Change Determination	Will run Change Determination against the specified tar...	Managed Equipment	

Notice that you can add more actions and edit or delete existing ones with the icons to the right. Click *Apply* once you have selected Change Determination.

11 Click *Save* to preserve this event processing rule. The rule should now respond to the configured event, triggering the action you configured.

NOTE:

Backup and Change Determination automates backing up target devices.
Also: Change Determination's current default is to compare files even if the "Config Change" flag has not been modified.

Change Determination Defaults

By default, Change Determination can run against all devices without requiring the config change update flag be set or updated based on events tied to the Update Config Change Flag event processing rule/action.

To disable the manual run-ability of the Change Determination process, uncomment the property in `\owareapps\changemgmt\lib\cm.properties`(or add it to `\owareapps\installprops\lib\installed.properties`).

```
############################################
# Change Determination Flag
# Allows system to be flagged to only run
# change determination against devices we
# have received Config Change Event for.
# Default Behavior is to run change determination
# for All targets (the same as setting the below property = false)
#com.dorado.changemgmt.change.determination.require.config.events=true
```

Compliance and Change Reporting

The Compliance Policy Violationreport is seeded when you have ProScan/Change Management in Cruz. Inventory Compliance Attributes for reporting can also appear in report templates when you install ProScan. These report in-compliance or out-of-compliance, the last compliance date (when last compliant or not compliant), last config date (when configuration last changed), last checked date (when change was last determined).

You can also run the Change Determination Report that displays changes made to configurations.

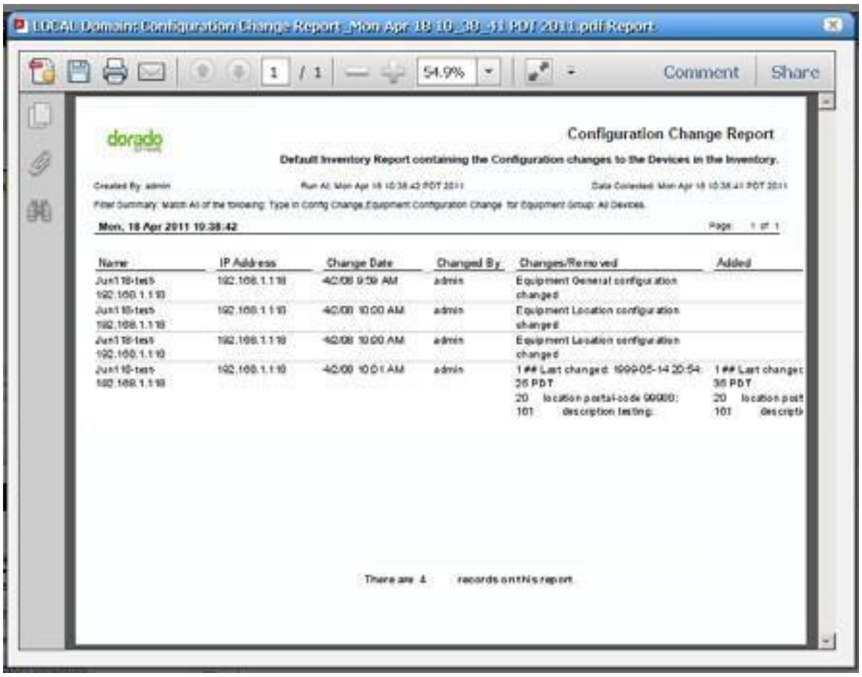
LOCAL Domain:Configuration Change Report_Mon Apr 18 10_38_41 PDT 2011.pdf Report
1 / 1
54.9%
Comment
Share
dorado

Configuration Change Report

Default Inventory Report containing the Configuration changes to the Devices in the Inventory.

Created By: admin   Run At: Mon Apr 18 10:38:42 PDT 2011   Data Collected: Mon Apr 18 10:38:41 PDT 2011

Filter Summary: Match All of the following: Type in Config Change.Equipment Configuration Change for Equipment Group: All Devices.

Mon, 18 Apr 2011 10:38:42   Page 1 of 1

Name	IP Address	Change Date	Changed By	Changes/Removed	Added
Jun118-test 192.168.1.118	192.168.1.118	4/2/08 9:59 AM	admin	Equipment General configuration changed	
Jun118-test 192.168.1.118	192.168.1.118	4/2/08 10:00 AM	admin	Equipment Location configuration changed	
Jun118-test 192.168.1.118	192.168.1.118	4/2/08 10:00 AM	admin	Equipment Location configuration changed	

Jun118-test
192.168.1.118
	192.168.1.118	4/2/08 10:01 AM	admin	1 ## Last changed: 1999-05-14 20:54:36 PDT
				20 location postal-code 99000;
				101 description testing;
					1 ## Last changec
					36 PDT
					20 location post
					101 descripti

There are 4 records on this report

See Generating a Report on page 267 for more about reporting capabilities.

The Change Determination Report displays detected changes based on a configuration change flag set when Cruz detects achange made to the device. Tosuccessfully execute this report, you must enable a scheduled Change Determination Process. The process must run before the reports has any contents. To run the process, go to the Schedules portlet, and schedule that change determination process.

Reporting Limitations

The Configuration Change Report only reports on incremental configuration changes discovered in the CD process. Simply making changes to configurations and backing them up in Cruz does *not* ensure these appear in *Configuration Change Reports*. They appear in reports only after running the CD process.

The *Configuration Change Report* includes a Filter that you can alter at runtime. By default, the report filters on *Type* only. If you want more filter criteria—like device IP, and/or date ranges—you must edit the Report filter. To edit the filter, in the Reports manager, right click the *Configuration Change Report,* and select *Open,* then edit the filter in the *Filter* screen by selecting that node on the left.

NOTE:

ArecommendedbestpracticeistoexecutetheCDprocessasanoperationrunagainstmultiple resources following ascheduled group backup of these resources. If you run backups everyday, the *ConfigurationChangeReport* thenshowsthedailychanges,untiltheyarepurgedfromthedatabase.

The application stores the specifics of what changed for future reporting.

How To:

Report on Change Determination

Follow these steps to produce regular change determination reports:

1. First, ensure the devices you want to scan are discovered, and send change notifications to the application server.

   Juniper JUNOS-based routers, for one example, provide configuration change information with an SNMP trap. The following configuration determines that configuration change traps are being sent to a host 192.168.1.24:

   ```
   trap-group test {
           categories {
               configuration;
           }
           targets {
               192.168.1.24;
           }
       }
   ```

   Check your vendor’s manuals to determine how to forward configuration change information to Cruz for your system. See Forwarding Configuration Change Commands on page 514 for others.

2. When Cruz receives a configuration change notification, in the JUNOS-based example, the device transmits an event (jnxCmCfgChange) to the Cruz mediation server. When received, this event automatically generates an event called redcellEquipmentConfigChangeNotification. Event history displays that notification.

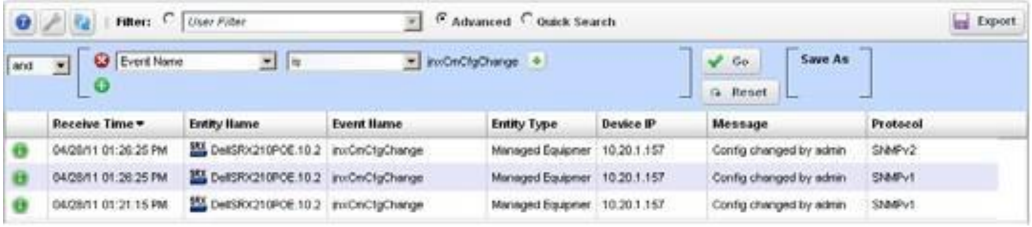

and | Event Name | is | jnxCmCfgChange

	Receive Time	Entity Name	Event Name	Entity Type	Device IP	Message	Protocol
●	04/28/11 01:26:25 PM	DellSRX210POE.10.2	jnxCmCfgChange	Managed Equipmer	10.20.1.157	Config changed by admin	SNMPv2
●	04/28/11 01:26:25 PM	DellSRX210POE.10.2	jnxCmCfgChange	Managed Equipmer	10.20.1.157	Config changed by admin	SNMPv1
●	04/28/11 01:21:15 PM	DellSRX210POE.10.2	jnxCmCfgChange	Managed Equipmer	10.20.1.157	Config changed by admin	SNMPv1

3. When CruzreceivestheredcellEquipmentConfig ChangeNotificationevent,itcaninitiate(if enabled) an event processing rule called *Configuration Change.*

   This processing rule triggers a flag in the Cruz database saying a change has occurred in the device’s configuration and that Cruz should run change determination against the device when requested.

4. When you run Cruz’s change determination process, it reviews the flag setting in the database and backs up a managed device if the flag indicates a change. This backup updates the Cruz label *Current* which is then compared to the Cruz *Change Determination* label. Cruz then writes the differences between the two labelled configurations to its database, where it is available for reporting purposes.

5. Once this occurs, the *Change Determination* label moves to point to the same configuration which is reflected by the *Current* label.

6 The report which can run to display these changes is Cruz's *Configuration Change Report*. It displays the name of the device in question, the IP address, date/time of change, who made the change, what was removed and what was added. You can schedule this report to run immediately after a Change Determination process too, so you can capture a history of changes.

Forwarding Configuration Change Commands

The following are setup commands for various other devices

Cisco

The following is the configuration to forward notifications to 192.168.1.176;

```
logging facility local0
logging 192.168.1.176
```

Thisconfiguration comes from executing the following commands on the CLI once you are logged in to privileged mode

```
conf t
logging on
logging 192.168.1.176
logging facility local0
end
wr mem
```

Netscreen

Netscreen'sforwardconfiguration changeinformation viasyslog. Tomake Change Determination work for Netscreen devices you must:

- Enable SyslogNotification event definition
- Enable Syslogon Netscreendeviceand forward Notificationsyslogfacility(which includes configuration change information) or Debugging facility (which forwards all facilities).

To do this, deploy something like this example configuration:

```
ns208-> get syslog
Syslog Configuration:
       Hostname: 192.168.1.176
        Host port:514
        Security Facility: local0
        Facility: auth/sec
        Traffic log: disabled
        Event log: enabled
        Transport: udp
        Socket number: 1030
       module=system: emer, alert, crit, error, warn, notif, info, debug
Source interface: ethernet1
Syslog is enabled.
```

10

Traffic Flow Analyzer

Thissection providesinformation about the Traffic Flow Analyzer (TFA) component, whichallows you to view and report on detailed traffic flow data within your network. TFA listens on UDP ports for packets in the NetFlow and sFlow family of protocols (this includes NetFlow implementations such as JFlow, NetStream, etc).

Using this data, Traffic Flow Analyzer can help you visualize network traffic, troubleshoot and anticipate bottlenecks.

Supported versionsinclude sFlowv5, and NetFlow v5 and v9, and IPFIX (which is essentially NetFlow v10) but not previous versions (for example, v2).

Typical packages come with a default limit to the number of registered exporters. Upgrade your license if you want to exceed the licensed limit. Tr affic Flow exporters are licensed separately from the Right to Manage (RTM) resource license count, so a license that includes 50 RTM does not necessarily let you have 50 Traffic flow exporters.

Cruz typicallylimits the number ofexportersavailablefor your system through licensing. Review your Traffic Flow Analyzer license setting in the license viewer under Product Licenses tab, the license detail section (MaxExporterCount =n, where “n” is the number of exporters licensed.) After reviewingrecommendations forperformance tuning and hardwaresizingin thisdocument, you can request a license allowing more exporters from your sales representative.

How does Traffic Flow work?

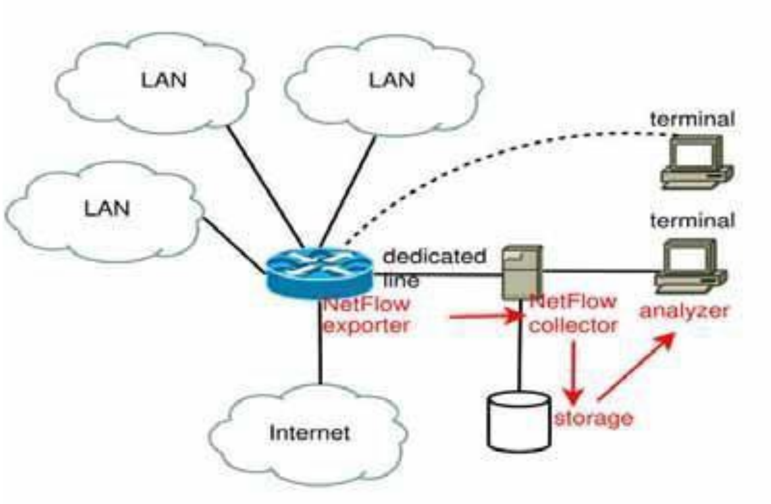

- The NetFlow/IPFIX/sFlow exporting device monitors the traffic that traversesing it
- The device becomes an Exporter of NetFlow/IPFIX/sFlow data.
- It exports the information to the NetFlow/IPFIX/sFlow Collector
- The collector stores, correlates and presents the information about
- Traffic bottlenecks in networks.
- Applications responsible for bandwidth utilization.

Aflowisaunidirectional stream ofdata between two network nodes. Thefollowingkeyparameters appear in flows:

- Source IP address
- Destination IP address
- Source port number
- Destination port number
- Layer 3 and 4 protocol type
- Input interface or port
- Output interface or port
- Source Autonomous System Number
- Destination Autonomous System Number
- Number of Bytes
- Number of Packets

Definitions

NetFlow —A traffic flow monitoring technology that is somewhat open-ended and thus it has variousimplementations bydifferent hardwarevendors. Sampling is optional, which means that the reported data might represent samples of the actual data or a full and detailed reporting of all traffic flow data.

IPFIX—Another traffic flow monitoring technology that is similar to NetFlow but is even more open-ended and allows for more flexibility in how it can be implemented by hardware vendors. Sampling is also optional for this technology.

jFlow — Juniper's implementation of NetFlow.

sFlow — Another traffic flow monitoring technology that has a clearly defined standard and thus there is verylittle leewayfor how hardware vendorscanimplement it. Sampling of traffic flow data is required in this technology, so a full reporting of the data is not possible and instead you would need to rely on estimated actual values.

Collector — Application listening on a UDP port for NetFlow/IPFIX/sFlow datagram.

Exporter — Network element that sends the NetFlow/IPFIX/sFlow datagram.

Endpoint—A network node that sends and/or receives flows. If the flowis merely passing through the node but was not sent nor received by the node, then this node would not be considered an endpoint (not for this particular flowanyways) but it might be considered an exporter (if it is registered as such).

Conversations—IP Communications between two networkendpoints. Forexample, ifendpoint A sends data to endpoint Z and then the data flow is reversed so that Z sends data to A, these flows are both part of the same conversation.

Flow — A flow is a unidirectional stream of packets between two network endpoints.

NOTE:

Counter sFlow packets do not appear as Traffic Flows, but essentially duplicate Performance metrics for interfaces. Such Flows monitor how data traverses between *two* endpoints. You can monitor interfaces with Performance monitors. See Performance Monitoring.

Setup Tasks

To set up TFA, you will need to configure one or more devices to send traffic flow data to Cruz and you will also need to register these devices as exporters. Each of these tasks begins from the Managed Resources portlet and the Traffic Flow Analyzer sub-menu that appears when you right-click on a device.

Creating TFA Configuration for a Managed Resource (Top-Level Device)

For any devices in which you would like to get visibility into what traffic is flowing through it, you will need to configure such devices to summarize and/or to sample this data and to send the flow information to Cruz. This can be done automatically for some device vendors and models by using the Create Configuration menu item. Note that not all devices and models are supported, and thus, manual configuration would be necessary for unsupported models. Manual configuration can often be accomplished by logging into the device either through CLI commands and/or through a web GUI. Please check the documentation for your device to see if any traffic flow protocols are supported, and if so how to enable this feature. For automatic configuration, follow these steps:

1 From the Managed Resources portlet, right-click on the appropriate device.

2 Select Traffic Flow Analyzer >Create Configuration.

3 popup box will appear that prompts for the parameters that are necessary to perform this action. These parameters are determined by the configuration parameters of the device model. Depending on the device model, these parameters might include sample rate, active timeout, etc.

4 After you have entered the appropriate parameters, click Execute.

5 If this operation was successful, then the device should be configured to send traffic flows to {product name} using one of the supported protocols (sFlow, NetFlow, IPFIX). Please note, however, that for some models it is also necessary for you to create configuration for TFA at both the device level (using the Managed Resources portlet) and also at the subcomponent level (using either the Ports or Interfaces portlet, which is described in the next section.

CreatingTFA ConfigurationforaPortorInterface(Subcomponent)

Depending on the device model, you might need to create the TFA configuration for both the top-level device and for one or more of its subcomponents (ports and interfaces). You will probably have to first configure the top-level device and then you can configure the subcomponent. To create TFA configuration for a subcomponent, follow these steps:

1 From either the Ports or Interfaces portlet, right-click on the appropriate device.

2 Select Traffic Flow Analyzer >Create Configuration.

3 A popup box will appear that prompts for the parameters that are necessary to perform this action. These parameters are determined by the configuration parameters of the device model. Depending on the device model, these parameters might include polling interval, sample rate, etc.

4 After you have entered the appropriate parameters, click Execute.

5 If this operation was successful then the port or interface should be configured to send traffic flows to Cruz using one of the supported protocols (sFlow, NetFlow, IPFIX).

Removing TFA Configuration from a Managed Resource, Port, or Interface

Toremove the TFA configuration for a device, whether this be the top-level device or a subcomponent (port or interface) follow these steps:

1 From the Managed Resources portlet, right-click on the appropriate device.

2 Select Traffic Flow Analyzer >Remove Configuration.

3 A popup box will appear indicating that no parameters are necessary to perform this action. Click Execute.

4 If this operation was successful, then the device should be configured to no longer send traffic flow data to Cruz.

Showing TFA Configuration for Managed Resource, Port, or Interface

1 From either the Managed Resources, Ports, or Interfaces portlet, right-click on the appropriate device.

2 Select Traffic Flow Analyzer >Show Configuration.

3 A popup box will appear indicating that no parameters are necessary to perform this action. Click Execute.

4 The TFA configuration for the device will be shown.

Registering an Exporter

After you have configured a device to send traffic flowinformation to Cruz, you must still register it as a traffic flow exporter in order for these flows to be processed. To register a device, follow these steps:

1. From either the Managed Resources portlet, right-click on the appropriate device.
2. Select Traffic Flow Analyzer > Register. This menu item will only be available if the device is not already registered, which means that the TFA Registered column should show an X.
3. A message will confirm that the registration succeeded and also the TFARegistered column should display a check mark indicating that the device is registered for TFA.
4. The system should then be ready to accept and process flow data from the device. If a device is not registered, the Register option appears in the menu.

Unregistering an Exporter

To unregister a device, follow these steps:

1. From the Managed Resources portlet, right-click on the appropriate device.
2. Select Traffic Flow Analyzer > UnRegister. This menu item will only be available if the device is already registered, which means that the TFARegistered column should show a check mark.
3. A message will confirm that the unregistration succeeded and the TFARegistered column should display an X indicating that the device is not registered for TFA.
4. The system will then not process any flows that are received from this device. It is recommended to remove the configuration from this device as well.

Showing Traffic

.

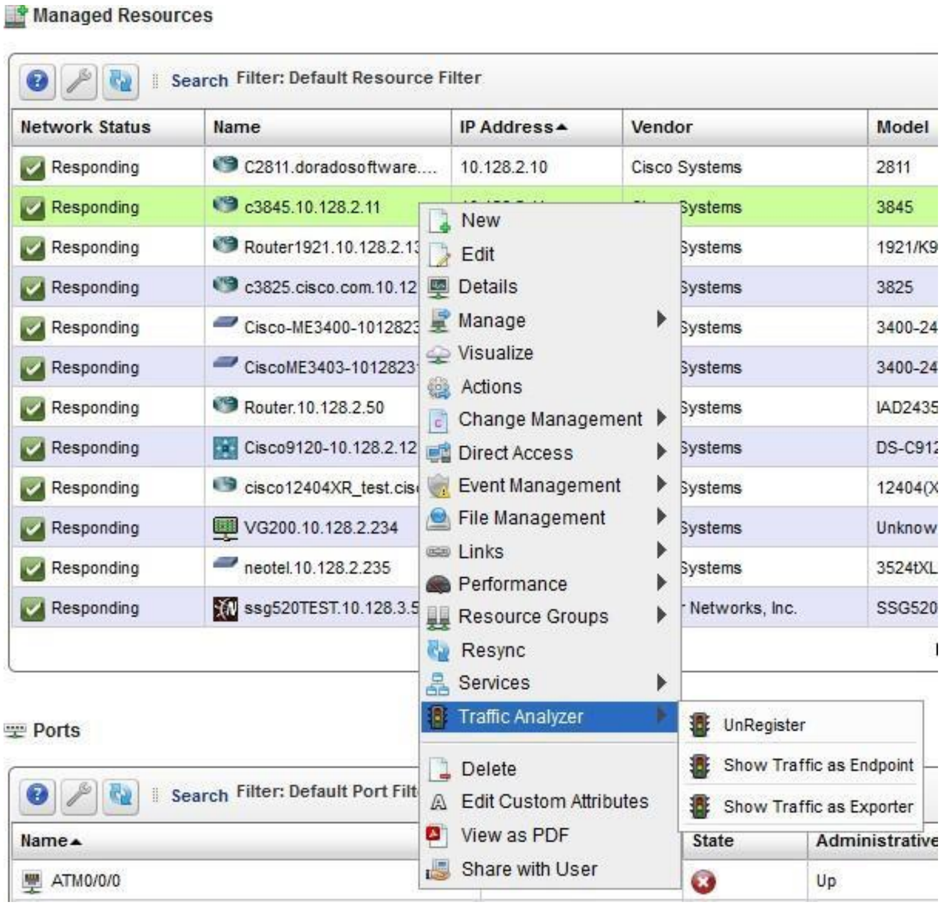

Network Status	Name	IP Address▲	Vendor	Model
Responding	C2811.doradosoftware....	10.128.2.10	Cisco Systems	2811
Responding	c3845.10.128.2.11		Systems	3845
Responding	Router1921.10.128.2.13		Systems	1921/K9
Responding	c3825.cisco.com.10.12		Systems	3825
Responding	Cisco-ME3400-1012823		Systems	3400-24
Responding	CiscoME3403-1012823		Systems	3400-24
Responding	Router.10.128.2.50		Systems	IAD2435
Responding	Cisco9120-10.128.2.12		Systems	DS-C912
Responding	cisco12404XR_test.cis		Systems	12404(X
Responding	VG200.10.128.2.234		Systems	Unknow
Responding	neotel.10.128.2.235		Systems	3524tXL
Responding	ssg520TEST.10.128.3.5		Networks, Inc.	SSG520

Name▲	State	Administrative
ATM0/0/0	✗	Up

Showing Traffic for a Managed Resource (Top-Level Device)

You can view the traffic flow data for any device in the Managed Resources portlet if you select a device and right-click and then expand the Traffic Flow Analyzer sub-menu. There are two *Show Traffic* menu options available from this sub-menu, both ofwhich will navigate to full-screen views of the expanded Traffic Flow Portlet and will show the traffic flow data for the selected device. *Show Traffic as Endpoint* is available for all managed devices and this option navigates to a page showing the traffic where this device was either the sender or the receiver of the flow (as endpoints are essentially senders and receivers combined). Show Traffic as Exporter is available only for registered exporters and this option shows all traffic going through the device. Once this portlet is shown, a drop-down list of available information types will also be available that will allow you to change the type of data that is displayed.

Showing Traffic for a Port or Interface (Subcomponent)

You can view traffic flow data for a port or interface by using the Show Traffic menu item that is available in either portlet. Within either portlet, select an entity and right-click and then select Traffic Flow Analyzer>Show Traffic. This will navigate to the expanded Traffic Flow portlet and the context will be set to show the data associated with the selected entity (if any data is available).

Showing Traffic from the Traffic Flow Page

Youcan view a variety of different types of TFA data by selecting the Traffic Flow menu item near the top of the screen. This page displays a portlet for each of the following: Exporters by Managed Equipment, Exporters by Subcomponent, Top 5 Applications, Top 5 Autonomous Systems, Top 5 Conversations, Top 5 Endpoints, Top 5 Protocols, Top 5 Receivers, and Top 5 Senders.

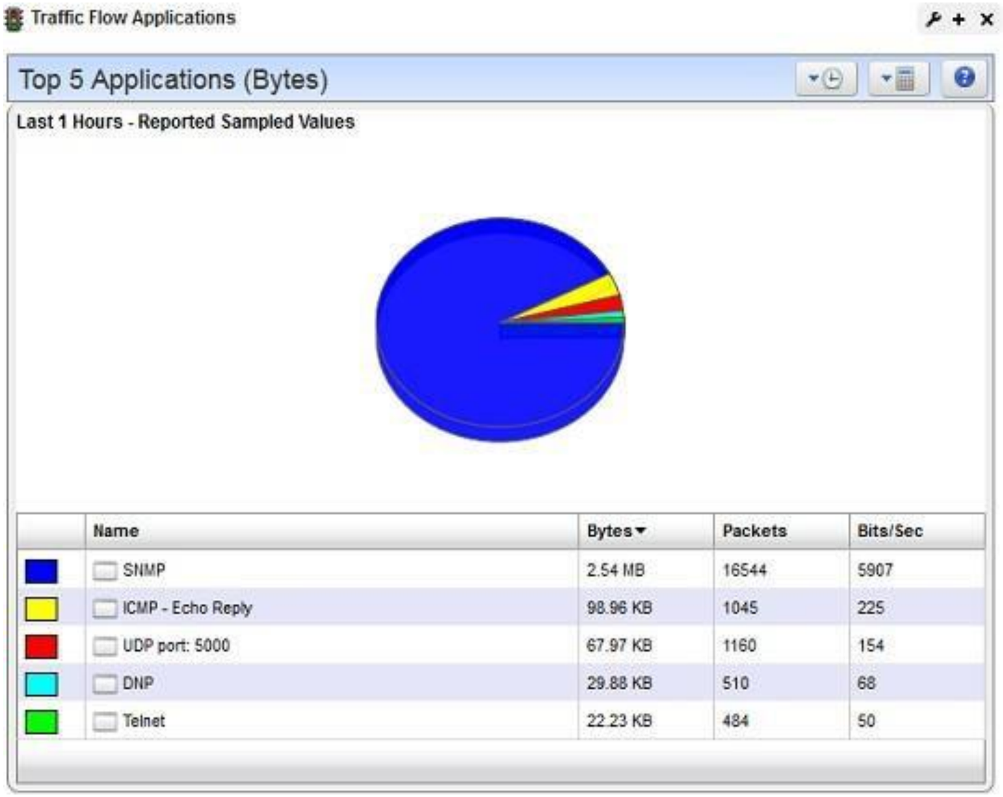
Traffic Flow Applications
Top 5 Applications (Bytes)
Last 1 Hours - Reported Sampled Values

	Name	Bytes	Packets	Bits/Sec
🟦	SNMP	2.54 MB	16544	5907
🟨	ICMP - Echo Reply	98.96 KB	1045	225
🟥	UDP port: 5000	67.97 KB	1160	154
🟦	DNP	29.88 KB	510	68
🟩	Telnet	22.23 KB	484	50

Traffic Flow Portlet

Traffic Flow Analyzer uses several types of portlets, one for each of the types of objects on which it reports. These are Applications, Autonomous Systems, Conversations, Endpoints, Exporters by Equipment Manager, Exporters by Subcomponent, Protocols, Receivers and Senders. Note that Endpoints, Conversations, Receivers, and Senders all have similar data. The way it works is that every flow has a sender and a receiver, and these are both considered endpoints. Also, if endpoint A sends a packet that is received by endpoint Z and then Z sends another packet back to A, these packets are both part of the same conversation between A and Z.

When you add one of the traffic flow analyzer portlets to a page, its summary, or minimized form appears. This displays a simple view containing a pie chart and a table showing the summarized collected data over the configured time period. Only the time frame (shown with a clock icon), the mode ofcalculation (shown with acalculator icon), and the flowdirection (ingress, egress, or both) can be changed in this view. These controls appear as dropdown buttons in the upper right corner of the portlet. Most minimized traffic flow portlets are limited to the top entities as measured by total bytes. The minimized Exporters by Equipment Manager portlet is not limited to 5 entries. It supports pagination if there are lots of exporters, so that all the necessary data can be shown.

Tosee more information about any item within the minimized traffic flow portlet, you can click on the name and this will navigate to the expanded traffic flow portlet, as shown in the context of the selected item. Another way to navigate to the expanded portlet is to click on the + sign in the upper right corner.

The Expanded Traffic Flow Portlet displaysan interactive graph. Youcan also Drill Down to details about components within this portlet by clicking on one of the links in the table below the graph.

NOTE:

The selected period determines whether data is present, especially if you have just started monitoring Traffic Flow. Choose theshortest period tosee data immediately (it still takes afew minutes to appear), and select longer periods only after monitoring has run for longer periods.

Expanded Traffic Flow Portlet

When you expand the portlet, a more complex interactive view appears. Initially, it displays a line graph for the selected time frame.

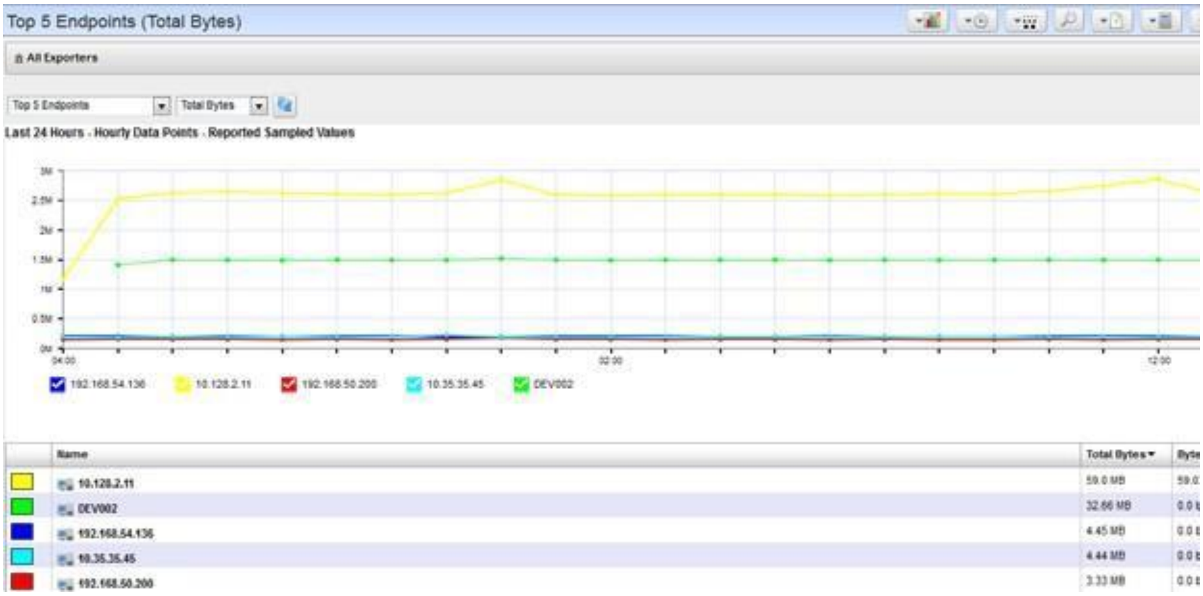

Top 5 Endpoints (Total Bytes)

All Exporters

Top 5 Endpoints | Total Bytes

Last 24 Hours . Hourly Data Points . Reported Sampled Values

	Name	Total Bytes ▾	Byte
🟨	10.128.2.11	59.0 MB	59.0
🟩	DEV002	32.66 MB	0.0
🟦	192.168.54.136	4.45 MB	0.0
🟦	10.35.35.45	4.44 MB	0.0
🟥	192.168.50.200	3.33 MB	0.0

The graph on the upper portion of the portlet shows the data points for every time-slice and the chart below this graph shows the aggregate for all data points covering the selected time frame.

The following controls appear in the title bar:

Select Chart Type—Lets you change the chart type. Available chart types include *Pie, Line, Bar, Stacked Bar* and *Column*. Note that the pie chart only displays aggregate data while the other chart types display time-slice data within the selected time frame.

Select TimeFrame—Several built-in relative time frames are available including *Last 15 Minutes, Last Hour, Last 6 Hours, Last 24 Hours, Last 5 Days, Last 30 Days and Last 12 Months. you can also select Custom to display the Custom Time Frame popup screen.*

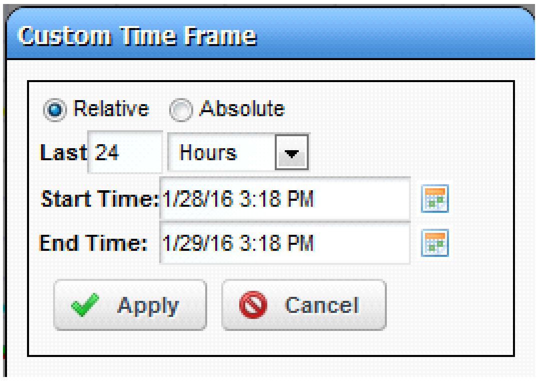

Custom Time Frame

● Relative ○ Absolute

Last 24 Hours

Start Time: 1/28/16 3:18 PM

End Time: 1/29/16 3:18 PM

When this screen is shown, first select either Relative (wherein the time frame starts a certain period of time ago and includes everything up to the present) or Absolute (wherein the time frame includes everything between a certain arbitrary start time and end time that can both be in the past). If Relative is selected, you must enter a relative time value and time unit. For example, Last 8 Hours, Last 120 Minutes, Last 45 Days, etc. If Absolute is selected, you must enter an absolute start and end time. Click Apply for these settings to take effect.

Data for Last 15 Minutes shows minute-by-minute data and this type ofdata typicallyappears after about 2-5 minutes of collection (after the exporter is first registered and flow packets are received from this device). Data for Last 24 Hours shows hourly data by default, and this type of aggregate data is computed after each hour has passed. For example, aggregate data for 9:00 AM will only be computed after the 9:00 hour has passed (9:59 AM has rolled over to 10:00 AM). Likewise, daily data is computed once the day has passed into the next (11:59 PM hasrolled over to 12:00 AM). Changing the time frame displays time-slice data points, each of which represent a certain time range, such as a single minute of data or a full hour of data. The granularity of the time-slice data points is determined initially by the length of the selected time interval, but you can then use the "Select Granularity of Data Points" option to display different time-slice data points for the same time frame.

The granularity of time-slice data points is determined initially by the length of time frame (the duration from the start time to the end time), according to the following thresholds:

Less than or equal to 60 minutes: 1 Minute data points

Less than or equal to 12 Hours: 10 Minute data points

Less than or equal to 3 Days: Hourly data points

Less than or equal to 30 Days: Daily data points

More than 30 Days: Weekly data points

Select Granularity of Data Points —Lets you change the granularity of the time-slice data points that are shown in the graph. The options for granularity are 1 Minute, 10 Minute, Hourly, Daily, and Weekly. The way this works is that the less granular data (representing longer time intervals) are aggregates of the more granular data. For example, a 10-minute data point representing 6:20 would be an aggregate of all 1 minute data from 6:20 to 6:29. Likewise, an hourly data point from 6:00 would be an aggregate of all 1 minute data from 6:00 to 6:59. When you select a time frame, the length of the interval determines the granularity of the data points that are shown at first, but you can change this. Forexample, the time frame Last 24 Hours will by default show hourly data points, but you can then see more granular data by selecting 10 Minute from this menu. It would even be possible to show 1 Minute data for this same time frame, but this might cause the web browser to perform poorly because there would be hundreds of data points.

Search — Displays a search dialogue to find specific traffic data.

Select Report Type—Lets you change the number of results that are shown and also whether to show the most or least significant results. Youcan choose between Top 5, 10 or 25 and Bottom 5, 10 or 25.

Select Mode of Calculation — Lets you change the mode of calculation for data values. NetFlow and sFlow data are often sampled, which means that the values reported by the exporters should be understood as representing only a fraction of the total traffic. You can choose to show either the reported sampled values or the estimated total values. The estimates are calculated by multiplying the sampled values by the sampling rate, asreported bythe exporter in one of the following ways:

- sFlow sampling, where the sampling rate is in the packet (usually 512 or 1024 or some power of 2)
- NetFlow V5, where the sampling rate is in the header of the packet
- NetFlow V9 and IPFIX with a single sampling rate for the entire device. In these cases, an options datagram contains the sampling rate for the device.
- NetFlow V9 and IPFIX with a sampling rate specific to an interface. In these cases, an options datagram contains a sampling rate and a sampler ID. Flows then contain a sampler ID to associate to this. When flows say the sampler ID is 0 then this feature is disabled for this interface.

Estimated total bytes and packets are calculated by multiplying the reported sampled values by the sampling rate. This works differently for sFlow and NetFlow. NetFlow packets contain fields for the reported sampled bytes and packets. For example, if the packet reports that the sampled bytes are 45 and the sampling rate is 100 then the estimated total bytes will be 4500. Likewise, if the sampled packets are 3 then within this same example the estimated total packets will be 300. sFlow packets do not contain fields for the sampled bytes or packets.

Instead, sFlow works by sampling packets that are going through the exporter and attaching these sampled packets to sFlow packets for export. So, the reported sampled bytes is computed byadding the total bytes of the sampled packet that is contained within the sFlow packet and the reported sampled packets is naturally 1 for each sFlow packet that is exported. The estimated total bytes and packets is computed by multiplying each of these numbers by the sampling rate that is reported in the sFlow packet. Forexample, if the sampled bytes is 10 and the sampling rate is 512 then the estimated total bytes will be 5100 and likewise the estimated total packets will be 512 since this is the sampling rate.

NOTE:

For devices that are not sampling traffic data (they are reporting on 100% of the traffic going through the device), the estimated values will be the same as the raw values.

Select flow direction—Lets you change the flow direction of the data that is queried. The options are Ingress, Egress, and Both Ingress and Egress, which is the default. This notion is related to the way in which the traffic is handled by the exporter and how the exporter was configured at the subcomponent (port or interface) level. When a flow goes through a device, it must go into some subcomponent (ingress) and out some subcomponent (egress). In some cases, only ingress flows are reported on by the exporter and in some cases only egress flows are reported. Sometimes both ingress and egress are reported, but the byte and bits/second values might be affected by whether the flowis ingress or egress, as reported by the exporter. In some cases, these values can have different data compression rates and thus it would be misleading to count them together. Also, sometimes you need to know how much data is going into or out of a port or interface. This option allows you to sort this out by focusing only on ingress or egress if you need to do so.

Traffic Flow Snapshots — Load or save a snapshot (preserved views) of traffic flow.

Export Data — Lets you export the data in the current view to a file. You can either save the current view to a PDF file or export the data to a CSV file. If you select Export to PDF, the resulting file will show a screen shot of the graph and the chart as shown on the screen. If you select Export to CSV, the resulting file will contain the data points for the time-slice data in the graph shown on the screen. The exception to this is if you are viewing a pie chart - in which case there is no time-slice data and the file will instead contain the same aggregate data as shown in the chart that appears on the screen below the graph. You can retrieve the generated file in the *My Alerts* area at the lower left corner of the portal.

NOTE:

When you export a line graph to CSV, the resulting file will have the times (minute-by-minute, hourly, etc.) in the header row, but the spreadsheet program you are using (i.e. Microsoft Excel) will apply default formatting to these values. You can apply your own custom formatting as needed.

NOTE:

The Export to PDF menu item is not available when canvas-based line charts are enabled. If you need to use this feature turn off Canvas Line Charts in the Application Settings.

Settings—Allows you configure the way Cruz stores and queries traffic flow data. Youcan configure how to retain data, based on collection/rollup intervals.

Below the title bar a navigation bar displays the context path. See Drill Down on page 526 below, for more about this.

Below that navigation bar a row containing the following controls appear:

Entity Type—Selects the type ofentity to report on (Exporters by Equipment Manager, Exporters by Subcomponent, Applications, Conversations, End points, Senders, Receivers, Autonomous Systems, and Protocol).

Attribute — Selects which attribute to graph (Bytes, Packets, Bits/Sec).

Refresh — Refreshes the screen (runs the report) applying any new settings.

NOTE:

You can check/uncheck by clicking on the colored squares in the legend below these graphs. This reveals/conceals lines connected to the labelled item.

Drill Down

you can “drill down” into a report by clicking on one of the links in the table. This displays a detail view of the selected entity and the name of the entity appears in the navigation bar.

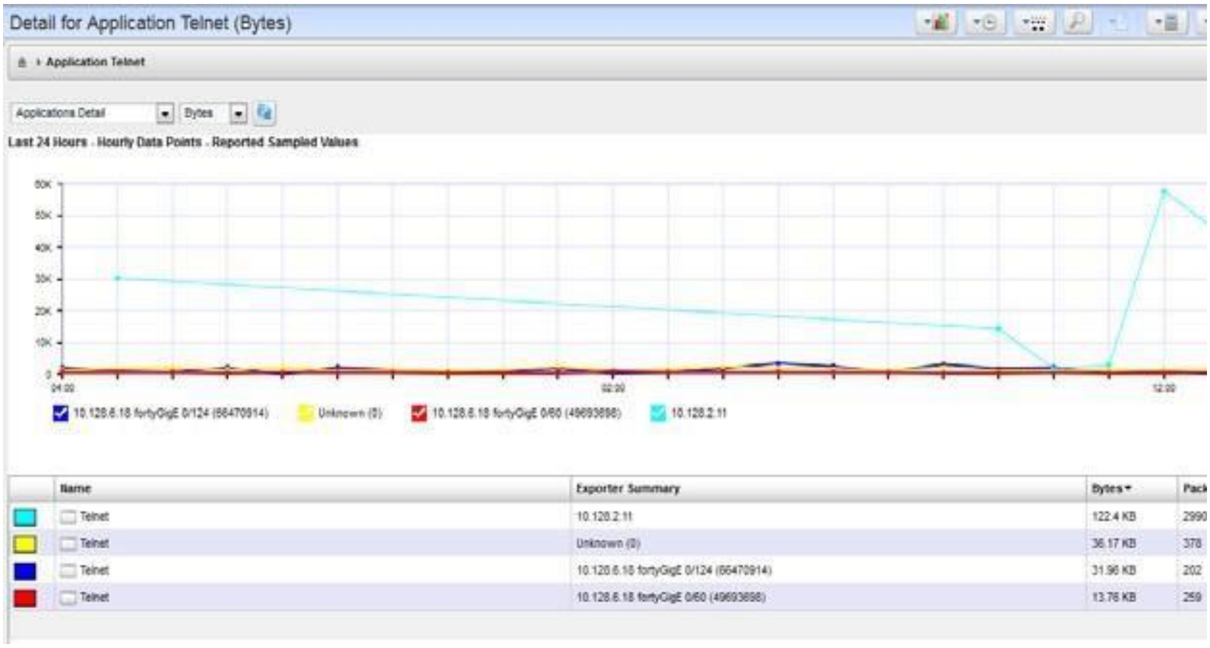

Detail for Application Telnet (Bytes)

⌂ › Application Telnet

Applications Detail ▾ Bytes ▾

Last 24 Hours . Hourly Data Points . Reported Sampled Values

60K
50K
40K
30K
20K
10K
0
04:00
02:00
12:00
☑ 10.128.6.18 fortyGigE 0/124 (66470914)
☑ Unknown (0)
☑ 10.128.6.18 fortyGigE 0/60 (49693898)
☑ 10.128.2.11

	Name	Exporter Summary	Bytes ▾	Pack
🟦	□ Telnet	10.128.2.11	122.4 KB	2990
🟨	□ Telnet	Unknown (0)	36.17 KB	378
🟦	□ Telnet	10.128.6.18 fortyGigE 0/124 (66470914)	31.96 KB	202
🟥	□ Telnet	10.128.6.18 fortyGigE 0/60 (49693898)	13.76 KB	259

When a detail view appears, the entity type appears as in the title bar. You can change to a “Top/ Bottom n” report of a different type, then click re fresh to display a report of the top entities that apply to the current detailed entity. Each time you click into the detail view for an entity, this adds the entity to the context, so that the report is generated by applying all of the filters that have been added to the context.

When you add an entity to the context, there will be fewer entity types available in the drop down list. Forexample, if you start from Exporters by Equipment Manager and then click on the name of an equipment manager in the list, this will then add this entity to the context. If you then change the entity type to Exporters by Subcomponent, then only the subcomponents within this equipment manager that are also flow exporters will be shown. Keep in mind that you might have to change the flow direction option to "Both Ingress and Egress" for all possible subcomponents to show in this view. You could then change the entity type to another available option such as Applications, Protocols, etc. for a more detailed look at what flows are going through the selected exporter. Another way of drilling down is to start with something other than an exporter, such as Senders, Receivers, etc. and clicking into the detail view for a specific entity that is shown in that view. There are many permutations of drilling down into different entity types that will work to show exactly what you are interested in. This process can continue until the Conversation Detail view is reached. This is the end of the line for drillingdown.

You can return to the Home context (in which no entity is selected, and all data is available and no "drill down" options have been added to the list) at any point by clicking the house icon to in the

upper left, to the left side of the "drill down" list.

Search

Search by clicking on the Search (magnifying glass) icon in the title bar. Type any string in the next screen to search through the traffic data. A list of all entities found matching the string appears below it.

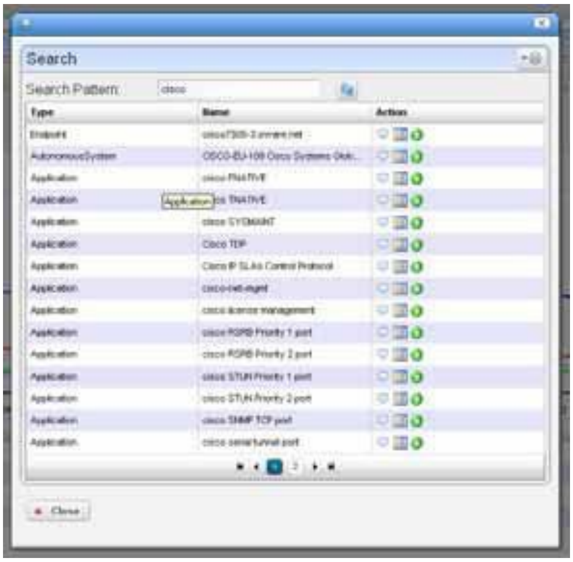

Search

Search Pattern: cisco

Type	Name	Action
Endpoint	cisco7305-3.vmware.net	
AutonomousSystem	CISCO-EU-109 Cisco Systems Glob...	
Application	cisco FNATIVE	
Application	cisco TNATIVE	
Application	cisco SYSMAINT	
Application	Cisco TDP	
Application	Cisco IP SLAs Control Protocol	
Application	cisco-net-mgmt	
Application	cisco license management	
Application	cisco RSRB Priority 1 port	
Application	cisco RSRB Priority 2 port	
Application	cisco STUN Priority 1 port	
Application	cisco STUN Priority 2 port	
Application	cisco SNMP TCP port	
Application	cisco serial tunnel port	

Entity found in the search support the following actions:

View Top Conversations — Displays the top n conversations for the selected entity.

Show Detail View — Displays a top-level detail view of the selected entity.

Add to Current View — Adds the entity to the current view and drills down to it.

Edit — Allows you to edit information about the entity. This is currently only supported for applications. If you click this icon, then another screen will be shown that allows you to edit the name and/or the keywords of the application.

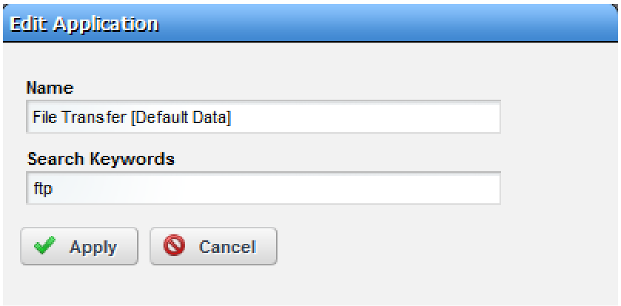

Edit Application

Name: File Transfer [Default Data]

Search Keywords: ftp

This is useful if you have a custom application that has traffic flow data associated with it. By default, the names for such applications has the format "(protocol name) port: (port number)". For example, "UDP port: 5001". If it is the case that you see an application entry that has this default name or has the wrong name and you would like to change it, you can do so through this screen. Also, you can enter keywords that will help with future searches. Click Apply to save the changes and return to the search screen.

NOTE:

The *Settings* button (the gearinthe upper right corner) lets youconfine the search by types (*All, Applications, Protocols, Autonomous Systems, Endpoints)).* Note that you cannot search by exporters through this screen. If you select one of the two items for exporters in from the entity type drop down list, then this willshowapaginated list ofallexporters (by equipmentmanagerorbysubcomponent, depending on whichitem wasselected). Ifyouwanttosearch foraspecificexporter byequipment manager, then you can gotothe Managed Resources portlet and useitsbuilt-in searching capability and then use the menu item to Show Traffic as Exporter. Likewise if you wanttosearch for aspecific exporter by subcomponent, then you can go to either the Ports or Interfaces portlet (both are different types of subcomponents) andalsouseitssearchcapability andmenuitemto Show Traffic for the subcomponent entity.

Traffic Flow Snapshot

This portlet lets you display Traffic Flowyou configure and save asa snapshot in a portlet visible on any Cruz page. It is, in effect, a portlet that permanently displays the Expanded Traffic Flow Portlet, beginning with the selected snapshot.

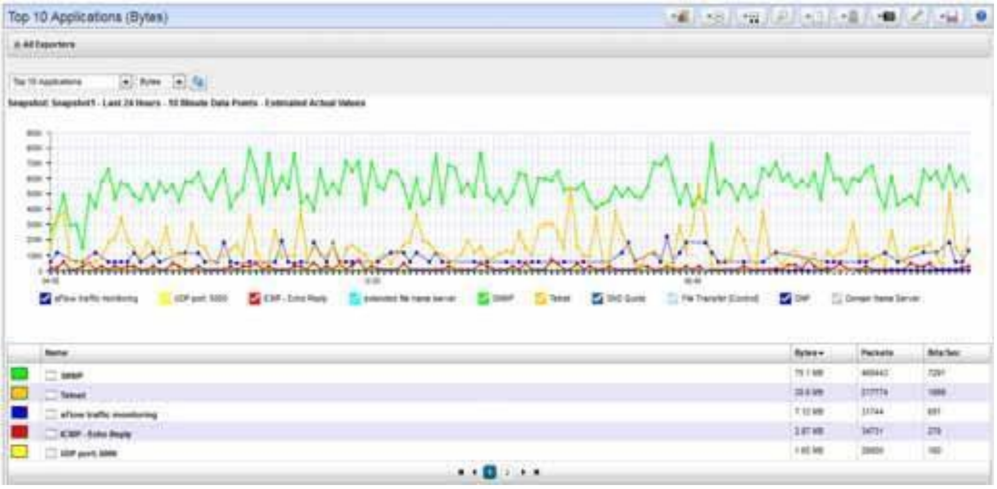

	Name	Bytes ▾	Packets	Bits/Sec
🟩	□ SNMP	75.1 MB	[illegible]	7291
🟧	□ Telnet	[illegible]	[illegible]	[illegible]
🟦	□ sFlow traffic monitoring	7.12 MB	11744	691
🟥	□ ICMP - Echo Reply	2.87 MB	[illegible]	279
🟨	□ UDP port 3000	1.65 MB	[illegible]	160

After adding this portlet to a page, use the selector to choose which snapshot you want to appear. Refresh the portlet with the double arrows to the right of the units displayed. You can also change what appears, the units, the time interval, and so on, just as described in Expanded Traffic Flow Portlet on page 522.

Settings

You can edit the configuration options by clicking on the Settings (wrench) icon.

Several of the fields shown on this screen are for the number of days to keep each granularity of traffic flow data. Initially, all traffic flows are associated with a specific minute and thus become 1-minute data in the database. These flows then aggregate (rollup) to larger degrees of granularity. 1-minutedata aggregates to 10-minute intervals, whichaggregate to hourly, which aggregate to daily, which aggregate to weekly data. It is important to consider how long you need to keep each granularity of data and to set the values accordingly. If you are saving data for longer than you need it, then this can fill up your database.

On this screen there are also fields that help you optimize queries. You can set the maximum number of rows per rollup table, which can be helpful for improving query performance. Also, you can determine whether or not the queries should use the cache tables. Using cache tables is faster, but there is a possibility that not all of the detailed flow data is in the cache tables, so it is sometimes helpful to un-check this box, which will use the raw data tables instead.

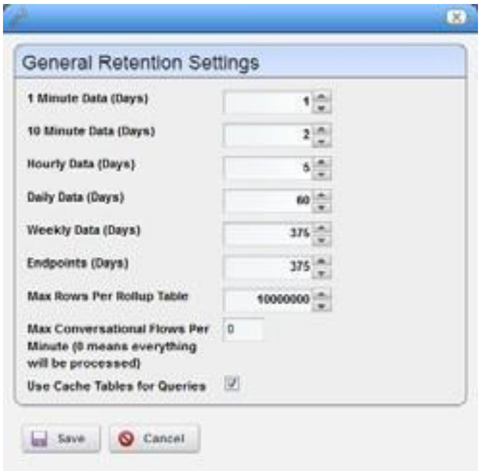

General Retention Settings

1 Minute Data (Days)	1
10 Minute Data (Days)	2
Hourly Data (Days)	5
Daily Data (Days)	60
Weekly Data (Days)	376
Endpoints (Days)	375
Max Rows Per Rollup Table	10000000
Max Conversational Flows Per Minute (0 means everything will be processed)	0
Use Cache Tables for Queries	☑

Sometimes the volume of traffic flow data collected from the exporters can overwhelm the database. Cruz collects traffic flowpackets (sFlow, NetFlow, IPFIX) fromthe registered exporters and aggregates the data into 1-minute flows differentiated by exporter, protocol, application, and conversation (including sender and receiver for both IP and autonomous system). Every minute, these flow records are inserted into the database.The number of distinct flows, differentiated by each of these fields, can sometimes add up to more row inserts than the database can handle in a single minute. In such cases, the only way for the system to function is for it to only insert the conversation-based flows that are the most significant (highest byte totals) and allowfor the less significant flows (low byte totals) to be aggregated into a single flow that represents all other conversations.

This can be done through the “Max Conversational Flows Per Minute” feature. This feature is disabled by default, but if enabled it will reduce the number of inserts to the database per minute while preserving the most significant flows and while preserving the overall byte and packet totals. What this means is that if you enter a number for “Max Conversational Flows PerMinute” then you are configuring Cruz to keep that many flows per minute. So, this many flows will be inserted into the database as differentiate d by exporter (by equipment manager and by subcomponent), protocol, application and conversation (sender and receiver) but any flow above this number will lose their conversation data and willonlybe differentiated bythe other threefields (exporter, protocol, and application). It does this by first ranking all flows by estimated total bytes and then taking the highest N and inserting each of these into the database as they were reported by the device. The rest of the flows are then aggregated into the “Other” category and inserted into the database.

For example, flows with estimated total values of 3, 15, 44, 89, 248, 510, 746, 1038, 4313, and 9755 and "Max Conversational Flows Per Minute" of 5 would find the top 5 to be 510, 746, 1038, 4313, and 9755. These flows would have their sender and receiver data (IP and AS) saved in the DB but the values 3, 15, 44, 89, and 248 would be added together into the "Other" category, but their protocol, application, and exporter data would be preserved. So, these smaller flows would only actually be aggregated if they were all for the same exporter, protocol and application.

To enable this feature, enter a value for Max Conversational Flows Per Minute. Enter the highest number of conversational flows that you believe your system is capable of handling per minute. Note that a value of 0 disables this feature so that everything will be processed.

Domain Name Resolution

Initially the Name column ofsenders, receivers, endpoints and conversations shows the IP address. There is a task that is configured to execute periodically to resolve these IP addresses to domain names. This works if a DNS server is available within the network. When an IP address is resolved to a domain name, this name is shown in the Name column when viewing senders, receivers, endpoints and conversations. By default, this task executes every 24 hours. To configure how this task is executed, navigate to the Schedules portlet and search for an item where the description contains the text "Resolve hostnames using DNS". From here you can edit the frequency of the execution of this task and you can also disable this feature if you wish.

There is also an option to perform this DNS look up for public IP addresses but to exclude private IP addresses. To disable private IP DNS resolution while keeping this feature enabled for public IP addresses, add the following line to installed.properties and restart the application server:

NetFlowRetentionMBean.EnabledPrivateIPDNSLookup = false

You can also manually execute this task by navigating to the Actions portlet and searching for an item where the name contains "Resolve DNS Hostnames". From here you can execute this action on demand rather than waiting for the scheduled task to execute. When you execute this task manually, you have the option to re-do the DNS resolution for all endpoints, including those that were previouslyresolved. Caution: Executing this action with this option selected can be very time consuming and resource intensive.

Resolving Autonomous System (AS) Numbers

Cruz provides local resolution ofautonomous system numbers (ASN) based on static mapping of AS number registrations and it also performs re mote resolution of these names for AS numbers that are not in the local file. It also supports user overrides of the default AS number to name mappings. To do this, configure properties you can find in the \owareapps\trafficanalyzer\lib\ta.properties file. Remember, best practice is to override properties as described in Overriding Properties on page 112.

Traffic Flow Analyzer - Example

The following describes typical situations where flow is useful. When ports are over-utilized because of intermittent performance problems diagnosis of the problem sometimes difficult. Turn on flow traffic data collection to evaluate who, what applications, and so on, are responsible for the traffic on the affected ports. This avoids getting overwhelmed with collection of traffic going in all directions. Follow these steps to do this:

1 From the Resources monitor, select a desired device that has support for NetFlow/IPFIX/sFlow

2 Enable NetFlow/sFlow on most impacted devices that support NetFlow/IPFIX/sFlow. Also, register a number ofexporters to enable an efficient and scalable data collection environment.

NOTE:

You can disable NetFlow/sFlow and unregister exporters.

3 After NetFlow/IPFIX/sFlow hasbeenrunningfora while, verifythat bandwidth utilization is within expectation. This will help insure optimum performance of critical business applications.

4 Select the Top 5 Applications portlet (or add it to the page).

5 From the list of the Top 5 Applications, you'll typically see most bandwidth is being consumed by the key applications in our organization.

Alternative 1

6 Toensure bandwidth is not being hijacked by unauthorized or unwanted video or music streaming applications, select the Top 5 Conversations.

7 Often the top conversation is video streaming software.

8 To answer "Where and who is running this rogue application?" drill down into the conversation to see End points involved in the conversation. This identifies the user running the streaming application. You could now go and stop (or block) this rogue application.

Alternative 2

An alarm indicates port X is surpassing its threshold. If the port has become a bottleneck in the overall network bandwidth, we want to identify what applications are at cause, and who is responsible for running them.

1 Look in the Top 5 Traffic Flow Endpoints portlet.

2 From the list of the Top 5 Endpoints, you will typically see that port X is high on the list.

3 Expand the portlet and drill down into the port X endpoint to see what the top conversations are going through port X.

4 Drill down intoconversations to identify any unauthorized applications.

5 Drill down further to identify users of any unauthorized applications

6 Now, go stop them!

How To:

Create a Traffic Flow Report

You can create reports based on traffic flow data.

1 Create a new Report Template by right clicking the Report Templates portlet, selecting New > Table Template.

2 Name the report (here: Test Traffic Flow Applications Report).

3 Select a source in the Source tab. Here: Tr affic Flow Analyzer > Traffic Flow Exporter.

4 Notice that the Select your inventorycolumns paneldisplays the attributes availablebased on your traffic flow entity type selection.

5 Select Available columns and click the right arrow to move them to Selected. Follow these guidelines:

> You should always select the entity name column (Exporter Name for the entity type Traffic Flow Exporter, Application Name for the entity type Traffic Flow Application, etc.) There is also certain entity type specific columns that provide context for the entity that is being reported on. For example, Traffic Flow Autonomous System has an available column for ASN and Traffic Flow Conversation has columns for An Endpoint and Z Endpoint.
>
> Each template should be for either time-slice data (where the data values represent a specific slice of time within the total time range of the report) or aggregate data (where the data values represent sums or averages across the entire time range). Some columns are appropriate for one of these but not the other.
>
> Templates for time-slice data should include the Time column and one or more time-slice data columns, including Bytes, Packets, and Bits/Sec. For certain entity types the available data columns might include the three of these, with directional orientation "In" and/or "Out" and perhaps also "Total", which is the "In" and "Out" values added together.
>
> Templates for aggregate data should not include the time column because the time is understood to be the total time range of the report. Also, aggregate templates should include one or more aggregate data column, which includes Bytes (sum), Packets (sum), and Bits/Sec (avg). For certain entity types there might be directional ("In", "Out", "Total") versions of these aggregate attributes.
>
> Note that time-slice templates should not include aggregate data columns, and aggregate templates should not include time-slice data columns.

6 Arrange thecolumns and fonts as you like in the Layout tab.

7 Save the template.

8 Right-click and select New in the Reports portlet.

9 Enter a Name and Title for the report.

10 Notice that since this is the first report created since you made the Test Traffic Flow Applications Report template, that it is the Report Template already selected.

11 Select the desired filter conditions. Note that all traffic flow entity types have the same available filter attributes. If a value is not selected for an attribute, then a default value will be used:

> Calculation Type defaults to Reported Sampled Values.
>
> Data Point Type defaults to Hourly.
>
> Report Limit defaults to 10. Note that only the Equals operator works for this filter condition.
>
> Report Limit Type defaults to Top.
>
> Direction defaults to both Ingress and Egress.
>
> Time defaults to Within Last 24 Hours. Note that the only comparison operators that work for this filter conditions are Within Last, Is, Between, and After.
>
> Exporter Equipment Manager is not specified by default (meaning that datafor all exporters should be included), but this attribute can be used to filter by a specific exporter by equipment manager (which includes all subcomponents).
>
> Exporter Subcomponent is also not specified by default (again this means that data for all exporters should be included), but this attribute can be used to filter by a specific exporter by subcomponent (port or interface).

12 Test Traffic Flow Applications Report should appear in the Reportsportlet.

13 Right-clickand select Execute (noticing that you can also schedule such reports, even repeatedly).

14 Click the magnifying glass to the right of the Report Completed message in My Alerts to see the report.

15 Hover your cursor over the lower right corner of the report to see a set of icons that let you expand, zoom out and in, save, or print the report.

NOTE:

Traffic Flow Analyzer doesn't support comparison reporting.

Traffic Flow Analysis Life Cycle

The following diagrams Cruz's Traffic Flow Analysis life cycle

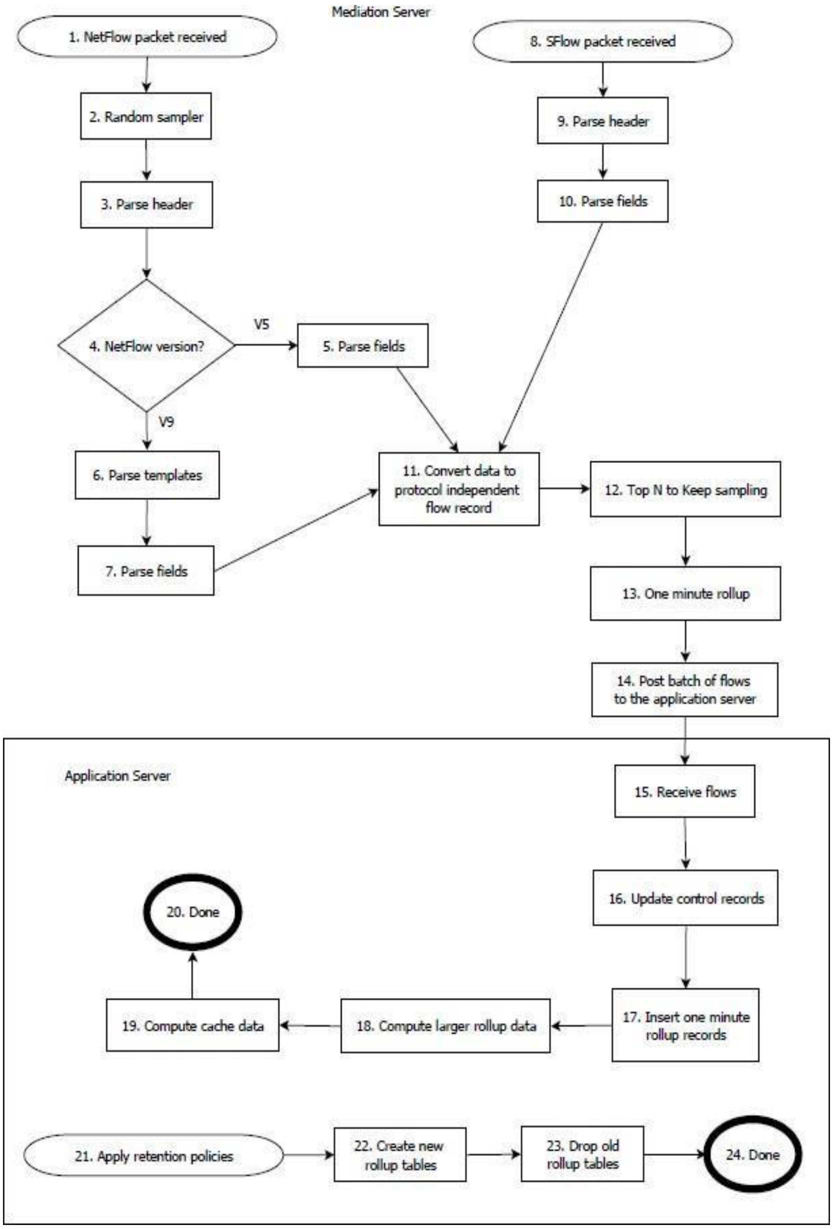
Mediation Server
1. NetFlow packet received
2. Random sampler
3. Parse header
4. NetFlow version?
V5
5. Parse fields
V9
6. Parse templates
7. Parse fields
8. SFlow packet received
9. Parse header
10. Parse fields
11. Convert data to protocol independent flow record
12. Top N to Keep sampling
13. One minute rollup
14. Post batch of flows to the application server
Application Server
15. Receive flows
16. Update control records
17. Insert one minute rollup records
18. Compute larger rollup data
19. Compute cache data
20. Done
21. Apply retention policies
22. Create new rollup tables
23. Drop old rollup tables
24. Done

Traffic Flow Analysis Lifecycle Legend

1 **NetFlow/IPFIX packet received** — Cruz received a NetFlow or IPFIX packet from a registered exporter. Cruzignores NetFlow packets from devices not registered asexporters.

2 **Randomsampler**—Cruz appliesrandom sampling to incoming packets. This currently applies only to NetFlow or IPFIX packets and not sFlow packets. The `NetFlowListenerMBean` attribute `SamplingRate` (1 by default) determines this behavior.

This value represents the average number of packets received before one is processed (the others are discarded). If `SamplingRate` is 10, then Cruz processes one of every 10 packets on average and discards the other nine. The default processes every received packet and does not discard any (at this step).

3 **Parseheader**—Cruz parses the NetFlow or IPFIX packet header. The header contains information like its version and the number of flows it contains.

4 **NetFlow version?** —Determines the packet version (NetFlow V5 or V9 or IPFIX).

5 **Parsefields**—Fields for a NetFlow V5 packet are determined by a standard template and parsed here.

6 **Parse templates**—For NetFlow V9 or IPFIX packets Cr uz must find the template that the packetheaderreferences. Cruzreceivestemplatesas NetFlowor IPFIX packetsfromthesame exporter sending flow packets. If Cruz has not yet received the referenced packet yet, then it sets the current flow packet aside until the template comes in. When that template comes in, it is parsed (there is a standard format for template packets) and then TFA will know how to parse the flow packets that reference it.

7 **Parse fields**—Cruz parses the fields of this flow packet using the header-referenced template.

8 **sFlow packet received**—Cruz received an sFlow packet from a registered exporter. Cruz ignores sFlow packets received from devices not registered.

9 **Parseheader**—Cruz parses the header of the sFlow packet. The header contains information like the version and the number of flows containedwithin the packet.

10 **Parse fields**—Cruz parses the fields of this packet according to the standard sFlow template.

11 **Convert data to protocol independent flow record** —TFA supports multiple traffic flow protocols (NetFlow, sFlow) but once Cruz parses the data within these packets, it does not depend on any specific protocol. Here, Cruz normalizes data into protocol-independent flow records.

12 **"Top N to Keep" filtering** —Cruz applies the "Top N to Keep" filtering here. This feature lets you set a maximum number of conversational flows per minute to keep, which in turn means that if Cruz receives more than this number in any given one-minute period, then it aggregates the rest into the "Other" category. Cruz ranks the received flows according to the number of estimated total bytes they report on and preserves all data only for the flows designated most significant by this measurement. It still preserves the total byte and packet data for the less significant flows, but the sender and receiver will be set to "Other".

13 **One-minuterollup** —Cruz collects and "rolls up" (aggregates) flows by conversation and at the end of every minute submits the resulting one minute rollup flows for further processing.

14 **Post batch of flows to the application server** —Most of this processing occurs in the mediation server (process or host). Once Cruz produces the normalized, protocol-independent flows, it adds them to a queue the mediation server posts to the application server. Once the application server receives these results, it inserts them into the database for later querying. You can monitor this queue through the JMX console > oware > service = NetFlowListenerMBean and the operation getQueueCount.

15 **Receive flows**—The application server receives flows from the mediation server.

16 **Update control records**—Every flow refers to control data, which includes the IP address of the sender and receiver, the AS number of the sender and receiver, the protocol and the application. Cruz updates the associated database records here.

17 **Insert one-minuterollup records**—Cruz inserts the one-minuterollup flows into the database so they can be queried later.

18 **Compute larger rollup data** —The one-minute rollup records contribute to larger rollup records, including those representing 10 minute, hourly, daily and weekly time intervals. At the end of every time interval TFA computes the appropriate rollup records.

19 **Compute cache data** — Sometimes the volume of data can be so high that queries summarizing it can take several minutes. To make querying more efficient, Cruz summarizes the data ahead of time and caches the results in the database.

20 **Done**— Done processing new traffic flow data.

21 **Apply retention policies** —Apply retention policies to add new rollup tables and/or drop old rollup tables.

22 **Create new rollup tables** — Cruz creates new rollup tables as necessary according to the settings in NetFlowRetentionMBean as specified in the attributes rolloverFrequencyValue1Min, rolloverFrequencyValue10Min, rolloverFrequencyValueHourly, rolloverFrequencyValueDaily and rolloverFrequencyValueWeekly.

23 **Drop old rollup tables**—The old rollup tables are dropped as necessary according to the data retention settings that can be edited by clicking on the wrench icon from the expanded TFA portlet.

24 **Done**—Done applying retention policies.

Best Practices: Performance Tuning Traffic Flow Analysis

Most Cruz packages limit Traffic Flow exporters through licensing, however, be aware that even small numbers of traffic flow exporters may overwhelm hardware resources. For example, even if you monitor as few as one to five nodes export data with a very high sampling rate and have high traffic volumes, the amount of data that needs to be inserted into the database may quickly exceed the insertion rate capacity of a single 7200 RPM disc. The performance limitation imposed by the demands of Traffic Flow Analysis is one motivation to tailor monitored flows to fit within the limits of your hardware.

You may also need to edit or fine-tune several default configurations to accommodate the volume of traffic flow data exported from the managed devices and to manage the resources (processing power, total memory) available to Cruz.

XML and properties files in `/owareapps/trafficanalyzer/server/conf/` contain these configuration settings. The `ta-service.xml` configures standalone environments, `ta-med-service.xml` configures distributed environments, and for additional configuration, edit `ta.properties`. Any edits require an application server and/or mediation server restart to take effect.

NOTE:

Best practice: Contact support for assistance before editing these files.

In `ta-service.xml` and `ta-med-service.xml`, you can edit the entry for the NetFlowListenerMBean, like the following:

```
<mbean code="com.dorado.netflow.NetFlowListenerMBean"
  name="oware:service=NetFlowListenerMBean">

  <attribute name="ReceiverUDPPort">9996</attribute>   <!-- # NetFlow UDP
  port -->

  <attribute name="ReceiverTopNToKeep">0</attribute> <!-- Number of rows
  per time bucket to keep -->
  <attribute name="RecvQueueMaxSize">10000</attribute>  <!-- # unparsed
  packets -->
  <attribute name="RecvQueueLowThreshold">80</attribute>  <!-- Low
  threshold for packet queue as pct -->
  <attribute name="UnparsedMaxSize">10000</attribute> <!-- # unparsed v9
  flows entries waiting on template -->
  <attribute name="TransportBatchSize">1000</attribute>  <!-- # NetFlow
  values (inserts) and session summaries (updates) -->
  <attribute name="TransportBatchTimeout">10000</attribute>  <!--
  milliseconds -->
  <attribute name="SpoolQueueMaxSize">10000</attribute>  <!-- # NetFlow
  session results -->
  <attribute name="SpoolBatchSize">50</attribute> <!-- # NetFlow session
  results flushed at a time to disk during overflow -->
  <attribute name="SpoolFileBufferSize">131072</attribute>  <!-- # bytes
  for I/O buffer -->
  <attribute name="SpoolFileName">@OWARE_USER_ROOT@/owareapps/
  trafficanalyzer/temp/tadata_spool.dat</attribute>
  <attribute name="SpoolMaxFileSize">41943040</attribute>  <!-- # bytes
  for disk allocation -->
  <attribute name="EnabledDNSLookup">true</attribute>  <!--#Enabed DNS
  lookup-->
  <attribute name="SamplingRate">1</attribute>  <!-- #Sampling Rate-->
  <depends>oware:service=NotificationProcessingMBean</depends>
  <depends>oware:service=HAServiceController</depends>
  <depends>oware:service=ClusterPrimaryDesignator</depends>
  <depends>jboss.j2ee:jndiName=RuleEngine,service=EJB </depends>
</mbean>
```

CAUTION:
Changes to this file do not persist if you upgrade. Best practice is to save a copy of the file elsewhere, and re-copy it into the correct location after upgrading.

Using random sampling to deal with high-volume scenarios

If the volume of incoming traffic flow data is be yond what can be processed, you can tune displays to optimize samplingRate. The sampling rate determines processing speed for incoming NetFlow packets (this does not apply to sFlow packets). The sampling rate should be 1 if you want to try to process everything. If the system is unable to keep up, then set this parameter higher, for example: 10, and restart the server. This makes Cruz process only 1 out ofevery 10 incoming NetFlow packet, at random. The object is to find the lowest value for this attribute that still allows the system to fully process all records that it tries to process.

To enable this feature, add the following like to installed properties with your chosen value of N to the right of the equals sign and the restart the applicationserver:

NetFlowListenerMBean.SamplingRate =

Advanced Performance Tuning

You might need to add entries to installed.properties to override the following system properties:

NetFlowListenerMBean. SpoolQueueMaxSize —The maximum number of traffic flow records storable in the spool queue during internal processing. Such processing occurs after Cruz receives them from the exporter and before they are saved to the database. Cruz temporarily stores such records in the spool queue if Traffic Flow Analysis receives more traffic flow packets from the exporter than it can process at once. This allows Traffic Flow Analysis to process these records later, packets arrive more slowly. If packets consistently arrive at a faster rate than Cruz can process them, the spool queue fills up and Traffic Flow Analysis discards them without processing.

ta.ThreadPoolSize — The default number of threads working to process traffic flow data.

ta.MaxThreadPoolSize — The maximum number of threads working to process traffic flow data.

Traffic Flow Database Advice

Traffic flow records average 300 bytes. Retention policies determine how long they stay in your database. So, 1-minuteflows, retained for 48 hours, 50% correlated (Correlation factors indicate the percentage of flows across all one minute buckets being aggregated that correlate based on conversation key data), mean 2.7 million flows per interval. If your retention keeps these 48 hours, then there are 2,880 retained intervals, and 7,776,000,000 potential rows. These rows times the 300-byte record size equal 2,173G in space, and 45,0 00 inserts per second. Typically, retention is for 1 minute, 10 minute, hourly, daily and weekly intervals, so for an overall picture of database needs, you would need to add all these intervals.

See Understanding Performance Monitoring on page 351 for more about calculating database space and hardware and tuning monitored data.

Traffic Flow Limitations

- If you receive no endpoint flow data, the Traffic Flow Analysis form appears empty when you select endpoints. Make sure you are receiving flow data before concluding the device or Cruz is defective.
- The flows received are just samples, and consequently are only as accurate as the polling interval and size of the flows sent. Fewer polling and smaller packets mean less accuracy. For example, If you only get 1 packet over the course of a file FTP it will not show the entire file size.

11

Actions and Adaptive CLI

The Actions Manager lets you manage actions like enabling monitors, file backups, resyncs and so on. These actions are typically limited in scope, and not that complex. On the other hand, it also manages Adaptive CLI (command-line interface) commands to run against devices which can be complex.

These commands amount to “mini-scripts” to query and configure those devices. In it, you can create commands to run against devices after the device driver has opened a connection to the devices. The driver handles logins, and general connection management. You can even initiate these actions with the application’s actions that target groups (see Discover Links for a Group of Devices, for example)—although if you delete a target group, such operations fail. Many drivers seed pre-configured command that appear listed when you first open this manager. For a brief overview of creating and using these, see How to: Create Adaptive CLI Examples on page 573.

Adaptive CLI’s Attributes capabilities let you insert variables in scripts. See Attributes on page 548 for the details. You can also assemble configurations made here as component Tasks to execute with other component Tasks. You can even use this capability to include Perl scripts within Cruz. See Perl Scripts on page 571.

NOTE:

You can have Actions maintain lists like ACLs, and when these change, in the Adaptive CLI script, push the updated list out to the appropriate devices.

Adaptive CLI commands let you map several vendor-specific commands to a single action, so you could, for example, query two types of devices throughout the network for their MAC addresses with a single action. Adaptive CLI actions can also help you debug more complex scripts that either query or configure devices.

The Adaptive CLI manager displays a list of *Configure* and *Show* commands (the *Command Type*) with a *Name*, *Description* and the *Last Run Date*. You can filter what appears in this manager with the fields at its top.

NOTE:

The contents of the Action Portlet vary, depending on the installed options in your package.

△ CAUTION:

Particularly for Adaptive CLI, and possibly for other Cruz capabilities, the level of access to devices must match the desired effect. If Cruz’s login to a device permits only read access, then Adaptive CLI configuration commands which require write capabilities will not be effective.

Using Adaptive CLI

Youcan quickly take a set of commands or configuration file snippet from a device, copy it directly into the Script editor, mark it up, and save it as a working CLI.

When using the CLI Format, The Adaptive CLI tool will prompt you to create new attributes based upon your script markup. This lets you quicklycreate a script and schema to create an ACLI. If you have attributes that are mainly simple String attributes, this is a very quick and automated approach.

Using Perl in Adaptive CLI

If you need conditional logic that goes beyond simple scripting, you can use Perl in Adaptive CLI. The example below checks to see if a String Attribute is empty (null) or not. If the String attribute (`ShowCmdString`) hascontent, the showcommand with `ShowCmdString` asaparameter goes to the device. Otherwise, the Perl script skips or excludes this statement.

Embedded CLI Example:

```
[IF ShowCmdString]
 Show [ShowCmdString]
[ENDIF ShowCmdString]
```

You could use the CLI format for the above example, but if you need to check attributes of other types, besides String, then you must switch to Perl.For example:

Boolean `myFlag` equals True:

```
if ($myFlag)
{
     …
}
```

Integer `myInt` greater than zero:

Example:

```
if ($myInt > 0)
{
     …
}
```

To check whether a string is a particular value—like from a valid values list entry assigned to the String attribute—then you must also use Perl. The CLI format only can test if the String exists. It cannot validate its value when populated. For example: EncapsulationType = “VLAN-CCC”, “VLAN-TCC”,… Youcannot do this checkwith the CLI Format: `[IF EncapsulationType = “VLAN-TCC”].` Instead, use a Perl script with a statement likethis:

```
If ($EncapsulationType eq “VLAN-TCC”)
{
     print “set encapsulation $EncapsulationType\n”;
}
```

If any attributes in your script are a List (Collection), the only way to loop through the list’s items during the Adaptive CLI execution is to use Perl. For example: Processing a List of Strings:

```
$count = 0;
foreach @MyCommandList)
```

```
{
   print (“$MyCommandList[$count]\n”);
   $count++
}
```

Actions Portlet

The Actions Portlet lets you manage actions like Adaptive CLI, backups, change management actions, and so on. The list of actions available to your system depends on the exact configuration you have installed. This portlet is the primary access point for Adaptive CLI editing.

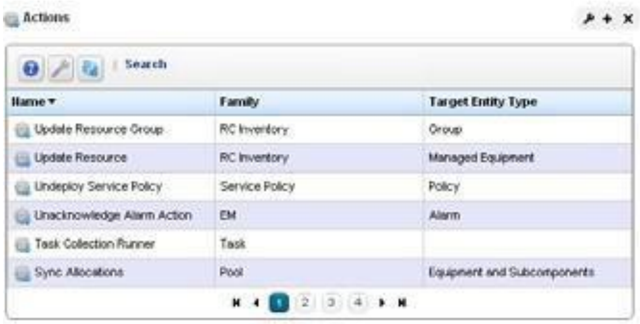

Name ▾	Family	Target Entity Type
Update Resource Group	RC Inventory	Group
Update Resource	RC Inventory	Managed Equipment
Undeploy Service Policy	Service Policy	Policy
Unacknowledge Alarm Action	EM	Alarm
Task Collection Runner	Task	
Sync Allocations	Pool	Equipment and Subcomponents

The summary portlet displays columns with the *Name*, *Family*, and *Target Entity Type* for the listed Action. The Family column describes the type of Action. Right-click and select *New > Adaptive CLI* to create a new action. You can also import/export actions and share actions with other users on your system.

△ CAUTION:

For Adaptive CLI to be fully functional, you may need to install it on your servers. Refer to the *Cruz Installation Guide* for Perl installation instructions.

To configure and schedule groups of actions, right-click in the Schedules portlet, and create an *Action Group*. This lets you run several actions and configure their order and targets.

Expanded Actions Portlet

The expanded portlet has the same right-click menu as the summary portlet, and addscolumns for *Description*, *Last Web Service ID*, *Access Level*, *Web Service Deployment,* and *Supports Groups.*

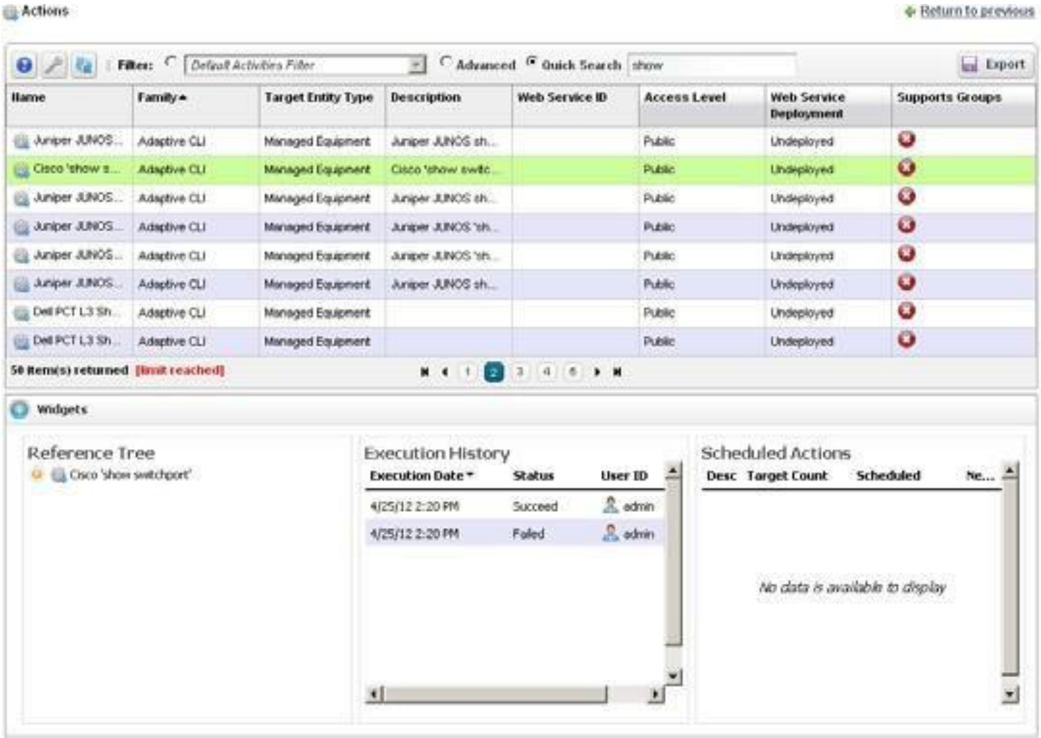

Filter: ○ Default Activities Filter ○ Advanced ● Quick Search show

Name	Family	Target Entity Type	Description	Web Service ID	Access Level	Web Service Deployment	Supports Groups
Juniper JUNOS...	Adaptive CLI	Managed Equipment	Juniper JUNOS sh...		Public	Undeployed	☒
Cisco 'show s...	Adaptive CLI	Managed Equipment	Cisco 'show switc...		Public	Undeployed	☒
Juniper JUNOS...	Adaptive CLI	Managed Equipment	Juniper JUNOS sh...		Public	Undeployed	☒
Juniper JUNOS...	Adaptive CLI	Managed Equipment	Juniper JUNOS 'sh...		Public	Undeployed	☒
Juniper JUNOS...	Adaptive CLI	Managed Equipment	Juniper JUNOS 'sh...		Public	Undeployed	☒
Juniper JUNOS...	Adaptive CLI	Managed Equipment	Juniper JUNOS sh...		Public	Undeployed	☒
Dell PCT L3 Sh...	Adaptive CLI	Managed Equipment			Public	Undeployed	☒
Dell PCT L3 Sh...	Adaptive CLI	Managed Equipment			Public	Undeployed	☒

50 item(s) returned [limit reached]

Reference Tree

Cisco 'show switchport'

Execution History

Execution Date	Status	User ID
4/25/12 2:20 PM	Succeed	admin
4/25/12 2:20 PM	Failed	admin

Scheduled Actions

Desc	Target Count	Scheduled	Ne...

No data is available to display

The expanded portlet also has snap panels to display Reference Tree connections between the selection and other elements within Cruz, as well as an Execution History panel listing *Device Name*(s), *Execution Date* and *Status* for the selected Action, and a Scheduled Actions panel cataloging any Schedules for the selected Action. Right-click a Schedule to edit, execute or delete it.

The Execution History snap panel displays history by device. Right-click to see the details of what occurred when the selected action ran against a device (*Execution Details).*

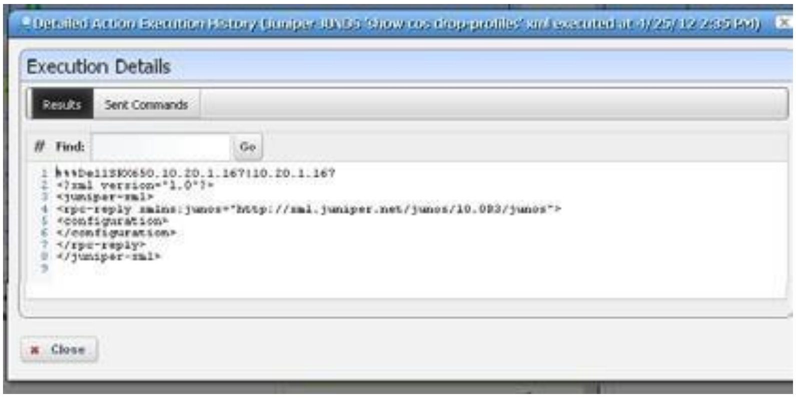

Detailed Action Execution History (Juniper JUNOS 'show cos drop-profiles' xml executed at 4/25/12 2:35 PM)

Execution Details

```
***Dell35X0650.10.20.1.167!10.20.1.167
<?xml version="1.0"?>
<juniper-xml>
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.0R3/junos">
<configuration>
</configuration>
</rpc-reply>
</juniper-xml>
```

The Execution Details panel displays tabs showing the *Results* of running an Adaptive CLI, and the *Sent Commands*.

You can also *View Job* to see a screen like the Audit Trail Portlet on page 150, or Delete to remove a listed Action record from the list.

Right-click menus on the Actions portlet can include the followingitems (these vary, depending on the Action’s family):

New/Edit —Lets you create or modify a selected action in the Adaptive CLI Editor, described below. Youcannot modify system-supplied Adaptive CLIs, but you can edit any that you create.

Execute—Execute the selected Action. This typically displaysa target equipment selector screen, and a screen where you can configure any parameters necessary for execution, then a screen like the Audit Trail Portlet on page 150. Cruz validates the parameters before executing the Adaptive CLI. If a parameter is invalid, for example a blank community name in the Juniper Community Adaptive CLI, Cruz logs a validation error to the audit trail. In this case the Adaptive CLI is not executed and leaves behind no history record.

Some Adaptive CLI scripts alsoletyou *Preview* what is sent the deviceinasubsequent screen. This does not appear in the execution of Targetless, and Multi-target Adaptive CLIs. Some actions are configured to target groups, too.

Details—Opens a screen displaying the Reference Tree, Execution History, and Action Details for the selected Action.

Web Services —You can elect to *Deploy/Undeploy* or *Export WSDL* to create a web service from the selected Action.

Deploy/Undeploy Web Service –Deploy or undeploy the selected activity as a web service. See Web Service Deployment Features on page 598 for more. Refer to the *Web Services Guide* for detailed information.

Export WSDL – This exports the WSDL for the selected activity. You must select the file name and location. Web Services Description Language (WSDL) is an XML format for the description of network servicesas a set of endpoints operating on messages containing either document-oriented or procedure-oriented information. See Web Service Deployment Features on page 598 for more.

NOTE:

If you have the optional Service Center installed, its Services and Work Orders automatically become Adaptive CLIs, and are therefore available to implement with Web Services.

History — Displays the history of the selected action.

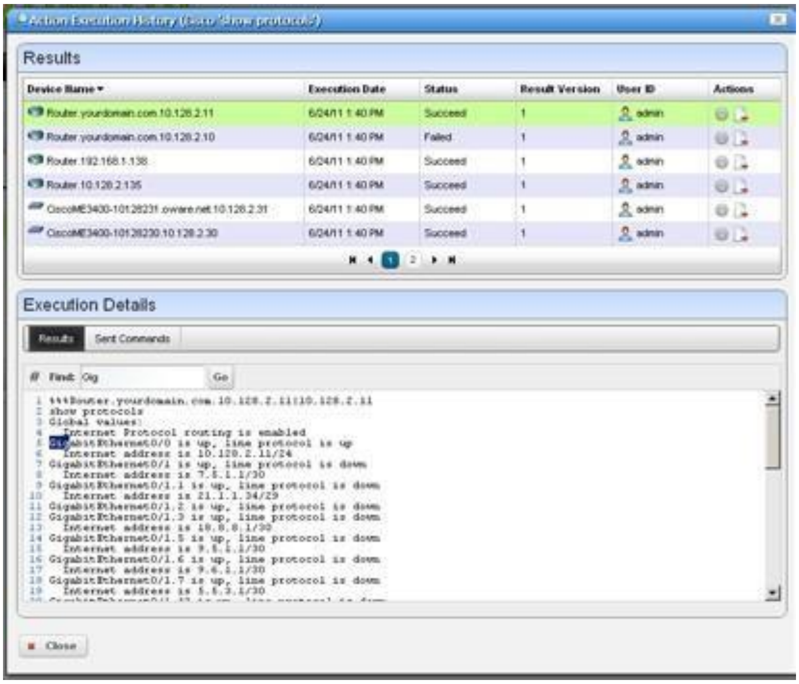

Action Execution History (Cisco 'show protocols')

Results

Device Name ▾	Execution Date	Status	Result Version	User ID	Actions
Router.yourdomain.com.10.128.2.11	6/24/11 1:40 PM	Succeed	1	admin	
Router.yourdomain.com.10.128.2.10	6/24/11 1:40 PM	Failed	1	admin	
Router.192.168.1.138	6/24/11 1:40 PM	Succeed	1	admin	
Router.10.128.2.135	6/24/11 1:40 PM	Succeed	1	admin	
CiscoME3400-10128231.oware.net.10.128.2.31	6/24/11 1:40 PM	Succeed	1	admin	
CiscoME3400-10128230.10.128.2.30	6/24/11 1:40 PM	Succeed	1	admin	

Find: Gig

```
***Router.yourdomain.com.10.128.2.11|10.128.2.11
show protocols
Global values:
  Internet Protocol routing is enabled
GigabitEthernet0/0 is up, line protocol is up
  Internet address is 10.128.2.11/24
GigabitEthernet0/1 is up, line protocol is down
  Internet address is 7.5.1.1/30
GigabitEthernet0/1.1 is up, line protocol is down
  Internet address is 21.1.1.34/29
GigabitEthernet0/1.2 is up, line protocol is down
GigabitEthernet0/1.3 is up, line protocol is down
  Internet address is 10.8.8.1/30
GigabitEthernet0/1.5 is up, line protocol is down
  Internet address is 9.5.1.1/30
GigabitEthernet0/1.6 is up, line protocol is down
  Internet address is 9.6.1.1/30
GigabitEthernet0/1.7 is up, line protocol is down
  Internet address is 5.5.3.1/30
```

In the *Results* (top of screen panel) click to select the device for which you want additional information, and the *Execution Details* panel displays the *Results* of execution in one tab and the *Sent Commands* in another.

Notice that you can *Find* text within a result (click *Go* to repeat the find). You can also see the bottom panel if you right-click a single execution within the *Execution History* snap panel in the Expanded Actions Portlet.

If you select two executions in the top panel (or in the *Execution History* snap panel and right-click), a comparison appears.

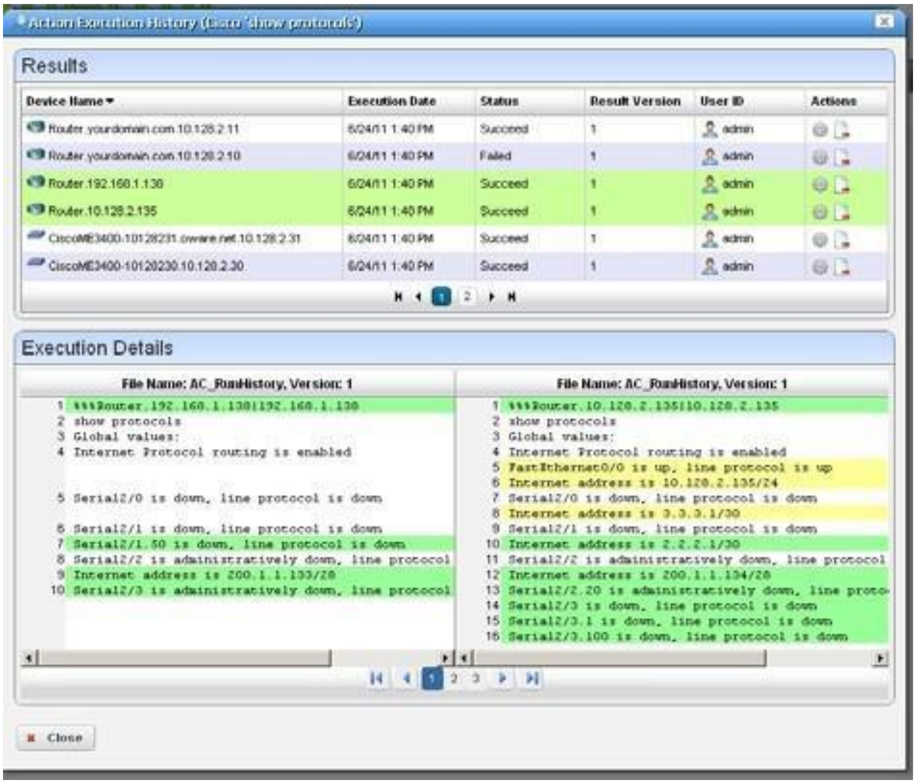

Results

Device Name ▾	Execution Date	Status	Result Version	User ID	Actions
Router.yourdomain.com.10.128.2.11	6/24/11 1:40 PM	Succeed	1	admin	
Router.yourdomain.com.10.128.2.10	6/24/11 1:40 PM	Failed	1	admin	
Router.192.168.1.138	6/24/11 1:40 PM	Succeed	1	admin	
Router.10.128.2.135	6/24/11 1:40 PM	Succeed	1	admin	
CiscoME3400-10128231.oware.net.10.128.2.31	6/24/11 1:40 PM	Succeed	1	admin	
CiscoME3400-10120230.10.120.2.30	6/24/11 1:40 PM	Succeed	1	admin	

Execution Details

File Name: AC_RunHistory, Version: 1

```
%%%Router.192.168.1.138|192.168.1.138
show protocols
Global values:
Internet Protocol routing is enabled


Serial2/0 is down, line protocol is down

Serial2/1 is down, line protocol is down
Serial2/1.50 is down, line protocol is down
Serial2/2 is administratively down, line protocol
Internet address is 200.1.1.133/28
Serial2/3 is administratively down, line protocol
```

File Name: AC_RunHistory, Version: 1

```
%%%Router.10.128.2.135|10.128.2.135
show protocols
Global values:
Internet Protocol routing is enabled
FastEthernet0/0 is up, line protocol is up
Internet address is 10.128.2.135/24
Serial2/0 is down, line protocol is down
Internet address is 3.3.3.1/30
Serial2/1 is down, line protocol is down
Internet address is 2.2.2.1/30
Serial2/2 is administratively down, line protocol
Internet address is 200.1.1.134/28
Serial2/2.20 is administratively down, line proto-
Serial2/3 is down, line protocol is down
Serial2/3.1 is down, line protocol is down
Serial2/3.100 is down, line protocol is down
```

This has the same color coding as you would see comparing configuration files. Lines that differ between the two Adaptive CLI results appear highlighted green. Lines that are missing in one, but that appear in another appear highlighted red. Added lines appear highlighted in yellow. Use the right/left arrows or the page numbers at the bottom of the screen to page through the side-by-side comparison.

Audit—Opens an Audit Trail Viewer for the selected Action. See Audit Trail Portlet on page 150 for details.

Show Last Results — Show the last execution details (like history for a single run).

Schedule — Schedule the selected Action. See Scheduling Actions on page 587 for details.

Delete — Remove the selected Action from the list.

Import/Export —Import or Export a file representations of the ACLIaction selected. Best practice is to export single actions, or actions that share a Schema, rather than exporting all actions at once.

If you have multiple ACLIs sharing the same Schema and change that Schema best practice is to retest the other ACLIs using that schema to ensure no unintended side effects occur.

Adaptive CLI Editor

This editor creates new Adaptive CLIs When you click *New*, or *Edit* after, selecting an existing command, the command editor screen opens. You can create *Configure Commands*, *External Commands*, and *Show Commands*.

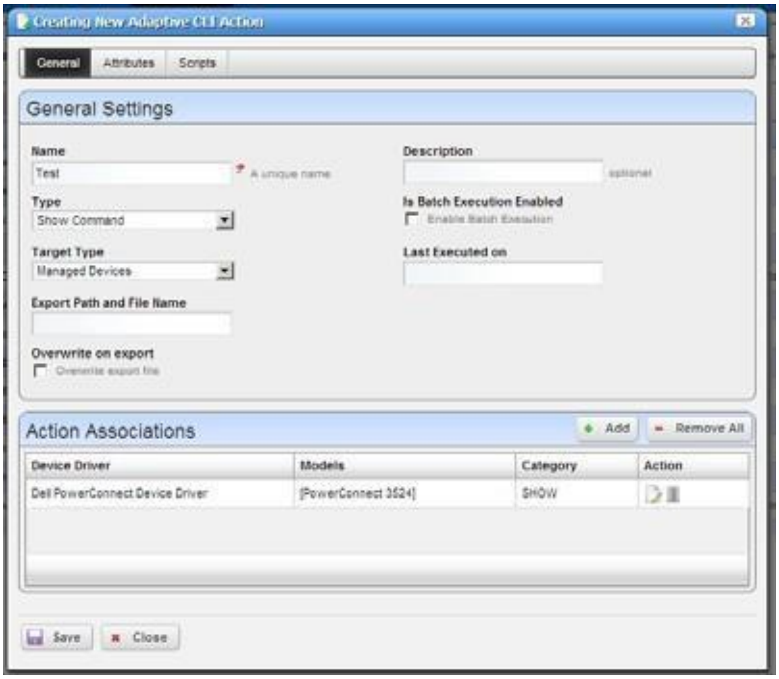

Name: Test * A unique name

Description: (optional)

Type: Show Command

Is Batch Execution Enabled: □ Enable Batch Execution

Target Type: Managed Devices

Last Executed on:

Export Path and File Name:

Overwrite on export: □ Overwrite export file

Device Driver	Models	Category	Action
Dell PowerConnect Device Driver	[PowerConnect 3524]	SHOW	

The editor screen has the following tabs (the ones that appear depend on the type of command you are editing):

- General
- Attributes
- Scripts

The Adaptive CLI Manager logs into devices in enable mode by default. For most configuration commands (and even some show commands), you must typically first set the device to its configuration mode. For example: The first Adaptive CLI Manager command must be `config t` (Juniper E-Series) or `edit` (Juniper M/T series) to initiate the router's config/edit mode. Cruz also validates entries. If saving fails, a red “X” appears next to required omitted entries.

Click *Save* to preserve the Adaptive CLI you have configured. Clicking *Close* does not save your configuration.

General

The following are parameters to configure in this panel:

Name — A unique identifier for this action. For example: “Retrieve MyDevice MAC addresses.”

> Fora new action to appear on the right-click Action menu, begin its name with the vendor name. For example, *Force10-showversion* would appear under Actions in that menu. Otherwise, it appears under and Adaptive CLI classification.

Description — A text description of the action.

Type—Select a type from the pick list (*CLI Configure Command, External Executable, Config File Generation* or *CLI Show Command*).

The *External Executable* command refers to a script. Making this an ACLI means Cruz can schedule such scripts or include them in a workflow. See External Executable on page 565 for more about these.

NOTE:

You can use Cruz's optional Proscan policies to scan the results of Adaptive CLI show commands for compliance, and trigger actions (alarms, e-mail, and so on) based on their contents. See .

See How to: Use Config File Generation on page 561 for more about that.

Target Type — Select a type of target from the pick list (*Card*, *Equipment and Subcomponents*, *Interfaces*, *Managed Devices*, *Ports*). Adaptive CLI targets can also be *None* (*Targetless)*. On execution, if you create an Adaptive CLI type with port target, then the selection view panel lets you choose ports. When the Adaptive CLI type is *External* then Target Type can be *None*; otherwise it is not an option

If you want the target to be a Cruz group after you right-click an Adaptive CLI, and select *Execute*, select *(Group Membership*) in the target selector filter. Click the icon to the left of the *Go* button, and to the right of the empty field to see available groups. Select the group(s) and click *Done*. Click the *Go* button. Control click if you want a subset of the group's devices, or simply click *Add All* and then click *Done.*

Export File Location—This is a file name and path (`C:\mypath\myfile.txt`) where you elect to store the result of an adaptive CLI execution. You must specify an extension for the file and may specify the variable `$IPAddress` in the filename for pattern substitution.

Overwrite on Export—Check to overwrite the result file. Thisoverwrites any existing results file with new results (if checked). If it is unchecked, any new results append to the exported file, with a time/date stamp and target-identifying information.

Is Batch Execution Enabled — Check to allow consolidation of related Adaptive CLI scripts, provided theassociateddevicedriversupports such consolidation whenprovisioningaservice. (Currently supported by the Juniper JUNOS driveronly.)

Batching is valuable for instances like the following: if an Adaptive CLI-provisioned service has 10 sub-services, Cruz runs commands for the first service, then if it's successful, commits, and logs off. Then Cruz repeats this procedure nine times more, logging on, committing and logging off for each command. If batching is turned on, then Cruz sends the 10 Adaptive CLIs to the device as a single unit before committing and logging off. (This logic does not apply if you are running a procedure against 10 devices.)

Batching is best practice for Juniper devices, since if one line of a command fails, the device rolls back the entire block of commands. Cisco devices typically skip and do not commit failing lines.

Last Executed On— Displays the last execution date. This is blank for new Adaptive CLIs.

Action Associations

Click the *Add* button to add associations to vendors and device models. For example, you can confine an Adaptive CLI to Dell devices, even to certain Dell models. When you right-click your discovered Dell device in the Managed Resource portlet, the associated Adaptive CLIs appear listed among the available actions you can request.

Attributes

Adaptive CLI commands let you configure modifiable *Attributes* as part of the comm and you send to the selected equipment.

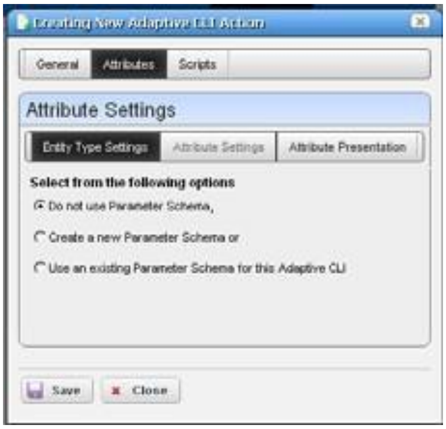

Select from the following options

● Do not use Parameter Schema,

○ Create a new Parameter Schema or

○ Use an existing Parameter Schema for this Adaptive CLI

Entity Type Settings

Use the radio buttons to select from the following options:

- Do not use Parameter Schema
- Create a new Parameter Schema
- Use an existing Parameter Schema for this Adaptive CLI

Sharing a schema rather than creating a new one with each Adaptive CLI lets you use the same attributes in complementary scripts. For example, one script may create an entity, while another removes it. In this case, the valid values, labels, and so on for the attributes are always going to be the same in both create and delete Adaptive CLIs. This means sharing the same schema is both safe and easy. Either script can mark unused attributes as “Not applicable.”

Do not use Parameter Schema

This option does not save a set of standard attributes to re-use later. Go directly to the Scripts tab to create this type of Adaptive CLI.

Create a new Parameter Schema

Click the *New* button and the schema screens appear.

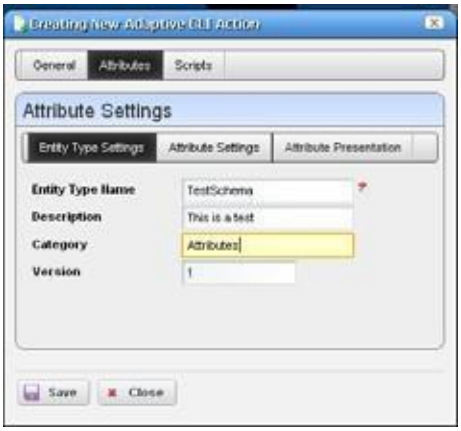

Entity Type Name: TestSchema *
Description: This is a test
Category: Attributes
Version: 1

Entity Type Settings

The *Entity Type Settings* tab has the following fields:

Entity Type Name — An identifier for the schema.

Description — A text description for the schema.

Category — A category for the schema.

Version — An automatically-created version number.

Attribute Settings

Click the *New Attribute* button and select the attribute type and open editor panel and configure the attribute. Configured attributes appear in a tree to the left of the editor panel. Click a listed attribute to edit it after it has been created.

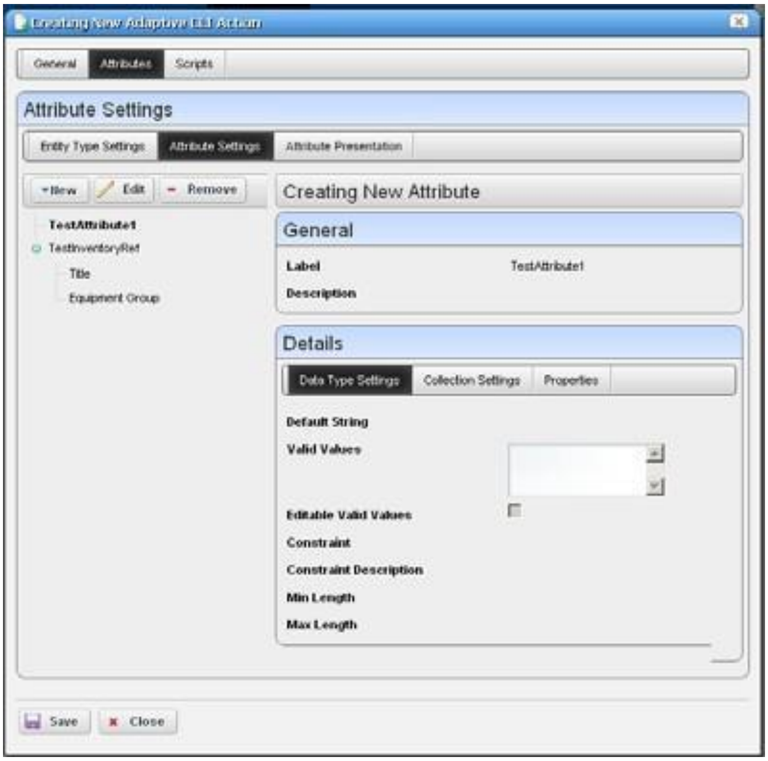

Label TestAttribute1

Description

Default String

Valid Values

Editable Valid Values □

Constraint

Constraint Description

Min Length

Max Length

The editor panel has the following fields:

Label—An identifier for the attribute. These can have spaces, but not underscores, unless your package is 7.2.3 or later, which supports both.

Description—A text description for the attribute.

The following tabs may appear, depending on the type of attribute you are configuring (some are absent). Additional fields may appear, depending on the attribute type you are configuring:

Datatype Settings

Default Value — An optional default value for the attribute.

Collection Settings

Is Collection? — Check to classify this attribute as a collection.

Allow Duplicate Values—Check to enable allowing duplicates.

Allow Reordering—Check to enable allowing reordering.

Collection Min/Max Length — Enter the minimum/maximum number of characters in this attribute.

Properties

Upper/Lower Case — Check to validate on case.

Case Insensitive — Validation ignores case.

Multi Line Text — Check to enable multiline text.

One Way Encrypt — Check to encrypt.

Truncate — Truncate the attribute.

Attribute Settings

You can create new attribute schemas. See Attribute Editor Panels below for information about different datatypes' fields. Once you create a set of attributes, they remain available for re-use as a schema, or collection of attributes. To identify schemas, enter the following fields:

Label — A unique, mandatory identifier for the collection of attributes.

Description — A text description of the entity.

Click *New* to create or select an attribute in the displayed tree and click *Edit* to open an editor where you can create or modify attributes. Select an attribute and click *Remove* to delete it from the list.

Attribute Editor Panels

The following panels appears, depending on the at tribute type selected from the pick list. The fields in the editor depend on this selection. Available types include *Boolean, Coded Value, Date, Decimal, IP Address, Integer, Long, Inventory Reference,* and *String.* The following fields appear for each of these types (omitting redundant fields):

NOTE:

Configure the data type of an attribute before you save a task. After attributes are in Scripts, you cannot change the data type.

Boolean

Default Value — Check for *True*.

Coded Value

Default Coded Value— Enter the default coded value. If an attribute a Coded Value then enter valid values in the format of NUMBER:Display Label. For example:

10:Hello World

20:Hello Moon

Without this pattern a validation error appears. Coded values become a Drop Down (Combo Selection) at runtime containing the Display labels within it (like Hello World, Hello Moon). Selecting one gives the script the numeric value (If users select Hello World, the value the script gets is 10)

The default appears by default in this list ofalternatives. Enter any other alternatives below this field in the *Valid Values*.

Valid Values—Enter a valid value in the line above the table of valid values, then click the green + to add the value entered to the list. Click the *Remove* icon (the red -) to delete a selected value. These must be formatted like the *Default Coded Value.*

Date

Default Value—Enter a default date or use date icon to display a calendar where you can select one. Click off the calendar to make it disappear.

Valid Values—Enter valid date values above the list and click the green plus to add them to the list.

Decimal

Default Value — Enter a single or range of default decimal values.

Constraints —Enter a range of acceptable numbers separated by a colon. For example, Constraints = 2:4096. At runtime, a field where you can enter numbers. validates that entered numbers are between 2 and 4096 when running the Adaptive CLI. If you enter a number outside this range, a validation message appears, and the attribute name turns red. You can also include several numbers or ranges separated by the pipe (|) character. If you specify a range, the lowest number must be to the left of the highest number. For example, 9:99|999:1010.

Valid Values — Enter valid decimal range values and click the green + (the red - removes them). You can manage these as described in Coded Value above.

IP Address

See also Validating IP Address Variables on page 553.

Default Value — Enter a default IP Address.

Valid Values—Enter valid values as described in Coded Value above. Check *IP Mask*, *Subnet*, *Allow 32 Bit Mask*, and *Allow Any Valid Ip* in the *Properties* tab if you want the values entered to be those.

Editable Valid Values — Check to enable editing of default or entered IP addresses.

Integer

Default Value — Enter a default integer.

Constraints—Enter a range of acceptable numbers separated by a colon. You can also include several numbers or ranges separated by the pipe (|) character. If you specify a range, the lowest number must be to the left of the highest number. For example, 9:99|999:1010.

Valid Values — Enter ranges of valid values as described in Decimal above.

Editable Valid Values—Check to enable editing of default or entered integer.

Long

Default Value — Enter a default long.

Constraints—Enter a range of acceptable numbers separated by a colon. Youcan also include several numbers or ranges separated by the pipe (|) character. If you specify a range, the lowest number must be to the left of the highest number. For example, 9:99|999:1010.

Valid Values — Enter ranges of valid values as described in Decimal above.

Inventory Reference

Select the *Reference Type* entity with the list that appears when you click the green plus (+), then use the side-by-side widget’s arrows to move available attributes from *Available* to *Selected*. You can change the *Reference Type* by deleting it with the red minus (-), then selecting a new type with the green plus.

String

Default String — Enter a default string.

Valid Values — Enter valid values as described in Coded Value above.

Editable Valid Values — Check to enable editing valid values.

Constraint —Enter the regular expression constraints, if any, on the string attribute.

Constraint Description—Enter the message to appear if the regular expression constraints are not met.

Min/Max Length—Enter the minimum/maximum number ofcharacters in a valid string.

Click *Apply* to accept your edits for the attribute or *Cancel* to abandon them.

Use an existing Parameter Schema for this Adaptive CLI

Select this, and a *Select Existing* button appears. Clicking this button opens a selector where you can select from previously configured attribute schemas (collections of at tributes) to use in the Adaptive CLI you are configuring.

Validating IP Address Variables

Programmatically, IP address attributes support four extended properties: IP_MASK, SUBNET, ALLOW_32_BIT_MASK, and ALLOW_ANY_VALID_IP. The state of the first two largely defines Cruz's responses.

IP_MASK—Determines whether Cruzacceptsan IPaddress ORasubnet/subnet mask. Thevalue accepted is an IP address attribute when false, subnet/subnet mask when true.

SUBNET—This property determines whether a subnet value must be provided or not, and controls display of the subnet portion of the widget. Valid subnet values are 1-31.

By default, when both of the above are false, the attribute only accepts valid IPv4 addresses. For example: 10.10.10.4

If IP_MASK is false, and SUBNET is true then Cruz accepts any valid IP address with a subnet specified. The address must be an IP within the specified subnet. Forexample, 10.10.10.4/24 is a valid entry whereas 10.10.10.0/24 is invalid since it represents the subnet id, not an actual address within the subnet.

If IP_MASK is true and SUBNET is false, then Cruz accepts one of the 32 valid subnet masks. The widget displays pick list for user to choose from. For example, 255.255.255.0

If IP_MASK is true and SUBNET is true, then Cruz accepts a subnet id (the first IP address within a subnet). For example, 10.10.10.0/24, with 10.10.10.0 as the first address within the subnet spanning 10.10.10.0 to 10.10.10.254. Entering an IP addresswithin the subnet, say 10.10.10.4/24, the attribute would convert that to 10.10.10.0/24

ALLOW_32_BIT_MASK—Validsubnet values are between 1 and 31. To extend this to support a 32-bit subnet, which is essentially a single IP address (10.10.10.4/32), set the ALLOW_32_BIT_MASK property.

ALLOW_ANY_VALID_IP —To accept either an IP address, IP address and subnet or subnet, then IP_MASKremainsfalse, SUBNET istrue. Withthe ALLOW_ANY_VALID_IP true, the subnet field is optional, and Cruz disables any requirement that a subnet id be specified. Basically, the only validation is that a valid IP address is entered. Forexample, in this configuration, 10.10.10.4, 10.10.10.4/24 and 10.10.10.0/24 would all be valid.

Attribute Presentation

This panel lets you reorder attributes by selectin g an attribute, then clicking the up/down arrows. Exporting or importing a reordered Action preserves the configured order of attributes.

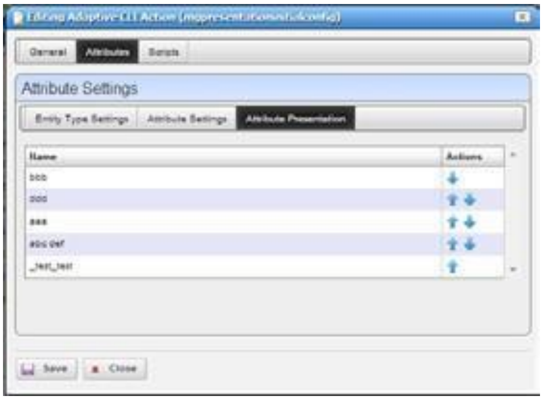

Name	Actions
bbb	↓
ddd	↑ ↓
aaa	↑ ↓
abc def	↑ ↓
_test_test	↑

If you change the attribute order after scheduling the Action, it does not affect the execution of the scheduled Action. Ifyou usethesame Entity Type(schema) foranother Actionattributes appearin the same order.

Scripts

This screen manages the Adaptive CLI scripts created to query (show) devices or configure them. Cruz runs only one script per target.

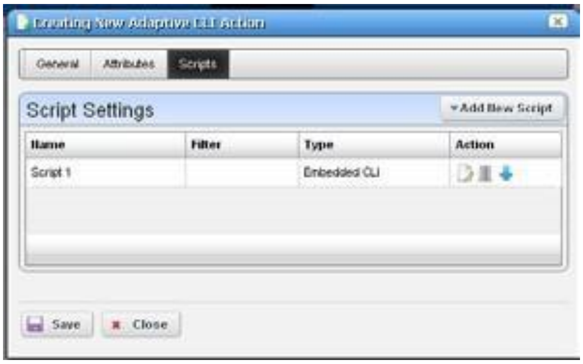

Name	Filter	Type	Action
Script 1		Embedded CLI	

Notice you can order multiple scripts with the arrow(s) to the right of a listed script. Only one schema of attributes exists for each Adaptive CLI, so the same attribute(s) appear when you construct each script.

Cruz uses the script’s filter to match the target. Forexample, imagine two scripts for which the first has filter = target.type = SWITCH, and the second has no filter. Then only SWITCH devices run the first script and quit. All remaining targeted devices do not run first script. Instead they run the second script since that script has no filter. Only one script runs on the selected target equipment. The ordering lets you to make the most efficient use of that one-run-per-target pattern.

Script Settings

Click *Add New Script* to create a new item in those listed at the top of this screen or select and item and click the *Edit* icon to its right to alter it. When you create a new script, you must select either *Embedded CLI* or *Perl*. Embedded CLI scripts are command-line interface (CLI) interactions. See Perl Scripts on page 571 for more about using Perl.

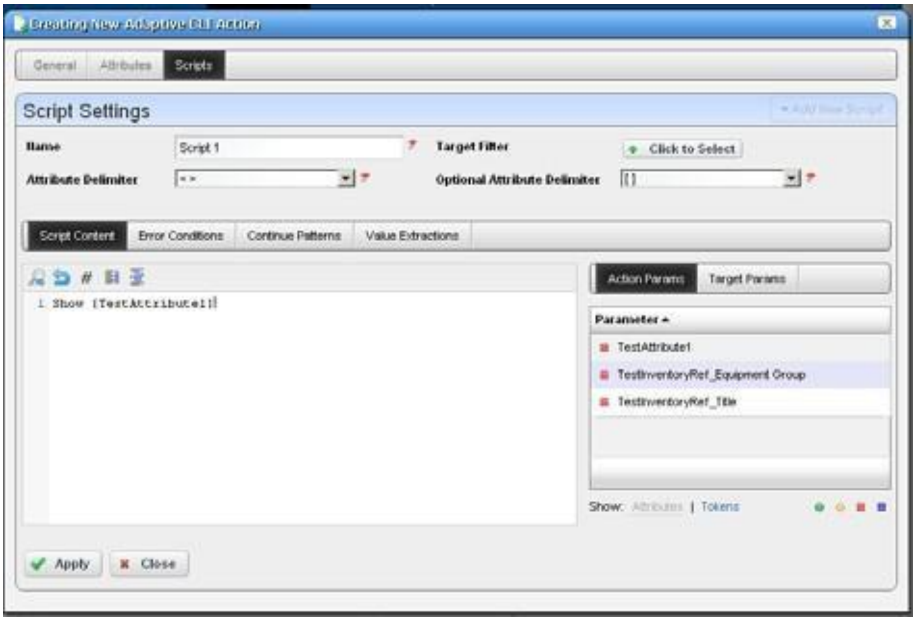

Name: Script 1

Target Filter: Click to Select

Attribute Delimiter: < >

Optional Attribute Delimiter: [ ]

```
1 Show [TestAttribute1]
```

Parameter
TestAttribute1
TestInventoryRef_Equipment Group
TestInventoryRef_Title

Clicking the *Delete* icon removes a selected item. Notice that the up/down buttons to the right of the list allow you to re-order selected item s (they run from top first to bottom last).

See Attribute Appearance and Validation for a description of what constitutes a valid attribute.

NOTE:

You must mark an attribute as required before adding it to the script. If you add an attribute before you markit as required, you must remove it from the script, markitasrequired, then re-addit. Insome browsers,afteraddingtheattributes youmustclickin thescriptscreentoensurethatthechanges persist.

Name — Enter an identifier for the script you are creating or altering.

Target Filter — Click the plus (+) to create a filter that describes the target for this script. For example, this filter could confine the action of the configured script to devices from a certain vendor, or only devices with an operating system version later than a certain number. Since you can have several scripts, those Adaptive CLIs with a single label (“Show Users,” for example) could thereforecontain severalscripts withsyntax appropriate to avarietyofdevices and operating systems.

CAUTION:

Adaptive CLI supports only filters that select the Managed Equipment type of device.

Attribute Delimiter — The delimiter(s) you select from the pick list here surround the attributes you designate as mandatory. See Adaptive CLI Script Language Syntax on page 570 for more about these.

Optional Attribute Delimiter — The delimiter(s) you select from the pick list here surround the attributes you designate as optional. See Adaptive CLI Script Language Syntax on page 570 for more about these.

All but *Delete* open a script editor with the following panels:

- Script Content
- Error Conditions
- Continue Pattern
- Prerequisite Validation
- Value Extraction

Script Content

On the left, you can enter text, *Search* by clicking the magnifying glass, and use *Cut, Copy, Paste, Undo, Jump to Line # , reformat* . The *Attributes* appear under *Target Params* on the right of this text entry screen. Double-click an attribute to insert it unless you are writing a Perl script; this feature does not work for Perl. Right-click the previously configured attributes in this panel to designate them as *Mandatory, Optional, Not Applicable* or *Non-Configuration* in a context menu that appears when you right-click.

NOTE:

Cruz does not send *Non-Configuration* attributes to the device with the script. These are comments that can serve to remind users of critical information. For example, you can make *Non-Configuration boolean attributes into a checklist for someone executing a script, and the history of this script can record whether* Cruz *made these checks when the script ran.*

Notice that the *Search* also permits Regular expressions.

You can also enter two types of script language here. See Adaptive CLI Script Language Syntax on page 570 for a description of the internal *If* capabilities. If you need more elaborate scripting, you can also use Perl scripts to send text to devices. See *Perl Scripts* on page 571 for a description of those capabilities.

Click *Apply* to accept the script you have configured.

NOTE:

Some versions of Firefox do not save attributes when you click Apply. Workaround: When you have added the new attribute into the script content click the cursor back into the script content, then click Apply.

Error Conditions

The error condition lets you con figure errors for your script.

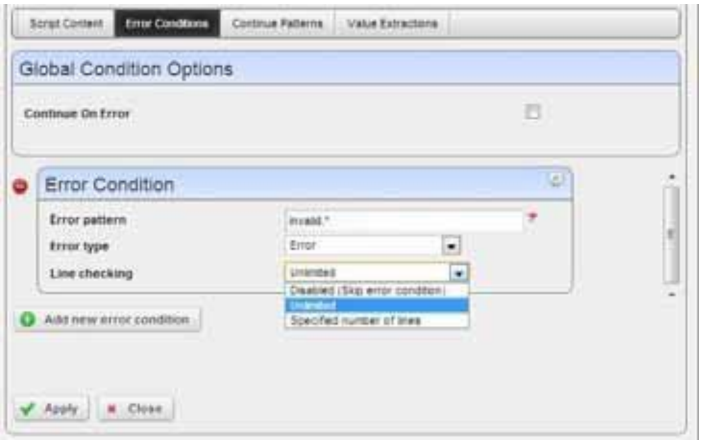

Continue On Error ☐

Error pattern: invalid.*

Error type: Error

Line checking: Unlimited

- Disabled (Skip error condition)
- Unlimited
- Specified number of lines

Check *Continue on Error* under the Global Condition Options, if you want the script to not stop when it encounters an error. Click *Add new error conditions* to configure a condition at the bottom of this screen with the following fields:

Error Pattern—Enter a regular expression for the error. Youcan also click the icon in the upper right corner to test the expression. See Regular Expression Testing on page 586.

Error Type — Select from the pick list of options (*Error, Warning, Ignore*).

Line checking —Select from the pick list (*Unlimited, Disabled (Skip error condition), Specific number of lines*). If you select a specific number of lines, enter the number of lines of the script output to check for the pattern specified, after each command execution. An error message is most likely to appear immediately right after the command is invoked.

Continue Pattern

Like Error Conditions, this screen lets you enter conditions to which script execution can respond.

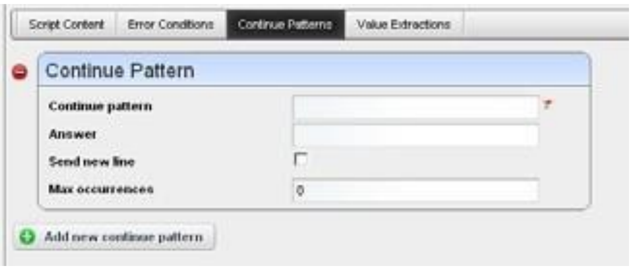

Continue pattern:

Answer:

Send new line ☐

Max occurrences: 0

The Continue Pattern editor operates like the Error Conditions editor but has slightly different fields.

Continue Pattern— If you expect the device output of a script to prompt to continue, you may add a *Continue Pattern* with a regular expression to parse. You can also click the icon in the upper right corner to test the expression. See Regular Expression Testing on page 586.

Answer— This field specifies the *Answer* to the *Continue Pattern* prompt.

Send New Line— For some devices, a single key response without a new line would be sufficient; in such cases, you may need to uncheck the *Send New Line* option.

Max Occurrences— Indicates the maximum number of times respond to a prompt. The default value zero (0) indicates no limit.

Prerequisite Validation

Some devices (JUNOS and Cisco XR) allow you to preview the result of an Adaptive CLI before actuallyexecuting it on devices that support such an advanced preview. Standard previewsimply discloses what Cruz sends to devices; advanced previewrequires a device response indicating the effect.

A *Prerequisite Validation* tab appears for a Show or CLI Configure Command Adaptive CLIs that lets you allow or prevent Adaptive CLI execution based on a regular expression match during such advanced previews.

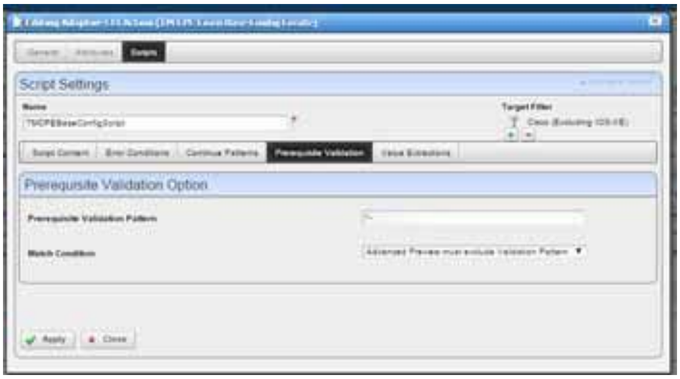

Name: TMCPEBaseConfigScript

Target Filter: Cisco (Excluding IOS-XE)

Prerequisite Validation Pattern: .*

Match Condition: Advanced Preview must exclude Validation Pattern

This tab provides options to select the following.

Prerequisite Validation Pattern—This is a regular expression that is a condition to prevent execution of the Adaptive CLI.

Match Condition—For a supported device, Match Condition lets you select whether execution should be prevented if a match occurs for the Prerequisite Validation Pattern on any line of the Advanced Preview Results obtained from the device. Alternatively, you may prevent execution if a match does not occur for any line of the advanced preview Result.

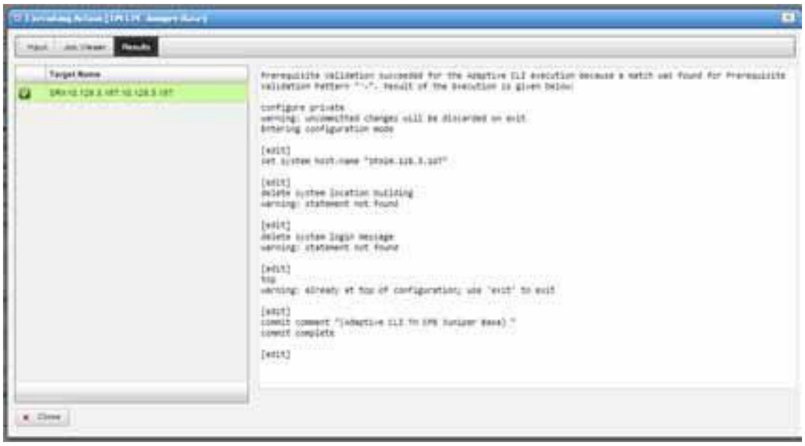

SRX-10.128.3.107 10.128.3.107

```
Prerequisite Validation succeeded for the Adaptive CLI execution because a match was found for Prerequisite
Validation Pattern ".*". Result of the Execution is given below

configure private
warning: uncommitted changes will be discarded on exit
Entering configuration mode

[edit]
set system host-name "SRX10.128.3.107"

[edit]
delete system location building
warning: statement not found

[edit]
delete system login message
warning: statement not found

[edit]
top
warning: already at top of configuration; use 'exit' to exit

[edit]
commit comment "(Adaptive CLI TM SPE Juniper Base)."
commit complete

[edit]
```

Fornon-supported devices, if you provide a Prerequisite Validation Pattern, on execution you are advised to edit the Adaptive CLI and remove the pattern before execution.

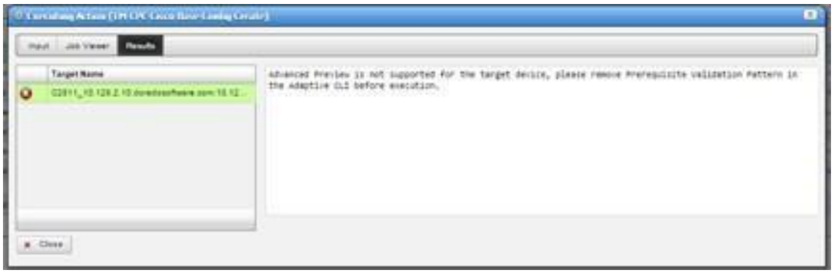

	Target Name
●	C2811_10.128.2.10.doradosoftware.com:10.12...

```
Advanced Preview is not supported for the target device, please remove Prerequisite Validation Pattern in
the Adaptive CLI before execution.
```

Prerequisite Validation does not appear for types other than CLI Show Command and CLI Configure Command.

Value Extraction

Tosupport Adaptive Service and Active Monitor functions, Adaptive CLI providesa wayfor the user to define output schema attributes. This ta b is active only if you have configured schema attributes to store values previously in the Attributes portion of this editor.

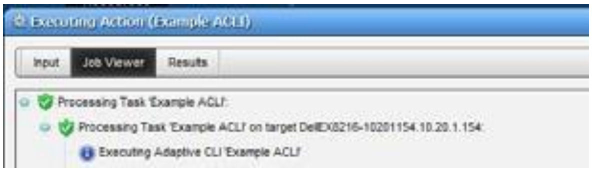

Executing Action (Example ACLI)

- Processing Task 'Example ACLI':
  - Processing Task 'Example ACLI' on target DellEX8216-10201154.10.20.1.154:
    - Executing Adaptive CLI 'Example ACLI'

This lets you *Add*, *Edit* or *Delete* extracted attributes, like Conditions' editor. To clarify configured *Attributes*, *Parse Algorithms*, and *Parse Expressions* accompany scripts, they appear in a table. Use the *Add* button to create more Value Extractions, and the *Edit* or *Delete* buttons to the right of listed patterns to alter or remove them.

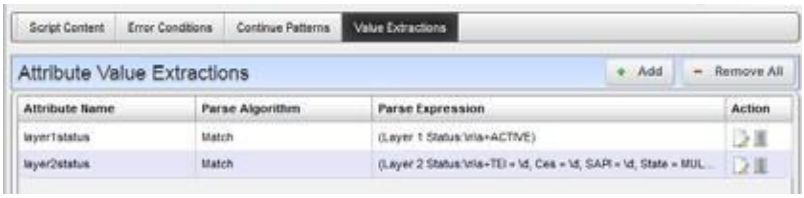

Attribute Value Extractions

Attribute Name	Parse Algorithm	Parse Expression	Action
layer1status	Match	(Layer 1 Status:\s\s+ACTIVE)	
layer2status	Match	(Layer 2 Status:\s\s+TEI = \d, Ces = \d, SAPI = \d, State = MUL...	

Configure Value Extractions with the following fields:

Attribute Name— This field specifies the name of the extracted attribute. To specify the output value of an attribute, select it from the provided list.

Attribute Type—The data type of the attribute extracted. Only schema attributes of simple type String, Integer, Long, Float, Double, and Boolean are available to choose from.

Parse Algorithm— Select from the pick list (*Extract, Match*). For match algorithm, the result is either *true* or *false* for the Boolean attribute type, 0 or 1 for numeric types, or *"true"* or *"false"* for String type.

NOTE:

Currently, Active Performance Monitor supports only numeric types, but you can configure extraction to produce numbers. See Example 5: Monitor Text Values on page 583 for an example.

Parse Expression— Enter a regular expression for Parse Expression and the Parse Algorithm (Extract or Match) used when evaluating the device output on a given script execution. Cruz matches the regular expression for sub-strings, so no need to provide a leading and trailing "match all" regular expression. (.*).

See Regular Expressions on page 495 for more information about what is these expressions can do. You can also click the icon in the upper right corner to test the expression. See Regular Expression Testing on page 586.

Click *Apply* to accept your edits or *Cancel* to abandon them. Click *Add new attribute extraction* to add more such patterns to your script.

Attribute Appearance and Validation

Invalid schema attribute names appear in the script in red italics. This indicates that you cannot use such attributes in the script.

Valid attribute names contain alphanumeric characters and underscore (_). They must begin with either an underscore or a letter [A-Za-z].

All blank space characters in the schema attribute name are converted to underscore (_) by default.

A schema attribute name that is invalid in Adaptive CLI may still be valid in other entities, so you can specify them in the schema, but they are not usable by Adaptive CLI.

Click *Apply* to accept your edits for the script or *Cancel* to abandon them.

Config File Template Generation

This feature lets you create configuration file templates with the Config File Generation Action. Like other Actions, you can associate these with a particular type of device in the initial panel of the editor.

You can use the Adaptive CLI Script Language Syntax within the configuration template you create. Optionally you can add two lines concealed in the generated script at the beginning of the file you create in the Scripts editor that determine some features of the configuration file itself:

```
filename=[myfilename] ; This is a sample comment
label=[mylabel]
```

You can also separate these from comments with a semicolon.

Youcannot enter encrypted passwords in a config fil e template if it contains an attribute delimiter or optional attribute tag. Youcan however create string variables as you would for other scripts and enter special characters when prompted for that variable as you execute the Action.

These specifications appear in the *Preview* and *Result* screens, but do not appear in the final configuration file produced. `filename` determines how to name the generated configuration file, and `label` applies the designated label, or creates a new label and applies it if it does not yet exist.

When you *Execute* this action, it requires you to select a target device, which can either be a database entry generated by right clicking in the Managed Resources portlet (select *New* and a device type), or an actual discovered device.

After generating the configuration, ifyou right-click that selected device, its *Details* panel's *History* tab displays the generated configuration file, and lets you right-click to view, edit or compare it as you would other configuration files . Youcan also right-click the *Execution Details* rows to viewany parameters you have entered in viewer's the *Scripts and Parameters* tab.

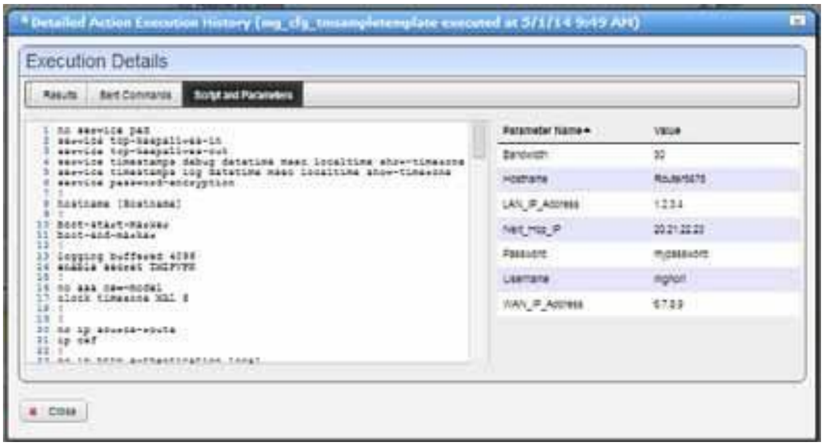
* Detailed Action Execution History (mg_cfg_tmsampletemplate executed at 5/1/14 9:49 AM)
Execution Details
Results
Sent Commands
Script and Parameters
Close

```
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
!
hostname [Hostname]
!
boot-start-marker
boot-end-marker
!
logging buffered 4096
enable secret [illegible]
!
no aaa new-model
clock timezone [illegible]
!
!
no ip source-route
ip cef
!
no ip http authentication local
```

Parameter Name	Value
Bandwidth	30
Hostname	Router5678
LAN_IP_Address	1.2.3.4
Next_Hop_IP	20.21.22.23
Password	mypassword
Username	[illegible]
WAN_IP_Address	6.7.8.9

The How to: Use Config File Generation that follows demonstrates how to use this feature.

How To:
Use Config File Generation

Youcan use Actions/Adaptive CLI to generate and push initial configurations to devices. Here are the steps to do so:

1. Create a new Action/Adaptive CLI by right clicking in the Actions portlet and selecting *New.*
2. In the subsequent editor screen, select Type *Config File Generation.*That action type includes as its script the Configuration File Template
3. Configure any attributes you want to add to the configuration file you generate as described in Attributes on page 548.

4 Add the configuration template to Scripts.

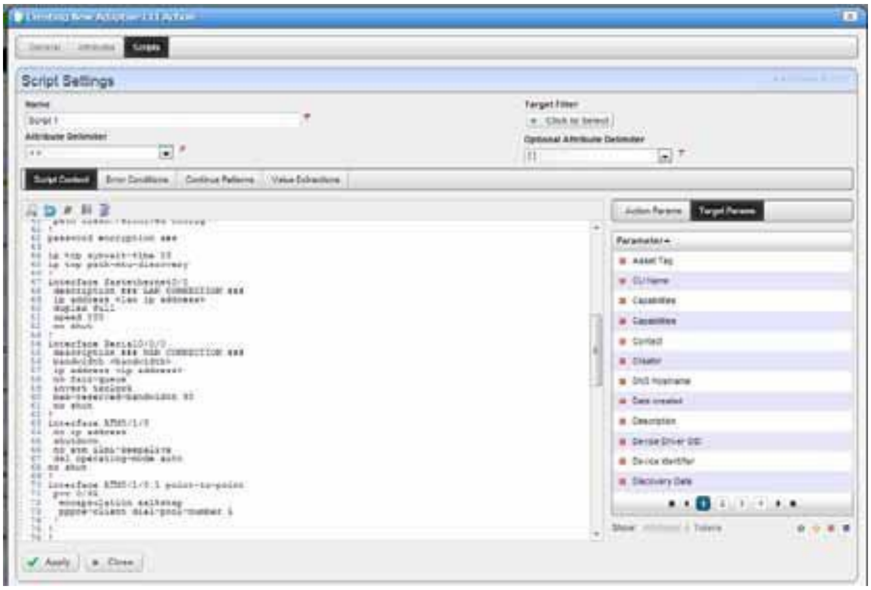

Name: Script 1

Attribute Delimiter: ++

Target Filter: Click to Select

Optional Attribute Delimiter: [[

```
!
password encryption aes
!
ip tcp synwait-time 10
ip tcp path-mtu-discovery
!
interface FastEthernet0/1
 description *** LAN CONNECTION ***
 ip address <lan ip address>
 duplex full
 speed 100
 no shut
!
interface Serial0/0/0
 description *** WAN CONNECTION ***
 bandwidth <bandwidth>
 ip address <ip address>
 no fair-queue
 [illegible]
 max-reserved-bandwidth 95
 no shut
!
interface ATM0/1/0
 no ip address
 shutdown
 no atm ilmi-keepalive
 dsl operating-mode auto
no shut
!
interface ATM0/1/0.1 point-to-point
 pvc 0/35
  encapsulation aal5snap
  pppoe-client dial-pool-number 1
 !
!
!
```

Parameter
Asset Tag
CLI Name
Capabilities
Capabilities
Contact
Creator
DNS Hostname
Date created
Description
Device Driver OID
Device Identifier
Discovery Date

5 Edit it, adding attributes, if necessary, and *Apply* and *Save*.

NOTE:

This editor validates entries. You cannot save the script unless you choose the correct Optional Attribute Delimiter.

To test this, create a new device in Redcell Inventory.

1 In Resources portlet, right-click and select *New*.

2 Select the Vendor and Model, Click *Create*.

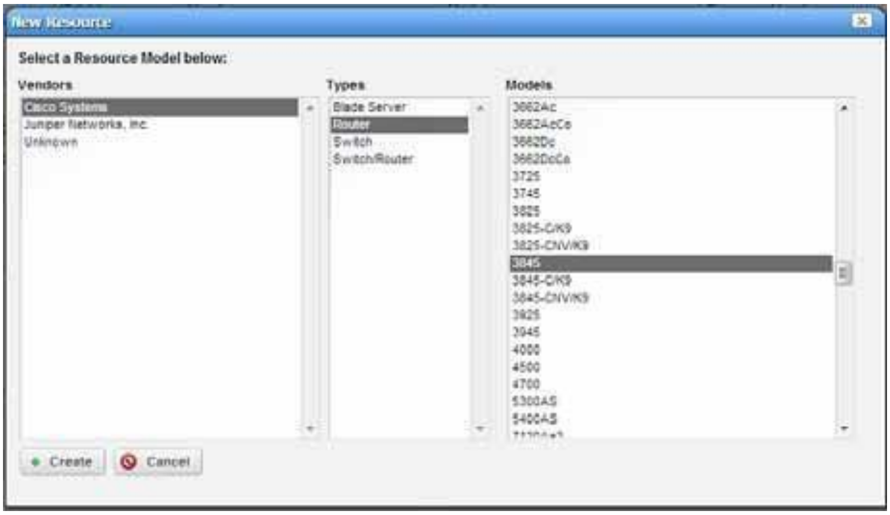

Select a Resource Model below:

Vendors:
- Cisco Systems
- Juniper Networks, Inc
- Unknown

Types:
- Blade Server
- Router
- Switch
- Switch/Router

Models:
- 3662Ac
- 3662AcCo
- 3662Dc
- 3662DcCo
- 3725
- 3745
- 3825
- 3825-C/K9
- 3825-CNV/K9
- 3845
- 3845-C/K9
- 3845-CNV/K9
- 3925
- 3945
- 4000
- 4500
- 4700
- 5300AS
- 5400AS

3 Populate attributes in the resource editor and *Save* the new device.

4 Execute the *Config File Generation* action you created in the previous steps with the newly created device as its target. When executing the action, you must enter the appropriate values.

To do this, right-click the newly created device in the Managed Resources portlet, select Actions, select the appropriate *Generate Config File* action and *Load Selected.*

5 Once you have selected the action, you must enter values in the attribute fields provided.

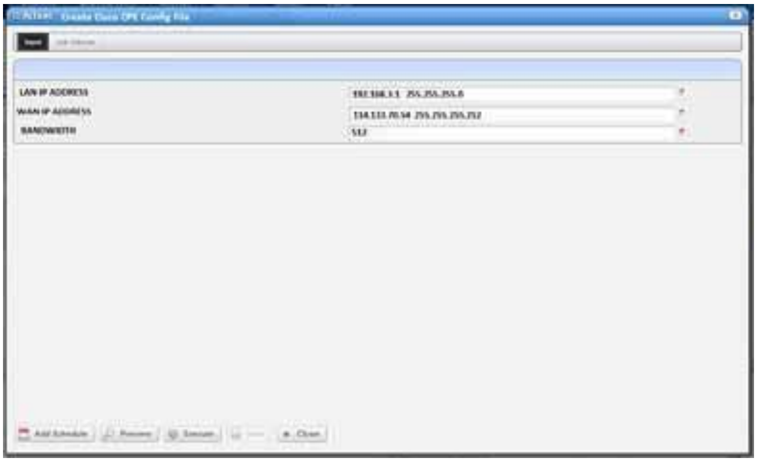

LAN IP ADDRESS 192.168.1.1 255.255.255.0

WAN IP ADDRESS 114.133.70.94 255.255.255.252

BANDWIDTH 512

6 Click on preview to see any changes this produces.

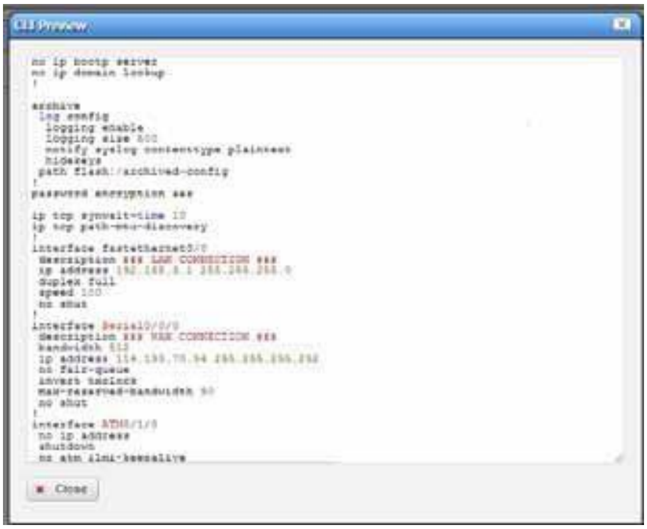

```
no ip bootp server
no ip domain lookup
!

archive
 log config
  logging enable
  logging size 500
  notify syslog contenttype plaintext
  hidekeys
 path flash:/archived-config
!
password encryption aes

ip tcp synwait-time 10
ip tcp path-mtu-discovery
!
interface FastEthernet0/0
 description ### LAN CONNECTION ###
 ip address 192.168.1.1 255.255.255.0
 duplex full
 speed 100
 no shut
!
interface Serial0/0/0
 description ### WAN CONNECTION ###
 bandwidth 512
 ip address 114.133.70.94 255.255.255.252
 no fair-queue
 invert txclock
 max-reserved-bandwidth 90
 no shut
!
interface ATM0/1/0
 no ip address
 shutdown
 no atm ilmi-keepalive
```

7 When you *Execute* this action, the audit trail screen also displays a *Results* panel that shows the configuration.

8 Oncecreated, the configuration fileappears in the device's *Details* panel's Execution History. You can right-click two selected configuration filesgenerated this way to compare them.

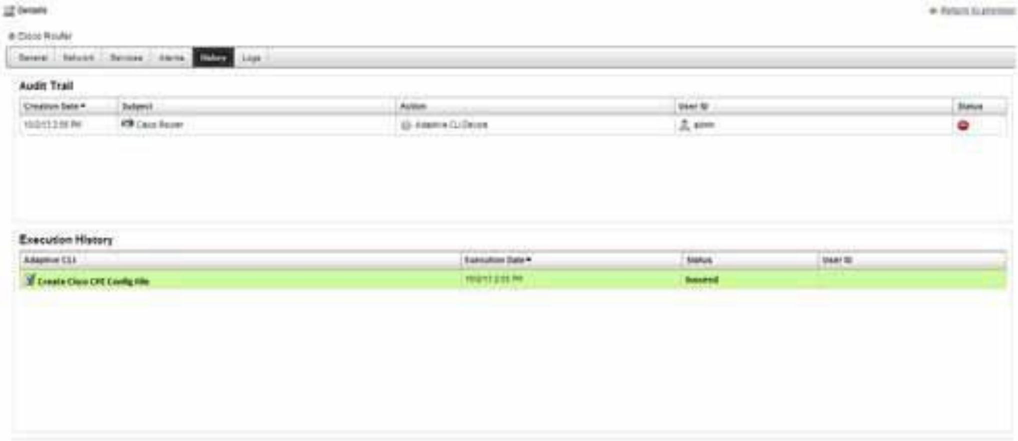

Audit Trail

Creation Date ▴	Subject	Action	User ID	Status
[illegible] PM	Cisco Router	Adaptive CLI Device	admin	●

Execution History

Adaptive CLI	Execution Date ▴	Status	User ID
Create Cisco CPE Config File	[illegible] PM	[illegible]	

9 Right click the Execution History row and select *Execution Details*. The attribute values appear.

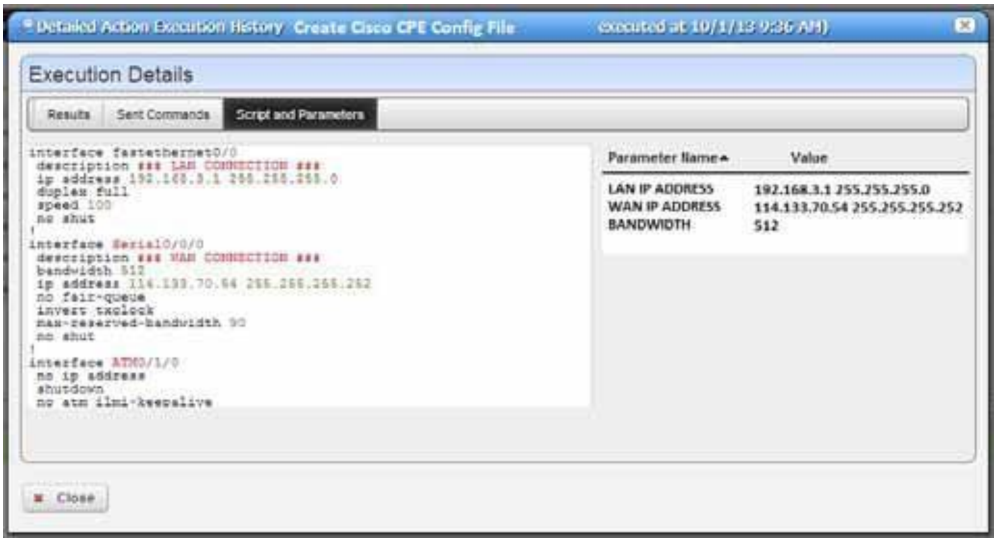
Detailed Action Execution History Create Cisco CPE Config File executed at 10/1/13 9:36 AM)
Execution Details
Results
Sent Commands
Script and Parameters
Close

```
interface fastethernet0/0
 description ### LAN CONNECTION ###
 ip address 192.168.3.1 255.255.255.0
 duplex full
 speed 100
 no shut
!
interface Serial0/0/0
 description ### WAN CONNECTION ###
 bandwidth 512
 ip address 114.133.70.54 255.255.255.252
 no fair-queue
 invert txclock
 max-reserved-bandwidth 90
 no shut
!
interface ATM0/1/0
 no ip address
 shutdown
 no atm ilmi-keepalive
```

Parameter Name ▴	Value
LAN IP ADDRESS	192.168.3.1 255.255.255.0
WAN IP ADDRESS	114.133.70.54 255.255.255.252
BANDWIDTH	512

NOTE:

Anywhere you can see the generated configuration file in the Configuration Files portlet or in a device's Details screen, you can right-click and export it. You can also import files as configurations in the Configuration files portlet. This may be handy to have a baseline to compare to generated files. Just select both files and right-click to compare.

External Executable

External executable Adaptive CLIsessentiallyrun external scripts fromthe Cruz environment. For example, you could run the DOS `dir`command (and schedule its execution). Makesure you select External Command as the Type of Adaptive CLI in the editor when you create an Adaptive CLI that refers to an external command. Also, make sure the `Net::Telnet` package is installed with Perl.

NOTE:

Toruntargetless external commands ina distributed environment, youmust addthecurrent application server's IPaddress toamediation partition's routing entry. Theexternal command runs atoneof mediation servers that belongs to that mediation partition. Any external scripts mustexistonall mediation servers within that mediation partition, and the directory path mustbethesame. Best practice istouseashared networkdrive(orcloud)wheneveryouneedtoaccess files frommultipleservers.

You can execute external commands with a device as target, using device attributes as input parameters to the Adaptive CLI script. See some of the Seeded External Scripts on page 567.

Audit Trail

When you execute a script, the audit screen displays information about it.

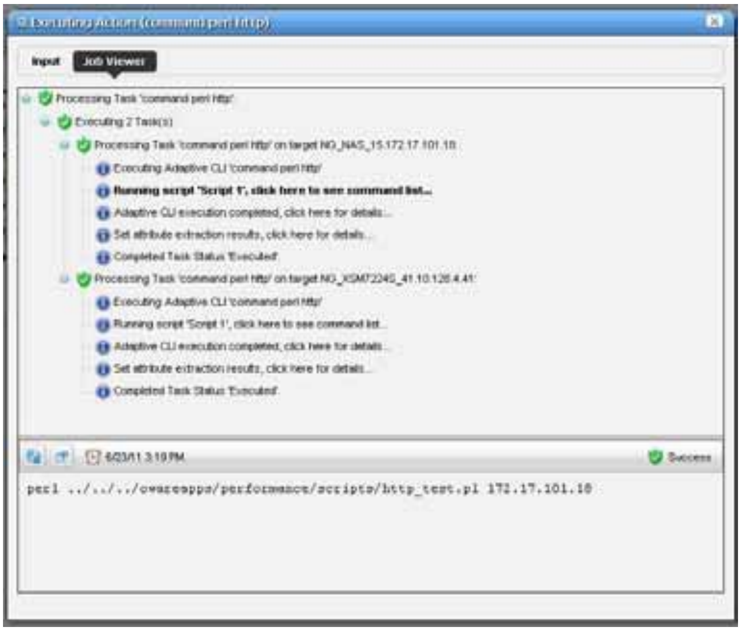

```
Processing Task 'command perl http'
  Executing 2 Task(s)
    Processing Task 'command perl http' on target NG_NAS_15.172.17.101.18
      Executing Adaptive CLI 'command perl http'
      Running script "Script 1", click here to see command list...
      Adaptive CLI execution completed, click here for details...
      Set attribute extraction results, click here for details...
      Completed Task Status 'Executed'
    Processing Task 'command perl http' on target NG_XSM7224S_41.10.128.4.41'
      Executing Adaptive CLI 'command perl http'
      Running script 'Script 1', click here to see command list...
      Adaptive CLI execution completed, click here for details...
      Set attribute extraction results, click here for details...
      Completed Task Status 'Executed'
```

```
perl ../../../owareapps/performance/scripts/http_test.pl 172.17.101.18
```

By default, this screen often conceals the *info* circles in this screen. To see them, click the icon next to the refresh icon to open the message level selector and check the *info* circle level of reporting, then click *Refresh* to see those blue circles.

Results

Cruz stores the results of running a script as lines the Execution Details snap panel. Right click the particular command run in the snap panel at the bottom of the Expanded Actions Portlet. Tabs show the Results, Sent Command, and Script and Parameters. When viewing a script run the results of running it appear target device-by-device.

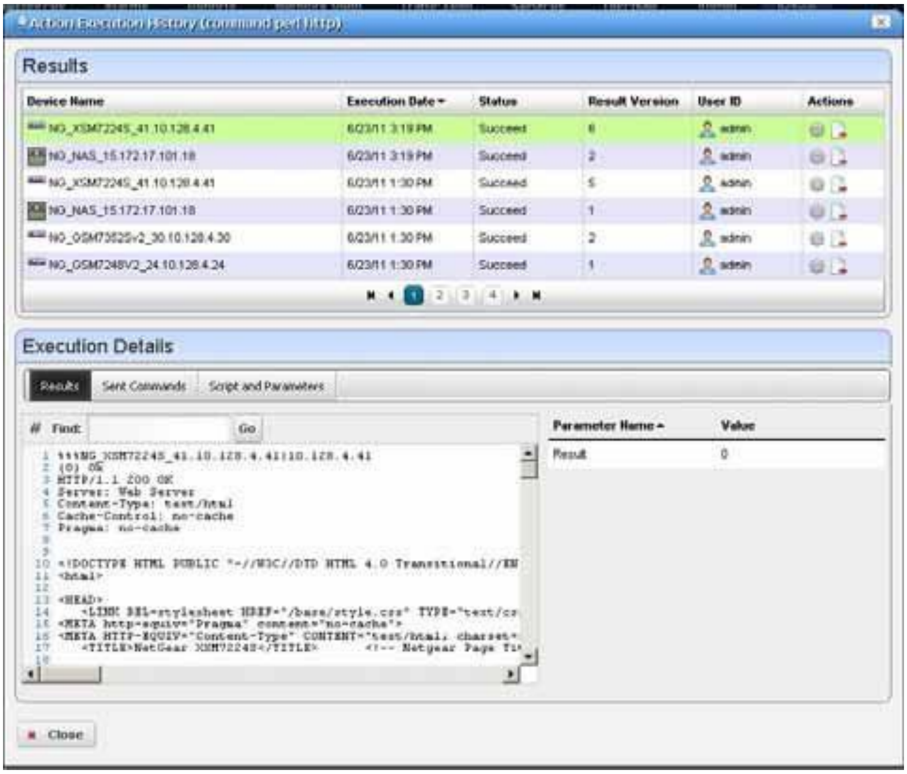

Device Name	Execution Date	Status	Result Version	User ID	Actions
NG_XSM7224S_41.10.128.4.41	6/23/11 3:19 PM	Succeed	6	admin	
NG_NAS_15.172.17.101.18	6/23/11 3:19 PM	Succeed	2	admin	
NG_XSM7224S_41.10.128.4.41	6/23/11 1:30 PM	Succeed	5	admin	
NG_NAS_15.172.17.101.18	6/23/11 1:30 PM	Succeed	1	admin	
NG_GSM7352Sv2_30.10.128.4.30	6/23/11 1:30 PM	Succeed	2	admin	
NG_GSM7248V2_24.10.128.4.24	6/23/11 1:30 PM	Succeed	1	admin	

```
$$$NG_XSM7224S_41.10.128.4.41$10.128.4.41
(0) OK
HTTP/1.1 200 OK
Server: Web Server
Content-Type: text/html
Cache-Control: no-cache
Pragma: no-cache


<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.0 Transitional//EN
<html>

<HEAD>
    <LINK REL=stylesheet HREF="/base/style.css" TYPE="text/cs
<META http-equiv="Pragma" content="no-cache">
<META HTTP-EQUIV="Content-Type" CONTENT="text/html; charset=
    <TITLE>NetGear XSM7224S</TITLE>       <!-- Netgear Page Ti
```

Parameter Name	Value
Result	0

Results can also appear in the audit screen messages and in the *Results* panel of the *Action* job viewer screen.

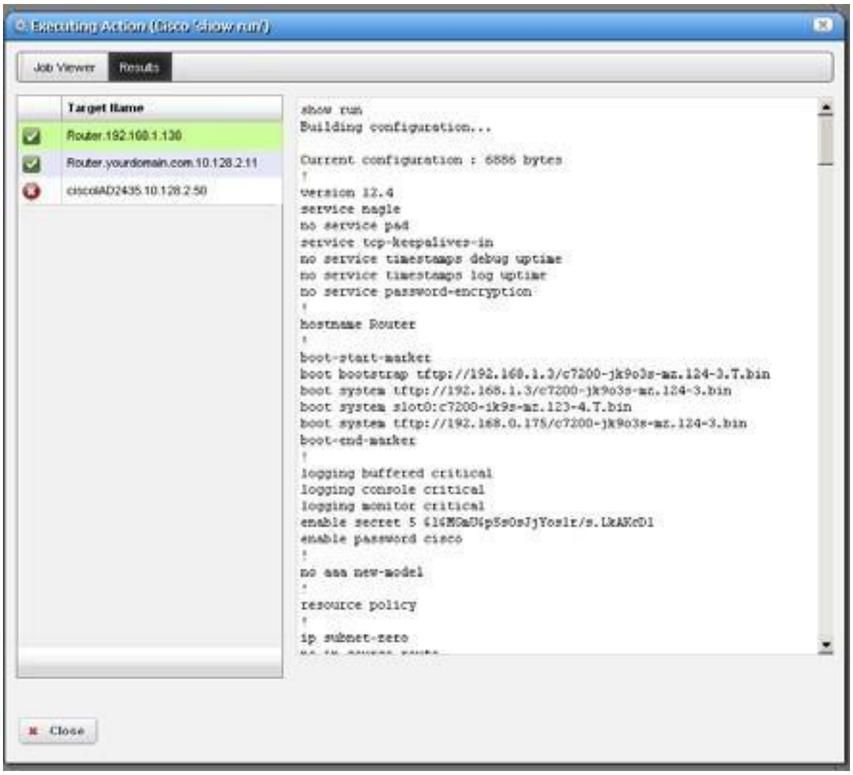

	Target Name
✓	Router.192.168.1.130
✓	Router.yourdomain.com.10.128.2.11
✗	ciscoAD2435.10.128.2.50

```
show run
Building configuration...

Current configuration : 6886 bytes
!
version 12.4
service nagle
no service pad
service tcp-keepalives-in
no service timestamps debug uptime
no service timestamps log uptime
no service password-encryption
!
hostname Router
!
boot-start-marker
boot bootstrap tftp://192.168.1.3/c7200-jk9o3s-mz.124-3.T.bin
boot system tftp://192.168.1.3/c7200-jk9o3s-mz.124-3.bin
boot system slot0:c7200-ik9s-mz.123-4.T.bin
boot system tftp://192.168.0.175/c7200-jk9o3s-mz.124-3.bin
boot-end-marker
!
logging buffered critical
logging console critical
logging monitor critical
enable secret 5 $1$MGaU$pSs0sJjYos1r/s.LkAKcD1
enable password cisco
!
no aaa new-model
!
resource policy
!
ip subnet-zero
```

Youcan also extract parameters for these external commands as is described in Value Extraction on page 559.

Seeded External Scripts

Severalexternal perl scripts come with Cruz as examples of the kind ofcommands you can execute (and Monitor, see Adaptive CLI on page 389 in Performance Monitoring). These are in `\owareapps\performance\scripts` under the installation root.

Torun these, the scripts panel in the Adaptive CLI editor should contain something like the following:

```
perl ../../../owareapps/performance/scripts/http_test.pl
```

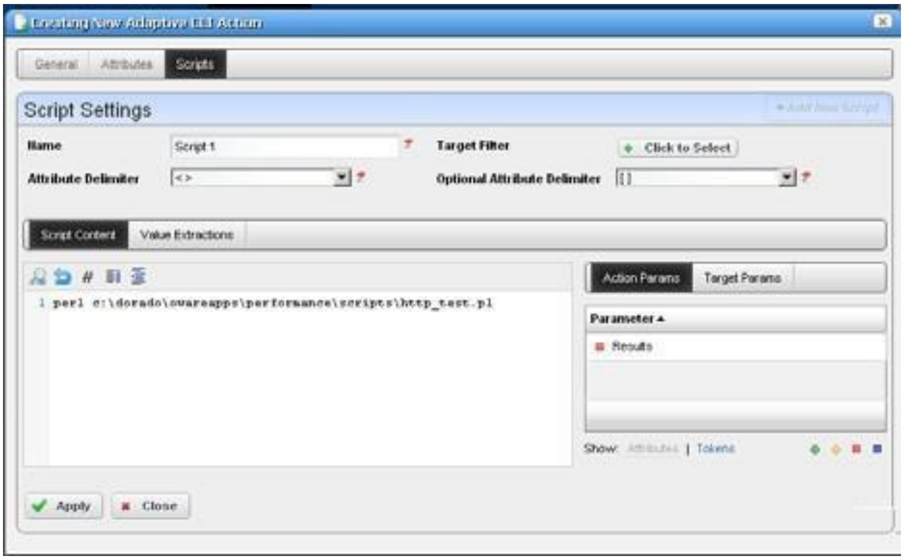

Name: Script 1
Attribute Delimiter: < >
Optional Attribute Delimiter: [ ]

```
perl c:\dorado\owareapps\performance\scripts\http_test.pl
```

Parameter: Results

Notice that these also include a parameter (*Result*) that contains values extracted.

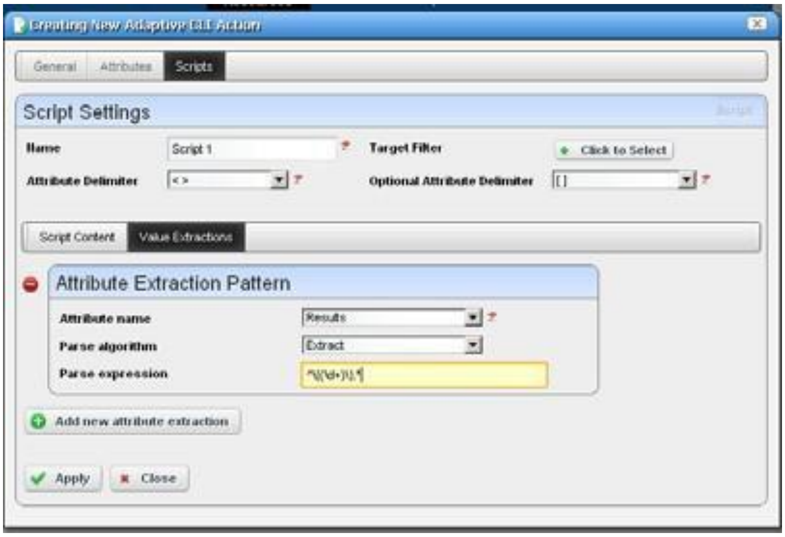

Name: Script 1
Attribute Delimiter: < >
Optional Attribute Delimiter: [ ]

Attribute Extraction Pattern

Attribute name: Results
Parse algorithm: Extract
Parse expression: ^\((\d+)\).*

Set up attribute extraction in the *Values Extraction* tab of the script editor.

Script Names and Functions

Linux installations must have the Net::Telnet package installed with Perl.

common.pl — Common functions defined for scripts in this directory.

dns_test.pl — Check if DNS can resolve the specified host name.

finger_test.pl—Check if the finger service is running on a specified host.

ftp_test.pl — Check the FTP service is running on a specified host.

http_test.pl—Check the HTTP service is running on a specified host.

nntp_test.pl—Check if the NNTP service is running on a specified host. (Public NNTP server to test: news.aioe.org)

peping_test.pl—Check if a target is pingable from the specified remote host.

pop3_test.pl — Check if the POP3 service is running on a specified host.

smtp_test.pl — Check if the SMTP service is running on a specified host.

telnet_test.pl—Check if the TELNET service is running on a specified host.

See Create a Monitor for an External Script on page 384 for more specifics about monitoring these.

If you have a clustered installation, then every server in the cluster must have scripts installed to the paths Adaptive CLIs using them specify.

How To:

Make an Adaptive CLI Run an External Script

The following demonstrates Adaptive CLI running an example Windows external batch file (script) that includes command line parameters. Follow these steps to create this training example:

1. Right click in the *Actions* portlet to create a new Adaptive CLI.
2. In the editor's *General* tab, enter a name (here: Test HelloWorld.bat), and select *External Command* as the type.
3. Click the *Attributes* tab, and create Hello World Schema with two string attributes (in the *Attribute Settings* sub-tab). Here, we make *Command1* and *Command2*.
4. In the *Script* tab, make the Hello World Batch File command (the example name), whose contents are:

   ```
   c:\HelloWorld.bat [Command1] [Command2]
   ```

   Both command line parameters are optional in this example, but you can create such scripts where parameters are required before the script will run. Select a parameter and click the Tokens at the bottom of the screen to arrange that.
5. After you have finished configuring the script, click *Apply* and *Save*.
6. Before running this Adaptive CLI example, you must create a batch file called `c:\HelloWorld.bat`. Here are its contents:

   ```
   @ECHO OFF
   ECHO -----------------------------------
   ECHO * Hello World!   *
   ECHO -----------------------------------
   ECHO ...then display the appended command(s)
   ECHO ...
   ECHO    %1
   ECHO ...
   ECHO    %2
   ```
7. After creating and saving `c:\HelloWorld.bat`, right click Test HelloWorld.bat in the Actions portlet, and select *Execute*.
8. You must select a target device before going further, even though this script does not require one. Select any device.
9. A screen offering to let you specify Command1 and Command2 appears. For the sake of this example, any string you enter works. We'll enter XXX and YYY.

10 Click the *Execute* button.

11 The Job Viewer screen appears displaying the command line you have specified (`c:\HelloWorld.bat XXX YYY`) in informational messages. Youtypicallyhave to configure the Job Viewer so these appear. They are concealed by default.

12 Finally, the *Results* screen appears with the device you specified on the left, and the result of the batch file run on the right.

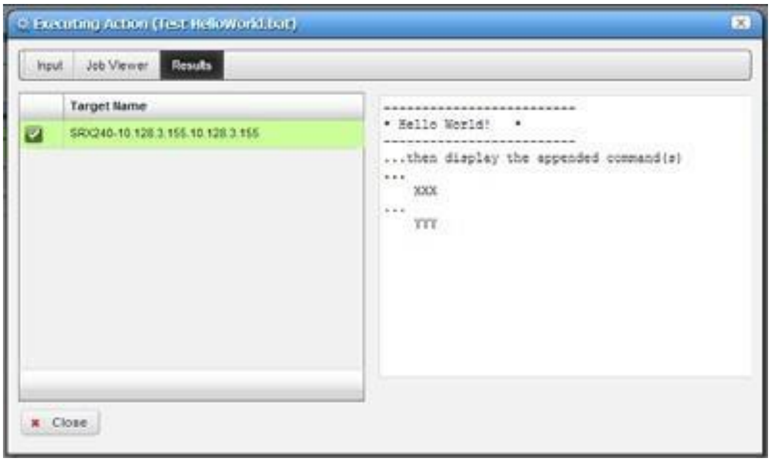

	Target Name
☑	SRX240-10.128.3.155.10.128.3.155

```
-------------------------
* Hello World!   *
-------------------------
...then display the appended command(s)
***
    XXX
***
    YYY
```

For further practice, try running a script of your own, or one of the seeded Perl Scripts (see Seeded External Scripts on page 567).

Adaptive CLI Script Language Syntax

The following is the Adaptive CLI scripting language syntax:

- CLI script is a line-based syntax. In other words, each line’s syntax has to be completed.
- CLI script supports primarily two features: Attributes and Conditional Blocks.

Attributes

Each attribute in the script is marked by a delimiter. The following delimiters are supported:

```
<> [] {} () $ % @ #
```

Think of Attribute delimiters as a pair of open/close markers surrounding a variable name. For single character Attribute delimiters, there is no closing marker (the close marker is empty).

Examples of Attributes are:

```
<var>, [var], {var}, (var), $var, %var, #var, @var
```

The default mandatory delimiters are <>, and the default optional delimiters are `[]`, but you can change those default settings. That means an Attribute variable like <var > may represent a mandatory or an optional Attribute depending on what are set as delimiters.

NOTE:

Single delimiter symbols require a space after the attribute. These do allow values immediately before thesymbol. Perlrequires aspaceaftertheattribute,ortheattribute’sclosingdelimiter,butvalues immediately before single delimiters works.

Here is an example of a command line with a mandatory and optional Attribute:

```
show <mandatory> [optional]
```

If you set the `<mandatory>` Attribute to *interface* and do not set the `[optional]` one, then the resulting command would be this:

```
show interface
```

If you set the `<mandatory>` Attribute to *interface* and set `[optional]` to *brief* then the resulting command would be:

```
show interface brief
```

Conditional Blocks

Every line in the script is presumably a command to be sent to the device, except for lines that denote either a beginning or ending of a conditionalblock.

The begin conditional block marker is tied to an Attribute and has the following syntax:

```
<optional-open-delimiter> IF optional-attribute <optional-close-
  delimiter>
```

The end conditional block marker has the following syntax:

```
<optional-open-delimiter> ENDIF optional-text < optional-close-delimiter>
```

Here is an example ofa conditional block, where the Attribute delimiters are <>, optional delimiter is `[]`, and the conditional Attribute variable is `set`:

```
[IF set]
  execute this command
  and execute this command
[ENDIF set]
```

If the Attribute set has a value then the block is evaluated; otherwise, it is ignored. The text after `ENDIF`., that is `set` or whatever is not required, and it is ignored.

Nested conditional blocks are allowed.

Perl Scripts

This section describes the details of using Perl scripts within Adaptive CLI. See Using Perl in Adaptive CLI on page 540 for more about why to use Perl.

The Perl output goes to the selected target device. Typically, this means creating lines like the following:

```
println(“show $param”);
```

or

```
print(“show $param\n”);
```

You must specify parameters within the script (like `$param`) in the screen described in Attributes on page 548. Unlike its internal scripts, Adaptive CLI does not automatically create attributes. You must also manually configure created attributes to be *Mandatory*, or *Optional* in that screen.

A few things to remember when using Perl:

- The normal output of your Perl scripts (to stdout) are the commands sent to a device by this application.

- If your script produces an error message (to stderr), the job fails with that message and all script outputs are ignored. Youcan validate a script before sending any command to the device by using die(...) and warn(...) functions in Perl to produce error messages to stderr. Such messages trigger the script's failure.
- Forsuch scripts to operate correctly, you must have Perl installed on the directory path for all Cruz servers.
- Perl may not come with Cruz and must be on the server system for it to work with Adaptive CLI.
- Youcaninstall your versionof Perland set the PATHenvironment variableaccordingly so that one can run `perl -v` from the command line (where the Cruz server is to be started). Adaptive CLI invokes that same `perl`command.

  If for some reason Adaptive CL I, fails to invoke the default `perl` command, it reads the setting of `activeconfig.perl.exe=...` inside `owareapps/activeconfig/lib/ac.properties`, and uses that alternative command.

  Note that thedefault `activeconfig.perl.prefix=`settingin `ac.properties` is prepended to every Perl script. It basically forces the script to `use strict` mode and provides a convenient `println` method for the user. Knowledgeable Perl users can change this default behavior setting but should be careful about it. Remember, best practice is to override properties as described in Overriding Properties on page 112.
- The standard output (using `println`) of the Adaptive CLI Perl script represents the command set that is to be sent to the device. For convenience, a `println` subroutine is embedded with the script.
- Adaptive CLI with Perl scripts must contain valid Perl under the "strict" pragma (`use strict;`). If you import or migrate from a previous version a Perl script that does not pass this "strict" criterion, you must rewrite it for "strict" compliance before it can be successfully edited or copied.

NOTE:

WhenyouimportaPerlAdaptiveCLIthatdoesn'tpassstrict,youcanexecuteitwithoutproblems. However, you*cannot*edititatall,unlessyoufirsteditittopassstrict(oritwon'tevenletyousavethe changes).

The following Perl Example may be of interest.

Perl Example

The following is an example Perl script for Adaptive CLI:

```
#
# A script example for testing against a Cisco-XR machine.
#

# The following variables (attributes) are defined in the schema,
# and their values are assigned when the script
# is invoked from the Adaptive CLI (or Resources) manager.
# These variables will be declared with values and prepended
# to each script automatically.  Something like:
#
# my $FromPort=<some number>;
# my $ToPort=<some number>;
# my $Mtu=<some number>;
# my $Desc=<some text>;
#

print("config t\n");

foreach ($FromPort .. $ToPort) {
  my $Desc = "$Desc Port #$_";
  my $addr = 100 + $_;

  print("interface GigabitEthernet0/1/1/1.$_\n");
  print("description $Desc\n");
  print("ipv4 address 10.10.100.$addr 255.255.255.0\n");
  print("ipv4 unreachables disable\n");
  print("mtu $Mtu\n");
}

print("exit\ncommit\nexit\n");
```

How To:

Create Adaptive CLI Examples

The examples that follow may not work for your device. They are often created with a specific device, with specific syntax. Best practice is to telnet to the device you plan to target with your Adaptive CLI and test the command line there first. Then configure any extraction you plan to use based on that testing.

The following describes the basics of creating and using Adaptive CLIs.

Example 1 - Existing Show Run uses an existing, seeded Adaptive CLI to show protocols.

Example 2 - New Adaptive CLI describes making and using a new Adaptive CLI.

Example 3 - Adaptive CLI with Reboot shows you how to make an Adaptive CLI that requires rebooting the target device(s).

Example 4 - Adaptive CLI To Extract Upload/Download Speeds demonstrates Adaptive CLI that extracts information from the target device, then displays the results on a dashboard.

Example 5: Monitor Text Values demonstrates using and Adaptive CLI configured to monitor attributes with strings that indicate their status.

Some devices do not respond to commands unless they are in the correct state. For example, some Dell devices must not be in “Simple” mode to respond to Adaptive CLIs. Take account of this as you create Adaptive CLIs.

Example 1 - Existing Show Run

1 Adaptive CLI Manager haspre-seeded tasksand diagnosticcommands based upon the drivers you have installed. For example: The *Cisco 'show protocols'* command. Right-click and Select *Edit* to view and/or alter this Adaptive CLI.

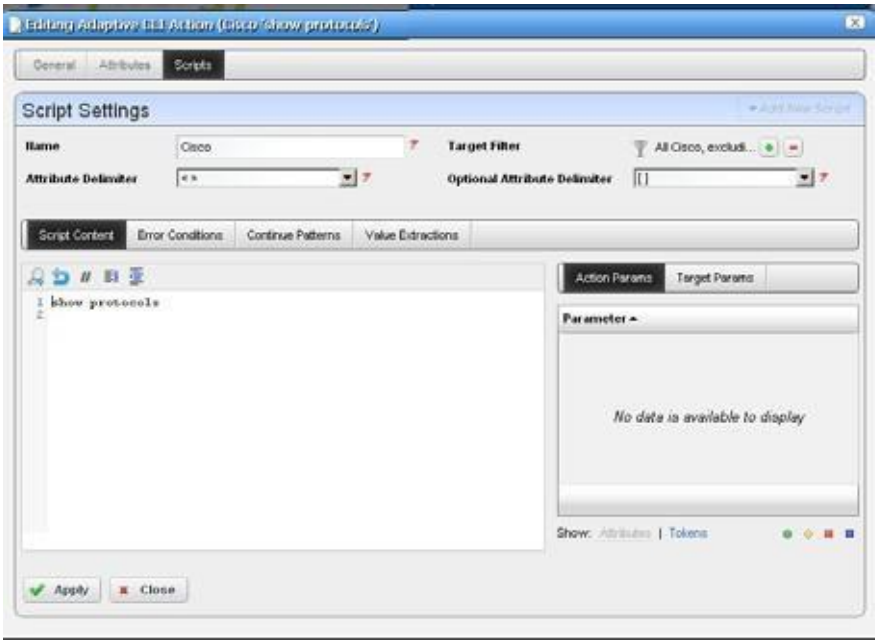

Editing Adaptive CLI Action (Cisco 'show protocols')

Name: Cisco

Target Filter: All Cisco, excludi...

Attribute Delimiter: < >

Optional Attribute Delimiter: [ ]

```
show protocols
```

No data is available to display

2 Click the *Edit* icon next to the Cisco script. The *Scripts* tab in this editor appears above, displaying the `show protocols` command to be sent target devices. Notice (in the upper right corner) that this Adaptive CLI filters so it applies to all Cisco devices excluding PIX.

3 Close the editor(s) and select this Adaptive CLI.

4 Right click to *Execute and* select the target equipment for this run in the next screen. The screen that appears is a standard Cruz equipment selector. The Adaptive CLI is valid only on devices that pass the Target Filter mentioned in step 2, but the selection here narrows the target devices for the Adaptive CLI.

5 An Audit trail screen tracks the execution progress

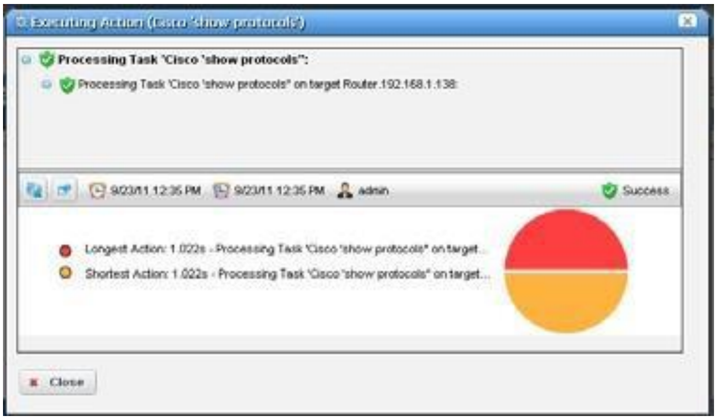

Executing Action (Cisco 'show protocols')

Processing Task 'Cisco 'show protocols'':

Processing Task 'Cisco 'show protocols'' on target Router.192.168.1.138:

9/23/11 12:35 PM 9/23/11 12:35 PM admin Success

Longest Action: 1.022s - Processing Task 'Cisco 'show protocols'' on target...

Shortest Action: 1.022s - Processing Task 'Cisco 'show protocols'' on target...

6 Select the Adaptive CLI you ran in the Expanded Portal, and right-click the execution run that appears in the *Execution History* snap panel at the bottom of thescreen.

7 Right-click and select *Execution Details*.

8 View latest results classified by the device you select on the left.

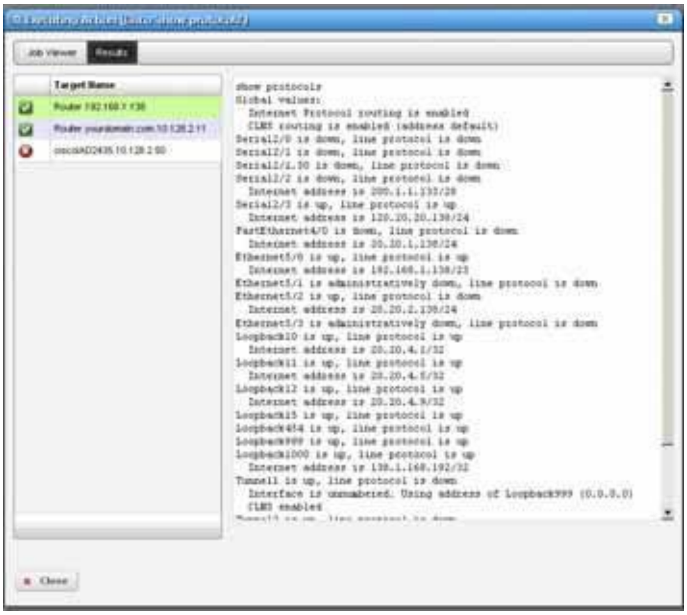

```
show protocols
Global values:
  Internet Protocol routing is enabled
  CLNS routing is enabled (address default)
Serial2/0 is down, line protocol is down
Serial2/1 is down, line protocol is down
Serial2/1.30 is down, line protocol is down
Serial2/2 is down, line protocol is down
  Internet address is 200.1.1.133/28
Serial2/3 is up, line protocol is up
  Internet address is 120.20.20.130/24
FastEthernet4/0 is down, line protocol is down
  Internet address is 20.20.1.130/24
Ethernet5/0 is up, line protocol is up
  Internet address is 192.168.1.130/23
Ethernet5/1 is administratively down, line protocol is down
Ethernet5/2 is up, line protocol is down
  Internet address is 20.20.2.130/24
Ethernet5/3 is administratively down, line protocol is down
Loopback10 is up, line protocol is up
  Internet address is 20.20.4.1/32
Loopback11 is up, line protocol is up
  Internet address is 20.20.4.5/32
Loopback12 is up, line protocol is up
  Internet address is 20.20.4.9/32
Loopback15 is up, line protocol is up
Loopback454 is up, line protocol is up
Loopback999 is up, line protocol is up
Loopback1000 is up, line protocol is up
  Internet address is 130.1.168.192/32
Tunnel1 is up, line protocol is down
  Interface is unnumbered. Using address of Loopback999 (0.0.0.0)
  CLNS enabled
```

9 Viewlatest results byright clicking in the *Execution History* snap-in of the expanded Action portlet. You can use the *Find* search box to find matches to strings within the results.

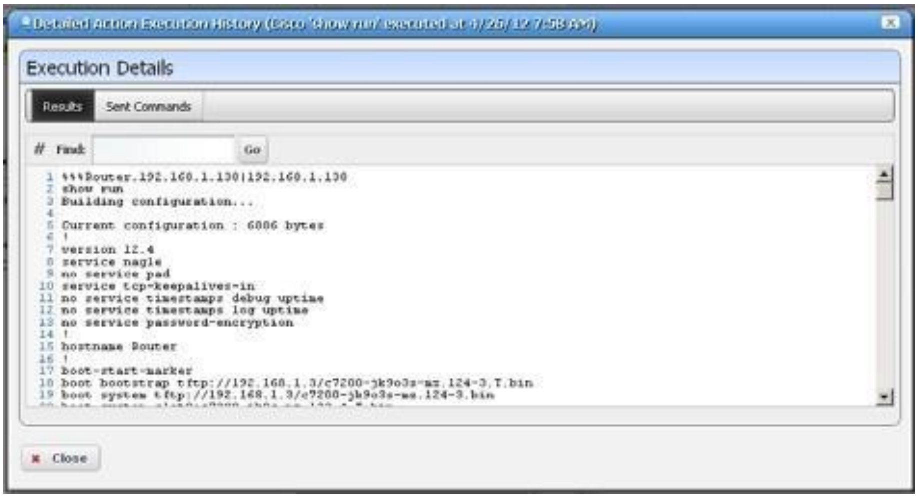

Detailed Action Execution History (Cisco 'show run' executed at 4/26/12 7:58 AM)
Execution Details
Results
Sent Commands
# Find:
Go
Close

```
$$$Router.192.168.1.130|192.168.1.130
show run
Building configuration...

Current configuration : 6806 bytes
!
version 12.4
service nagle
no service pad
service tcp-keepalives-in
no service timestamps debug uptime
no service timestamps log uptime
no service password-encryption
!
hostname Router
!
boot-start-marker
boot bootstrap tftp://192.168.1.3/c7200-jk9o3s-mz.124-3.T.bin
boot system tftp://192.168.1.3/c7200-jk9o3s-mz.124-3.bin
```

Click *Go* to see the next match.

10 You can also look in the *Sent Commands* tab to see what went to the device.

Example 2 - New Adaptive CLI

1 Create a new Adaptive CLI. Right-click and select *New.*

2 Name this (for example "TestACLI")

3 In the *Attributes* panel, create string attributes named *required* and *optional* after creating a new Parameter Schema (for example “test123”).

4 In the *Script* panel define the Attribute Delimiter (< >) and Optional Attributes Delimiter ([ ]) and enter the following three scripts:

```
show run
show <required>
show [optional]
```

Notice that the created attributes appear in the panel on the right of this screen.

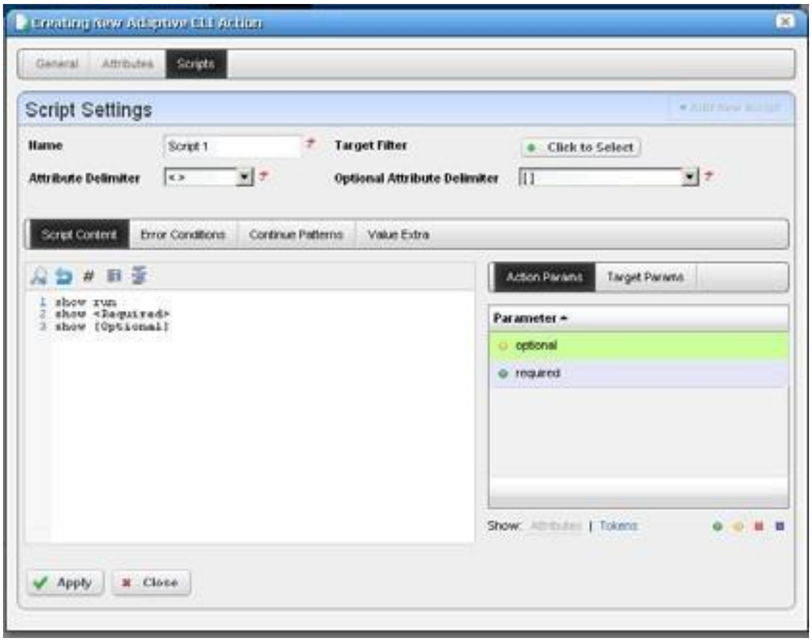

Name: Script 1

Target Filter: Click to Select

Attribute Delimiter: < >

Optional Attribute Delimiter: [ ]

```
show run
show <Required>
show [Optional]
```

optional

required

5 Select the attribute “required,” then click the *Required* icon (the green circle) in the lower right corner to of this screen to associate this icon with the Required attribute. Similarly, associate the *Optional* icon with the attribute “optional.”

Notice that you can double-click the attributes listed in the panel on the right, and they appear in the script editor at the cursor.

6 Save this Adaptive CLI

7 Execute it with *action* > *Execute.*

8 Notice that the attributes entered now are visible as inputs.

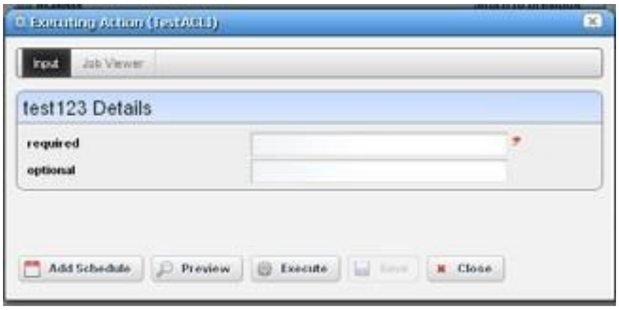

required:

optional:

When you enter values for these, they accompany the `show run` sent to the target devices. Notice that you *must* enter the required variable, or execution fails.

9 Select a target.

10 Click *Execute*. The `show run`, and any other required/optional run commands' results appear. These are searchable with the results screen.

Example 3 - Adaptive CLI with Reboot

The following describes how to set up multi-line ACLI with error/success tracking for a command sequence that requires reboot.

1 Create an example configure Adaptive CLI command (here *quickThenReboot*).

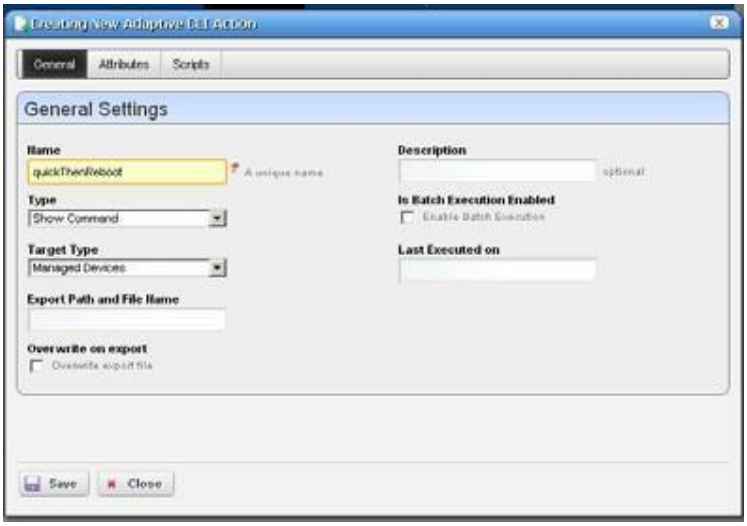

Creating New Adaptive CLI Action

General Settings

Name: quickThenReboot

Type: Show Command

Target Type: Managed Devices

Export Path and File Name:

Overwrite on export: □ Overwrite export file

Description:

Is Batch Execution Enabled: □ Enable Batch Execution

Last Executed on:

2 Separate commands into parts. First issue the command (here `show run`), then issue the reboot command with a parameter that allows a prompt return before actual reboot (a delay, for instance). If the first command fails the ACLI doesn't continue, so that makes using the reboot command second the solution.

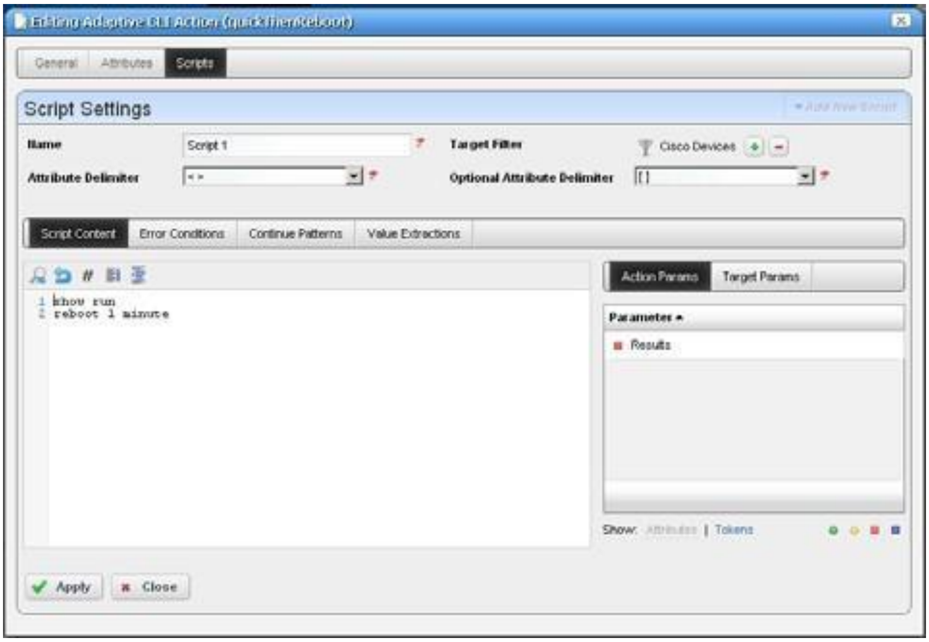

Name: Script 1

Target Filter: Cisco Devices

Attribute Delimiter: < >

Optional Attribute Delimiter: [ ]

```
show run
reboot 1 minute
```

In our example:

```
show run
reboot 1 minute
```

3 Cruz assumes commands are successful if a prompt appears without an error return. Default error tracking for most drivers provides all the error pattern matching you might need (testing the Adaptive CLI lets you know whether the device is addressed by a driver in “most”).

Use specific error pattern matching for cases where the driver does not detect the typical errors by default. As described in the Cisco Adaptive CLI Caveat on page 580, erroneous output appears if the error occurs on the reboot command.

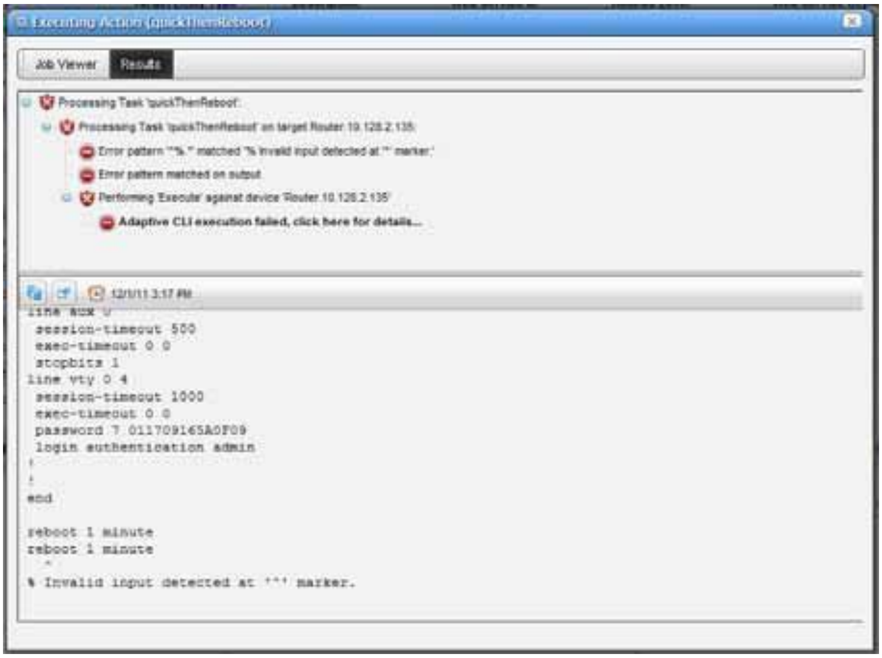

Processing Task 'quickThenReboot'
Processing Task 'quickThenReboot' on target Router 10.128.2.135
Error pattern "^%." matched "% Invalid input detected at '^' marker."
Error pattern matched on output
Performing 'Execute' against device 'Router 10.128.2.135'
Adaptive CLI execution failed, click here for details...

12/1/11 3:17 PM

```
line aux 0
 session-timeout 500
 exec-timeout 0 0
 stopbits 1
line vty 0 4
 session-timeout 1000
 exec-timeout 0 0
 password 7 011709165A0F09
 login authentication admin
!
!
end

reboot 1 minute
reboot 1 minute
  ^
% Invalid input detected at '^' marker.
```

4 When reboot is successful with a proper command sequence, the job screen displays the successful execution. See Cisco Adaptive CLI Caveat on page 580 for more about continuation prompts.

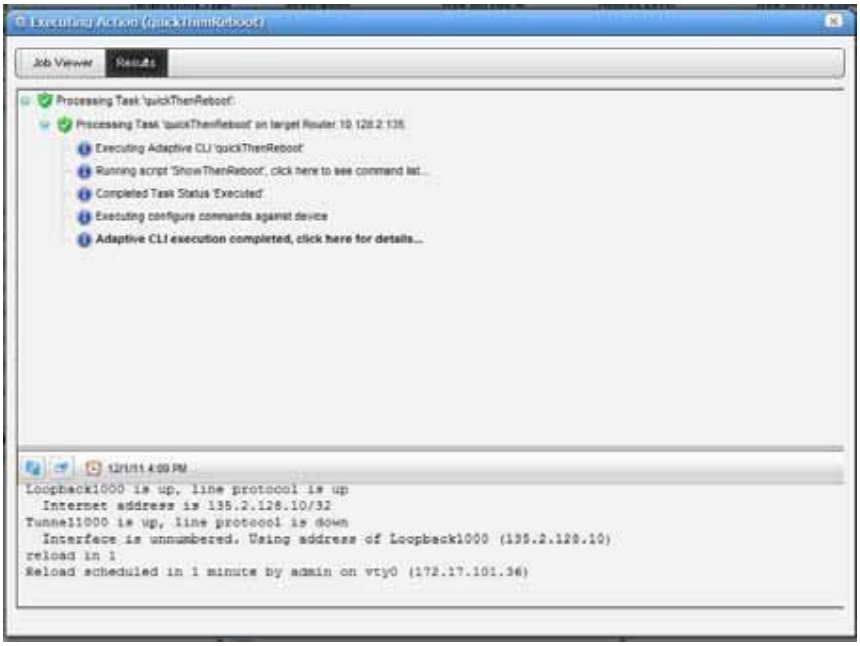

Processing Task 'quickThenReboot'
Processing Task 'quickThenReboot' on target Router 10.128.2.135
Executing Adaptive CLI 'quickThenReboot'
Running script 'ShowThenReboot', click here to see command list...
Completed Task Status 'Executed'
Executing configure commands against device
Adaptive CLI execution completed, click here for details...

12/1/11 4:00 PM

```
Loopback1000 is up, line protocol is up
  Internet address is 135.2.128.10/32
Tunnel1000 is up, line protocol is down
  Interface is unnumbered. Using address of Loopback1000 (135.2.128.10)
reload in 1
Reload scheduled in 1 minute by admin on vty0 (172.17.101.36)
```

5 **Continue Patterns** —The following Continue Patterns section is an addition to the above example. It looks for the `Proceed` prompt so the Adaptive CLI can issue a new line to force the reboot. But the shutdown command follows the next prompt, so the shutdown command must be in another continue pattern to force the last line before a pause in output to be the router's prompt. The patterns are `.*Proceed.*` and `.*SHUTDOWN in.*` allowing any characters before and after the keywords to match.

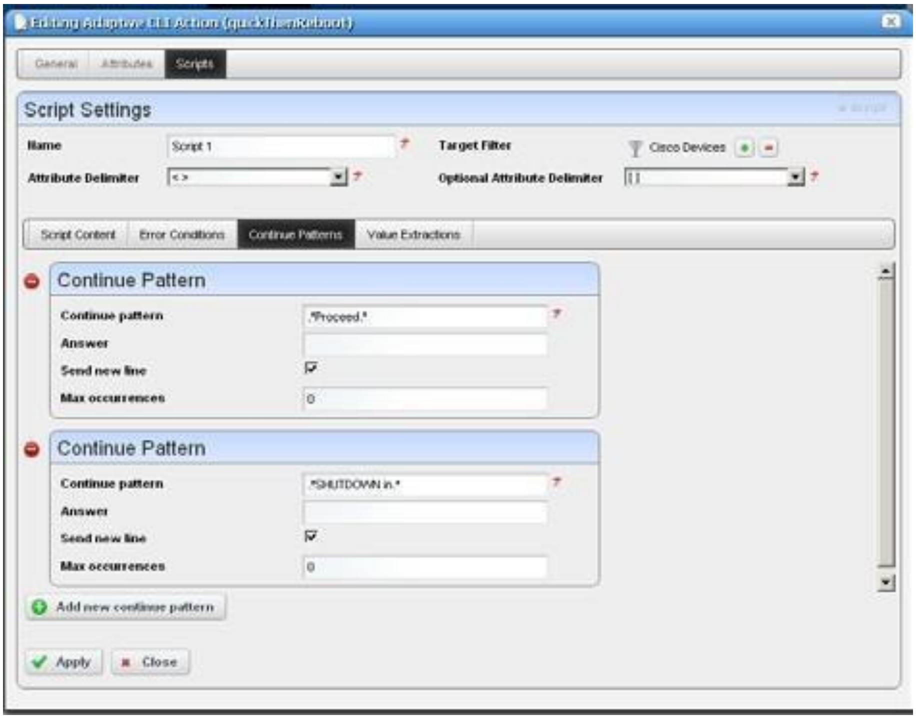

Name: Script 1
Target Filter: Cisco Devices
Attribute Delimiter: <>
Optional Attribute Delimiter: [ ]

Continue Pattern
Continue pattern: .*Proceed.*
Answer:
Send new line: ☑
Max occurrences: 0

Continue Pattern
Continue pattern: .*SHUTDOWN in.*
Answer:
Send new line: ☑
Max occurrences: 0

Alternatively, this example could have a third command after `reboot` to force a new router prompt but managing this problem with the continuation set seemed more straightforward.

Cisco Adaptive CLI Caveat

Ifyou have not saved configuration changes for Cisco routers, then this Adaptive CLI fails. Forsake of this example, such a failure is the correct response. However, if the Adaptive CLI needs multiple confirmations you can just add more with their responses as appears in the Continue Patterns described above.

Example 4 - Adaptive CLI To Extract Upload/Download Speeds

Thefollowing describesan example Adaptive CLI configured to extract upload and download ADSL speeds from a Cisco Router. To create this example, follow these steps:

1 Right-click to create a new Adaptive CLI in the Actions portlet.

2 Name it and configure the Adaptive CLI in the General screen. Since these are generic settings described elsewhere, the details do not appearhere.

3 Create attributes to extract. In this case, we configure Upload Speed, and Download Speed as integer attributes, with a name, description, and nothingelse.

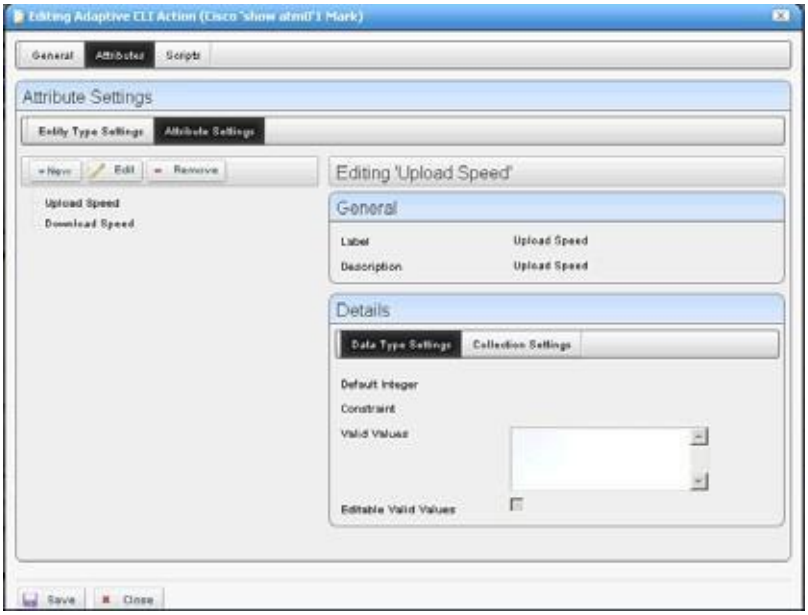

Label Upload Speed

Description Upload Speed

Default Integer

Constraint

Valid Values

Editable Valid Values □

Notice, however, that you could configure validation for extracted attributes ifyou liked in this screen.

4 Create a new schema for these attributes. Schemas are helpful if you are creating several Adaptive CLIs (create, destroy, update, and so on) with the same set of attributes. With schemas, you are sure the attributes are configured the same.

5 *Save* the configured attributes, click the *Script*panel

6 Enter the script. This extracts upload and download speeds from a Cisco device based on the output from this command (the script's contents):

```
show dsl int atm0 | inc Speed
```

Thiscommand shows dsl, grepping (`inc`) for the unique line beginning with `Speed`. The line for which this script searches looks like this:

```
Speed (kbps): 544 0 256 0
```

The attributes configured previously appear beside the script panel, but are not part of the script, even though that possibility might be useful for another Adaptive CLI. The current attributes are for extraction from the script results.

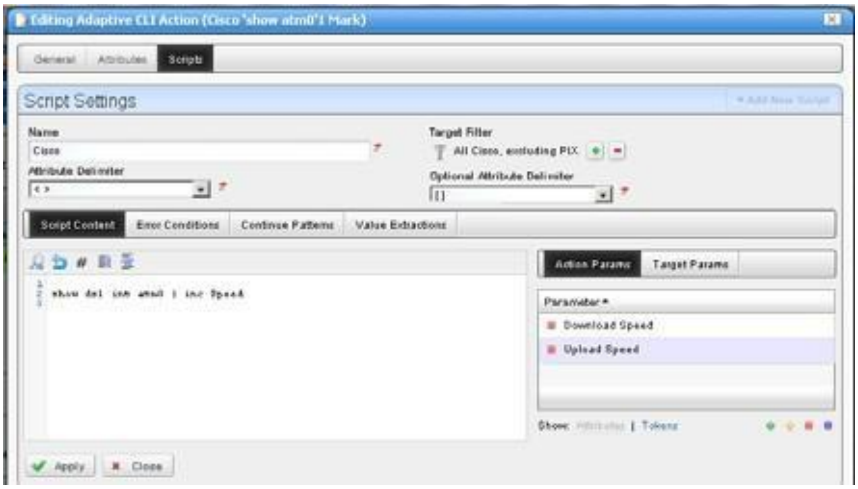

Name: Cisco

Attribute Delimiter: < >

Target Filter: All Cisco, excluding PIX

Optional Attribute Delimiter: [ ]

```
show dsl int atm0 | inc Speed
```

Parameter *
- Download Speed
- Upload Speed

NOTE:

The filter at the top of this panel can limit the devices scanned by the Adaptive CLI to extract data. If you haveaspecificdeviceorgroupofdevices againstwhichyouplantotestthis script, itwouldbeatime saver to create the filter first.

7 Click the *Value Extractions* panel within the Scripts screen and configure an extraction regular expression for each of the two values.

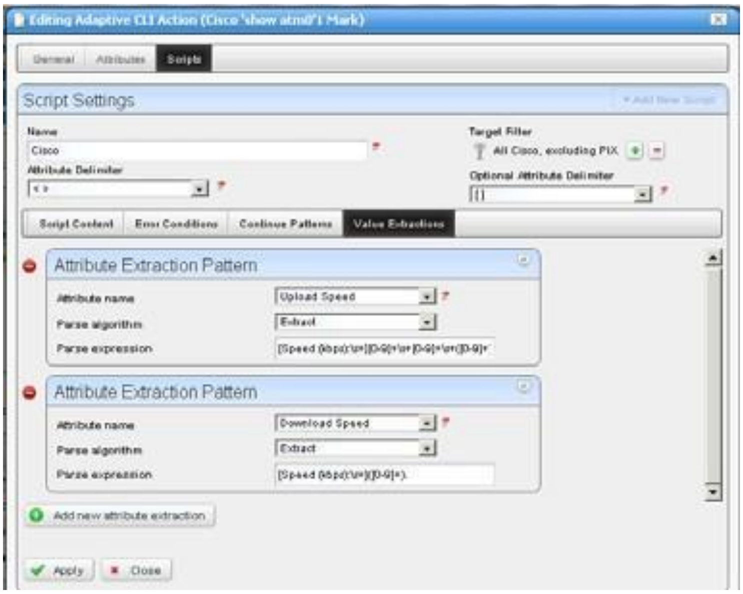

Name: Cisco

Attribute Delimiter: < >

Target Filter: All Cisco, excluding PIX

Optional Attribute Delimiter: [ ]

Attribute Extraction Pattern

Attribute name: Upload Speed

Parse algorithm: Extract

Parse expression: [Speed (kbps):\s*]([0-9]*\s*[0-9]*\s*([0-9]*

Attribute Extraction Pattern

Attribute name: Download Speed

Parse algorithm: Extract

Parse expression: [Speed (kbps):\s*]([0-9]*).

Click the green plus to add the second attribute.

With the pick lists, select an attribute, and that you want to extract (that is, within which you plan to store a value), then enter the regular expression to match its target value. Here are those attribute/regular expression pairs:

- Download Speed (the first integer in the output)
  `[Speed (kbps):\s+]([0-9]+).`
- Upload Speed (the third integer in the output)
  `[Speed (kbps):\s+][0-9]+\s+[0-9]+\s+([0-9]+).`

NOTE:

You can use free regular expression testers to debug these expressions. See Regular Expression Testing on page 586.

8. Apply the edits you have made to script and extractive regular expressions, then *Save* the Adaptive CLI.
9. Right-click the Adaptive CLI and *Execute* it.
10. Select the target device(s).
11. Confirm the execution. The screen that appears before you click *Execute* again would have fields if you had a script with input parameters.
12. The *Results* panel appears to advise whether the script ran successfully, displaying its output.
13. Click *Job Viewer* and arrange that panel so it displays informational messages by clicking the icon next to the date/time display. Check the checkbox next to the blue informational circle and click the *Refresh* icon to the far left.

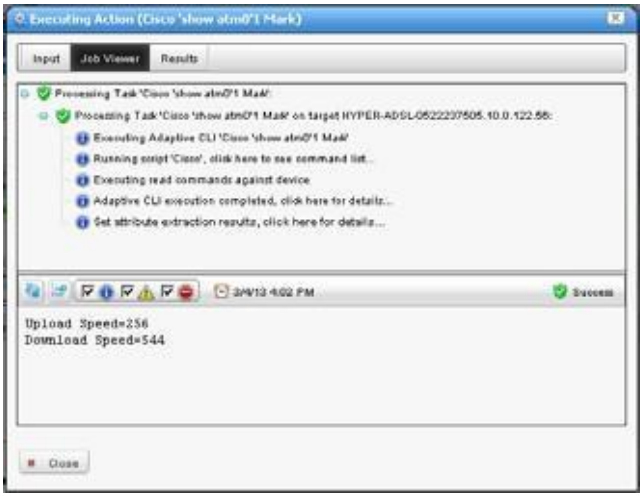

Processing Task 'Cisco 'show atm0'1 Mark'
Processing Task 'Cisco 'show atm0'1 Mark' on target HYPER-ADSL-0522237505.10.0.122.58:
Executing Adaptive CLI 'Cisco 'show atm0'1 Mark'
Running script 'Cisco', click here to see command list...
Executing read commands against device
Adaptive CLI execution completed, click here for details...
Set attribute extraction results, click here for details...

```
Upload Speed=256
Download Speed=544
```

14. Click the last informational message (*Set attribute extraction results...*) and the extracted attribute values appear in the data panel at the bottom of the screen.

Example 5: Monitor Text Values

Create an Adaptive CLI with the following to monitor layer 1 and layer 2 status:

- integer attributes: layer1status, layer2status
- Script to produce the output: `show isdn status`
  Here is the output to match:

```
Layer 1 Status:
 ACTIVE
Layer 2 Status:
 TEI = 0, Ces = 1, SAPI = 0, State = MULTIPLE_FRAME_ESTABLISHED
```

- Attribute Extraction Pattern:

```
layer1status/Match/(Layer 1 Status:\n\s+ACTIVE)
```

- For layer2status, the regular expression is like

```
(Layer 2 Status:\n\s+TEI = \d, Ces = \d, SAPI = \d, State =
 MULTIPLE_FRAME_ESTABLISHED)
```

Create a monitor to display the result of regularly running this Adaptive CLI on selected targets and display its result in a dashboard.

NOTE:

Don’t forget to enable the attributes in the monitor!

Monitoring Upload/Download Speeds

Once you have configured this Adaptive CLI, you can monitor its operation. Follow these steps to configure the monitor for the How to: Create a Monitor for the External Script Adaptive ACLI:

NOTE:

If you are testing, make the monitoring more frequent than you might in a production system so you can see if the data is available as expected. You can always change this after you have successfully tested the monitor.

1 Right-click in the Resource Monitors portlet to create a newmonitor.

2 Enter the default name, and interval for the monitor in the *General*panel.

3 In *Monitor Options*, select the Monitor Entities (target devices) with the green plus, and subsequent screen.

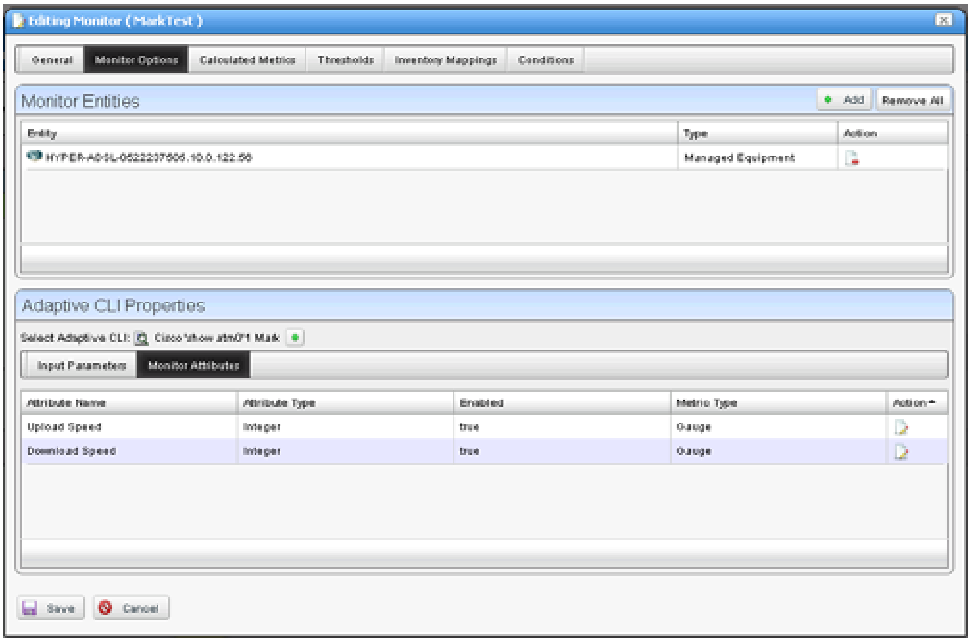

Entity	Type	Action
HYPER-ADSL-0522237505.10.0.122.56	Managed Equipment	

Select Adaptive CLI: Cisco 'show atm0/1 Mark

Attribute Name	Attribute Type	Enabled	Metric Type	Action
Upload Speed	Integer	true	Gauge	
Download Speed	Integer	true	Gauge	

4 In the same screen, elect to *Enable* the extracted Monitor Attributes with the editor icon to the right of the listed attribute. Notice you can also elect to report the attribute as a Gauge, Counter or Boolean. We selected Gauge.

5 Click *Save*.

6 Right-click the saved monitor to *View Monitor Data*.

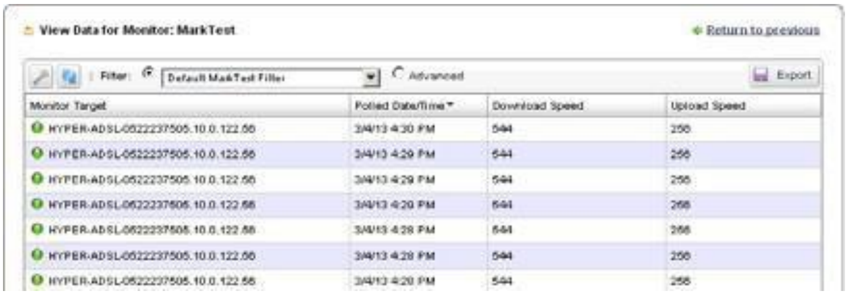

View Data for Monitor: MarkTest

Monitor Target	Polled Date/Time	Download Speed	Upload Speed
HYPER-ADSL-0622237505.10.0.122.56	3/4/13 4:30 PM	544	256
HYPER-ADSL-0622237505.10.0.122.56	3/4/13 4:29 PM	544	256
HYPER-ADSL-0622237505.10.0.122.56	3/4/13 4:29 PM	544	256
HYPER-ADSL-0622237505.10.0.122.56	3/4/13 4:29 PM	544	256
HYPER-ADSL-0622237505.10.0.122.56	3/4/13 4:28 PM	544	256
HYPER-ADSL-0622237505.10.0.122.56	3/4/13 4:28 PM	544	256
HYPER-ADSL-0622237505.10.0.122.56	3/4/13 4:28 PM	544	256

You may have to click the wrench icon to con figure the columns that appear, so this screen displays the extracted attribute information. Youshould see the extracted values displayed in a table.

Configure a Dashboard for Your Monitor

Finally, if you want to configure a dashboard to display your monitored data graphically, follow these steps:

1 Go to the Dashboard portlet, and right-click to create a Custom dashboard.

2 Enter the default data (name, retention policy, and so on) and configure the device and monitor selection by editing the panel(s) you want to display with its editor icon in the upper right corner.

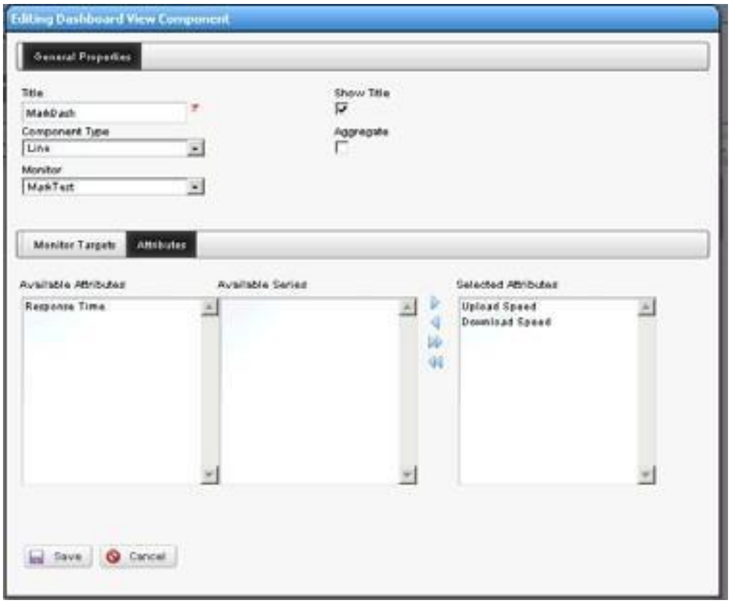

Editing Dashboard View Component

Title: MarkDash

Show Title: ☑

Component Type: Line

Aggregate: □

Monitor: MarkTest

Available Attributes: Response Time

Available Series:

Selected Attributes: Upload Speed, Download Speed

Notice that you can select not only the monitor, but also the target(s) and attribute(s) to display. Here, we haveselectedthe Upload/Download Speedattributesconfiguredinthe How to: Create a Monitor for the External Script Adaptive ACLI.

3 *Save* the configured dashboard.

4 Right-click the dashboard in the Dashboard Views portlet and view it in one of the options available (Full Screen/Popup).

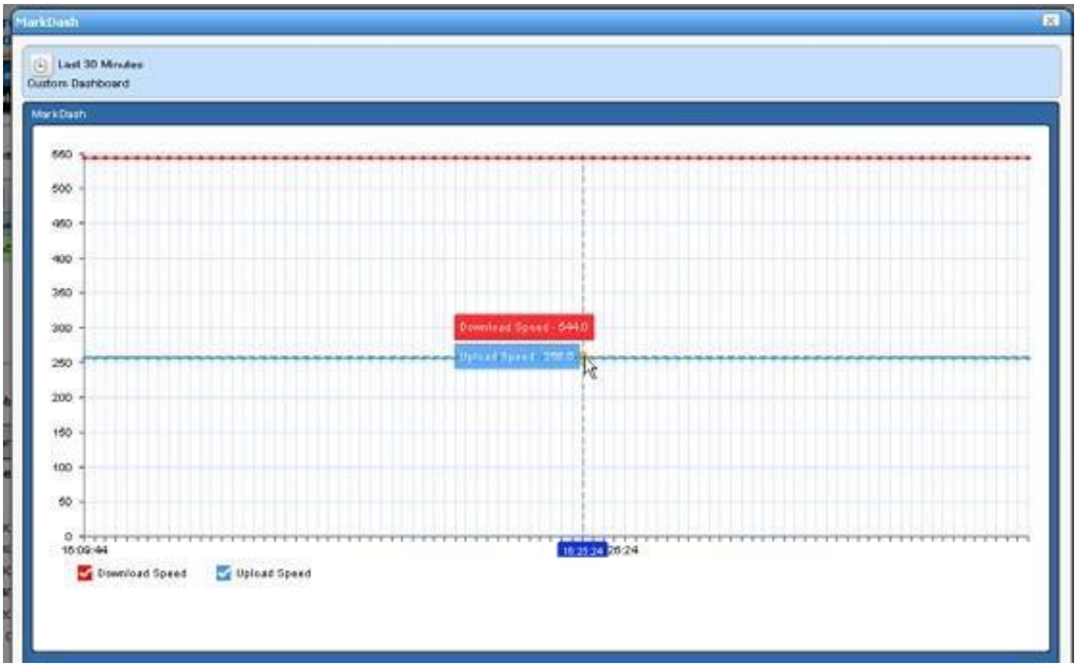

Notice that you can hover your cursor over a node in the graph and see all reported values for that node.

Regular Expression Testing

You can test regular expressions in Cruz by clicking the icon in the upper right corner of screens like the one for Attribute Value Extractions in the Adaptive CLI editor.

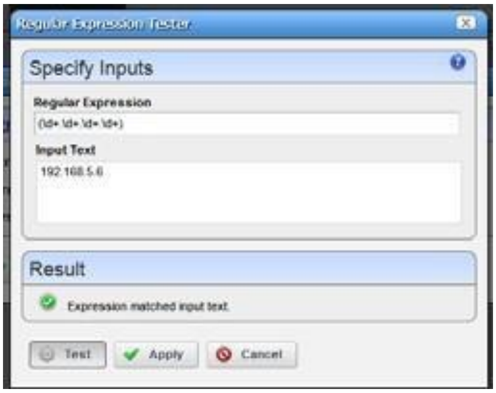

Regular Expression: (\d+.\d+.\d+.\d+)

Input Text: 192.168.5.6

Expression matched input text.

After entering the *Regular Expression* and the *Input Text*, click the *Test* button to see whether the expression extracts the text you want. Revise it if it does not match or click *Apply* to accept the expression and enter it in the editor.

Several applications, some free, are helpful to validate regular expressions too. These include websites (regexr.com, forone). These mayalsobe helpful when trying to matcha number and phrase in the Adaptive CLI and other output. In this example, we used Kodos to test various iterations of our regular expressions.

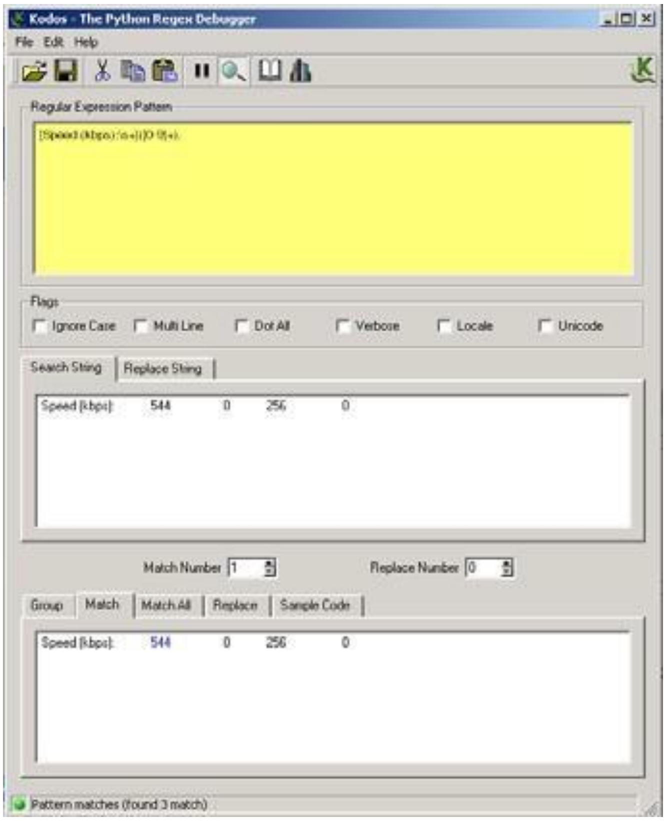

Regular Expression Pattern: [(Speed (kbps):\s+)([0-9]+).

Flags: □ Ignore Case □ Multi Line □ Dot All □ Verbose □ Locale □ Unicode

Search String:

```
Speed (kbps):      544      0      256      0
```

Match Number: 1 Replace Number: 0

Match:

```
Speed (kbps):      544      0      256      0
```

Pattern matches (found 3 match)

Enter the regular expression in the top panel (note the helpful hints from the online help), the output to scan in the middle panel, and the match appears in the bottom panel. Note: this application is not supported by Cruz developers.

Regular expressions include metacharacters to in struct the program how to treat characters it encounters. These include the following: ^, $, . , | , { , } , [,] , (,), *, +, ? , \. If you want to match one of these metacharacters, you must prepend a backslash (\). So to match a literal question mark, rather than instructing regular expression matching to match 0 or 1 of a previous expression, you must enter \?. To match a continue prompt that says `Proceed? (y/n)` you must escape three characters in regular match expression, like this `Proceed\? \(y/n\)`

Scheduling Actions

Youcan schedule actions with a right-click in the Actions portlet or the Schedules portlet. This opens an editor with the following screens:

- General
- Parameters
- Schedule

See Schedules Portlet on page 155 for more scheduling actions with that portlet. Schedules created in the Actions Portlet also appear in the Schedules portlet.

General

This screen lets you identify the scheduled item and its targets.

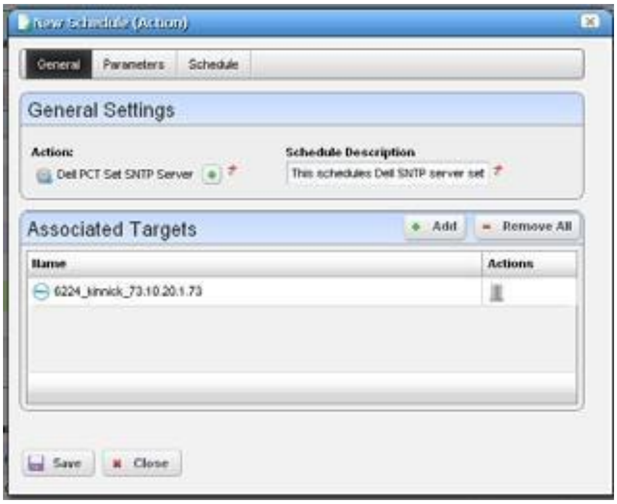

General Settings

Action: Dell PCT Set SNTP Server

Schedule Description: This schedules Dell SNTP server set

Associated Targets

Name	Actions
6224_kinnick_73.10.20.1.73	

This has the following fields:

General Settings

Action — Identifies the action being scheduled.

Schedule Description — Identifies the schedule.

Associated Targets

Click the *Add* button to select target equipment. You can remove listed equipment with the icon to the right of listed items or with the *Remove All* button.

Parameters

Thisscreen'sconfiguration depends on the selected action you arescheduling. Manyactions have no parameters, so this tab is disabled. Enter the parameters for the action you are scheduling.

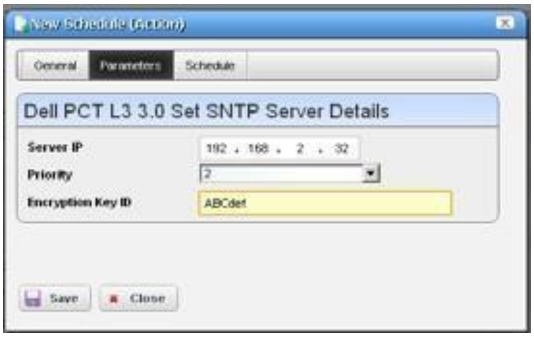

Dell PCT L3 3.0 Set SNTP Server Details

Server IP: 192 . 168 . 2 . 32

Priority: 2

Encryption Key ID: ABCdef

Hover the cursor over fields to make their description appear in a tooltip.

Schedule

This screen is a standard scheduler screen, as described in Schedules Portlet on page 155.

Comparison

Selecting (ctrl +clicking) two Adaptive CLI runs within the *Execution History* portlet lets you compare the two execution results. Right-click and select *Compare*.

Lines that differ between the two configurations appear highlighted green. Lines that are missing in one, but that appear in another appear highlighted red. Added lines appear highlighted in yellow. Use the right/left arrows at the bottom of this screen to page through the side-by-side comparison.

Active Performance Monitor Support

You can monitor Adaptive CLI execution results with Active Performance Monitor. To do this, you must select Adaptive CLI as the monitored type when creating a new performance monitor (see Resource Monitors on page 361), then select a target entities (with the *Add* button in the top panel) and a particular Adaptive CLI (with the green plus [+] in the Adaptive CLI Properties panel at the bottom of this screen. Click the *Edit* (page) icon to select the *Input Parameters* to monitor once you have selected an Adaptive CLI.

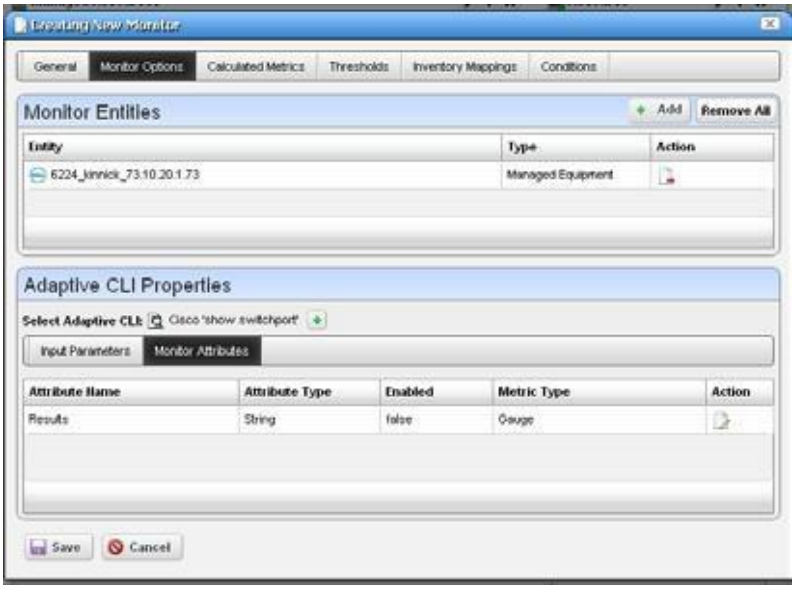

Monitor Entities

Entity	Type	Action
6224_kinnick_73.10.20.1.73	Managed Equipment	

Adaptive CLI Properties

Select Adaptive CLI: Cisco 'show switchport'

Attribute Name	Attribute Type	Enabled	Metric Type	Action
Results	String	false	Gauge	

The user can choose an Adaptive CLI to monitor and may have to configure both its input values and metric type for each output attribute. The Input data depends on what the Adaptive CLI attributes have configured.

Input Parameters

In Active Monitoring, all attributes of the schema appear in the *Input Data* for user-entered values. You must enter the data necessary for all selected targets' scripts. To enter data, click *Edit* and then enter values. Clicking *Apply* switches the panel back to read-only mode. You must click Save to preserve input or output data configurations.

Monitor Attributes

Configure Adaptive CLI output attributes for monitoring in this tab in the lower panel of the Monitor Editor screen. You can monitor only exposed attributes of numeric or boolean types. To change metric type, select the row and click the *Edit* button to its right.

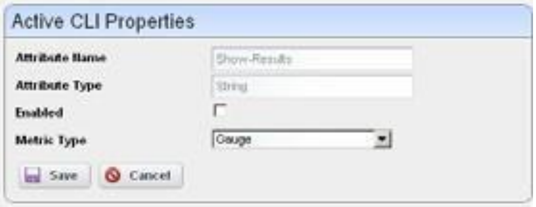

Active CLI Properties

Attribute Name: Show-Results

Attribute Type: String

Enabled: □

Metric Type: Gauge

An Adaptive CLI Properties screen appears that reminds you of the *Attribute Name*, and *Attribute Type*, where you can *Enable* the attribute monitoring, and select *Gauge*, *Counter* or *Boolean* buttons to the right of this panel to configure the metric type of the selected output data.

These attributes default to the metric type *Gauge*. Adaptive CLI is where you define these attributes, but you must select their metric type settings on this screen if it is something other than the default.

Click *Save* to preserve your configuration or *Cancel* to abandon it and close the editor screen.

How To:

Create a Monitor for an External Script

The following steps describe creating a monitor for an external command configured as an Adaptive CLI (ACLI). Several Perl scripts appear in this performance\scripts directory by default. You can try others in addition to the `http_test.pl` script in the example.

Create the Adaptive CLI

1 Right click in the Actions portlet, and create a new *External Command* ACLI

2 Make a new attribute schema with attribute: Status (integer)

3 In Scripts, enter the following as Script Content:

```
perl "[installation path]\owareapps\performance\scripts\
  http_test.pl"[_EquipmentManager_IP_Address]
```

The variable [_EquipmentManager_IP_Address] provides the target device's IP address, and comes from the *TargetParams* tab, where you can find other such variables. If you want to test this script on an HTTP process on a device not under management, just to see the outcome, enter a known URL instead of that variable (like www.testsoftware.com), and run the script to see its output. (You will still have to select a target managed object to run the script, even though it is not part of the command line.)

Since this is an example, use your [installation path] instead of those words.

4 In the Value Extraction panel enter the following:

```
^\{(\d+)\}.*
```

5 Click Apply

6 Click Save

7 Right click and *Execute* the ACLI to test it.

How To:

Create an Advanced Script Monitor Example

The following monitors an external Adaptive CLI example of setting up a simple process monitor using ACLI:

1 Make sure Perl is installed (and Windows has restarted after installing it), and check that the required libraries (Info.pm and WMI.pm) are in place. Your directory may vary; with 64-bit Strawberry Perl the locations are:

For Info.pm:

```
C:\strawberry\perl\vendor\lib\Win32\Process
```

and for WMI.pm:

```
C:\strawberry\perl\vendor\lib\Win32\Process\Info
```

The process folder is attached to this document with proper structure. Put it in C:\strawberry\perl\vendor\lib\Win32 and you are ready to go.

NOTE:

Here are the URLs where you can download these libraries:
http://search.cpan.org/~wyant/Win32-Process-Info-1.018/lib/Win32/Process/Info.pm
http://search.cpan.org/~wyant/Win32-Process-Info-1.019/lib/Win32/Process/Info/WMI.pm

2 Put process_check.pl in the proper directory. For Windows the default is

```
[installation root]\owareapps\performance\scripts.
```

3 In your actions portlet, import TEST_ACTION.xml.

4 In your monitors portlet, import PROCESS_UPTIME_MONITOR .xml.

5 Even though the monitor and Adaptive CLI do not technically need one, select any target a dashboard can track. This permits execution of the Adaptive CLI.

6 In your dashboard views portlet, create a new custom Monitor Dashboard for whatever device(s) you decided to monitor, you will see Status as one of the tracked metrics (1 for up, 0 for down). You can use it as you would any other metric in Cruz to track, graph, and so on.

By default, this script and monitor track whether notepad.exe is running, but you can have it track anything byediting the monitor. Go to Monitor Options > Adaptive CLI Properties, and you can edit the *Process Name* variable to be any other process.

Extra credit: Modify the script to track multiple applications.

process_check.pl

```
#!/usr/bin/env perl
use Win32::Process::Info;
$processname=$ARGV[0];
$found = 0;
$pi = Win32::Process::Info->new ();
@info =$pi->GetProcInfo ();   # Get the max
@info = grep {
  print $_->{Name};
  print "\n";
  if ($_->{Name})
```

```
    {
        if ($_->{Name} eq $processname)
        {
        $found = 1;
        }
    }
} $pi->GetProcInfo ();
if ($found == 1)
{
 print "Process " . $processname . " is running! 1";
}
else
{
 print "Process " . $processname . " is not running! 0";
}
```

TEST_ACTION

This action's name is *TestExternalScript* . It has two attributes, *Process Name*, a string, and *Status*, an integer. It stores the retrieved process' status in the *Status* integer and takes *Process Name* as a required input. It refers to the process_check.pl script as an external command in its *Scripts* tab. Here is the syntax:

```
perl [installation root]\owareapps\performance\scripts\process_check.pl
 <Process_Name>
```

In addition to referring to the script, this Adaptive CLI extracts the status from the script's run. Essentially it looks for 0 (down) or 1 (up) with the following regular expression in the *Value Extractions* tab:

```
(\d)$
```

Note: `[installation root]` is an example only, not a legitimate part of the path.

PROCESS_UPTIME_MONITOR

This monitor's name is *ProcessUptimeMonitor*. It refers to the TestExternalScript (TEST_ACTION) Adaptive CLI. Notice that the *Process Name* attribute defaults to notepad.exe, and the Monitor Attributes tab contains the *Status* attribute.

Monitor Dashboard

To see the result of your monitoring, create a custom monitor dashboard with the PROCESS_UPTIME_MONITOR as its target monitor, and the desired target device as its target device.

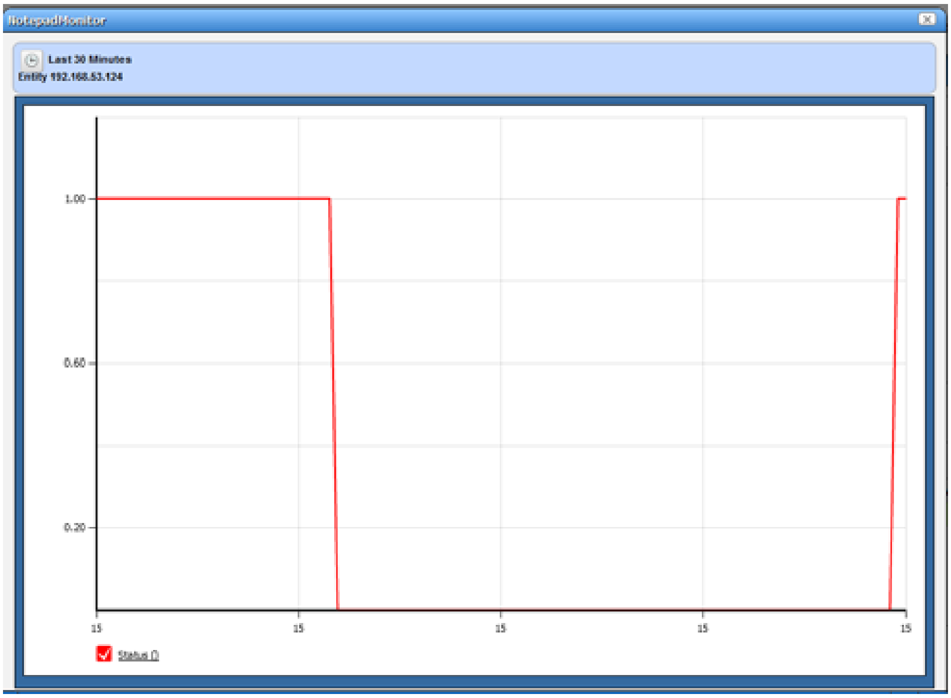

You can then see the process' activity over time when you launch the dashboard.

7 Look in Job Viewer for the results.

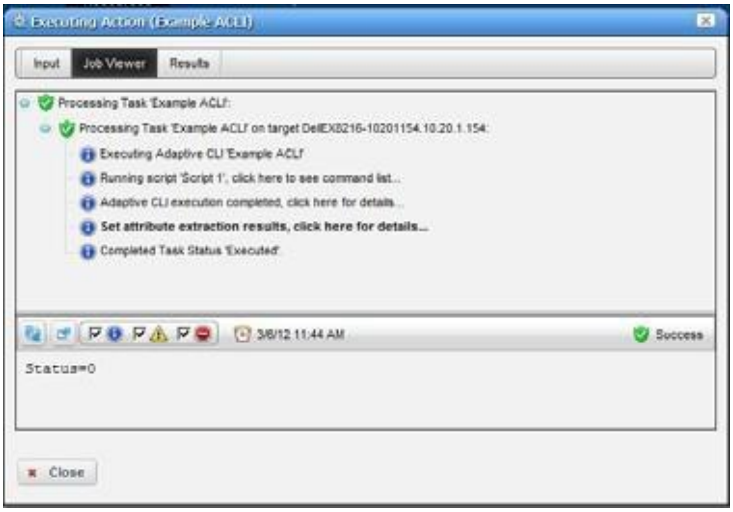

Processing Task 'Example ACLI':
Processing Task 'Example ACLI' on target DellEX8216-10201154.10.20.1.154:
Executing Adaptive CLI 'Example ACLI'
Running script 'Script 1', click here to see command list...
Adaptive CLI execution completed, click here for details...
Set attribute extraction results, click here for details...
Completed Task Status 'Executed'

3/8/12 11:44 AM Success

```
Status=0
```

Click *Set attribute extraction results, click here* to see the results appear in the bottom panel. Notice also that you must check informational messages for all these to appear, and that several additional sets of messages besides the extraction results appear.

Create a Monitor for the External Script Adaptive ACLI

Now that you have verified the script is working, you can create a monitor to see how this attribute is doing.

1. In the Monitors portlet, create a new ACLI Monitor
2. Uncheck *Update Network Status* (recommended since the ICMP monitor is already doing this)
3. You may want to test your monitor, in which case, change the monitoring interval to 30 seconds. Re-edit it to configure it with the interval needed for your production system.
4. In *Monitor Options* select your example monitor configured previously.
5. Confirm that *Monitor Attributes* displays the Status attribute configured previously.
6. In the *Conditions* tab of the Monitor Editor, create “Status Up” condition, with the severity of *Informational*, and check *Alert*.
7. Create a criterion which is Status = 0.
8. Save this condition
9. Create a new Condition called “Status Down”
10. The criterion is Status = 1
11. Apply and Save
12. Save your monitor.
13. Right-click to select *View Monitor Data*, and you can see the results of yourefforts.

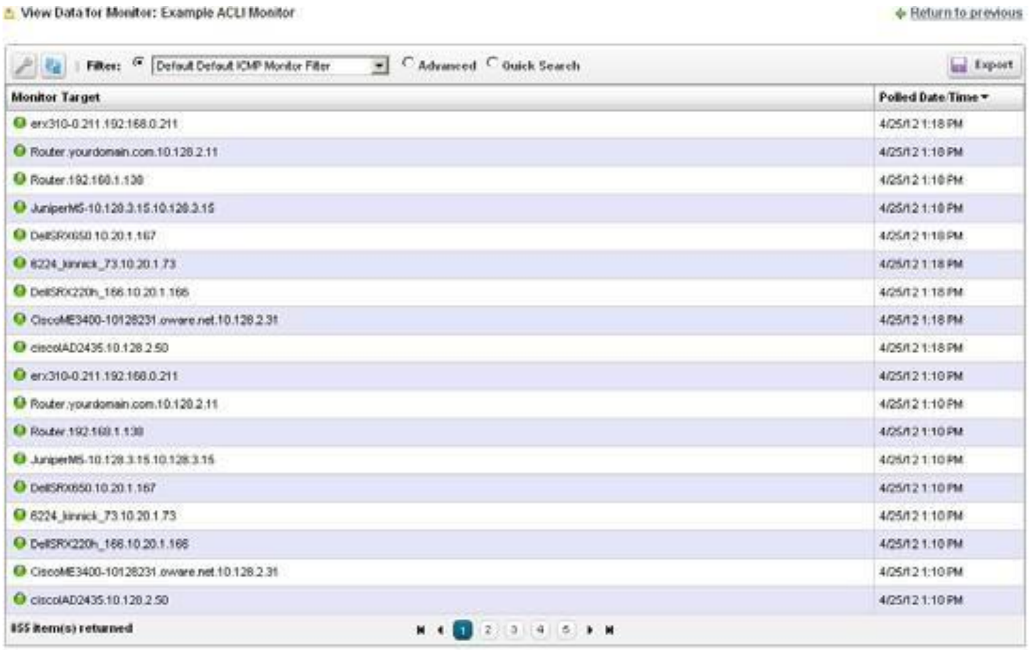

View Data for Monitor: Example ACLI Monitor

Return to previous
Filter: ● Default Default ICMP Monitor Filter ○ Advanced ○ Quick Search
Export

Monitor Target	Polled Date/Time
erx310-0.211.192.168.0.211	4/25/12 1:18 PM
Router.yourdomain.com.10.128.2.11	4/25/12 1:18 PM
Router.192.168.1.138	4/25/12 1:18 PM
JuniperM5-10.128.3.15.10.128.3.15	4/25/12 1:18 PM
DellSRX650.10.20.1.167	4/25/12 1:18 PM
6224_kinnick_73.10.20.1.73	4/25/12 1:18 PM
DellSRX220h_166.10.20.1.166	4/25/12 1:18 PM
CiscoME3400-10128231.oware.net.10.128.2.31	4/25/12 1:18 PM
ciscoIAD2435.10.128.2.50	4/25/12 1:18 PM
erx310-0.211.192.168.0.211	4/25/12 1:10 PM
Router.yourdomain.com.10.128.2.11	4/25/12 1:10 PM
Router.192.168.1.138	4/25/12 1:10 PM
JuniperM5-10.128.3.15.10.128.3.15	4/25/12 1:10 PM
DellSRX650.10.20.1.167	4/25/12 1:10 PM
6224_kinnick_73.10.20.1.73	4/25/12 1:10 PM
DellSRX220h_166.10.20.1.166	4/25/12 1:10 PM
CiscoME3400-10128231.oware.net.10.128.2.31	4/25/12 1:10 PM
ciscoIAD2435.10.128.2.50	4/25/12 1:10 PM

855 item(s) returned

Action Groups

The Action Groups Portlet lets you configure several actions with different targets, order their execution, and execute actions against these groups of targets all at once.

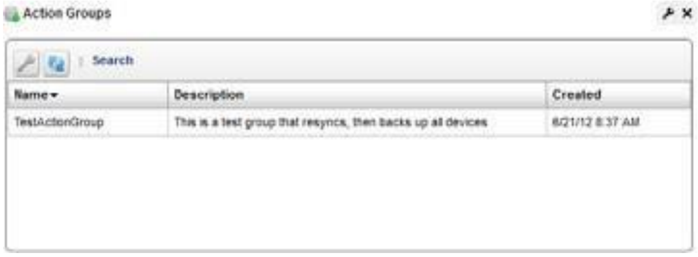

Name	Description	Created
TestActionGroup	This is a test group that resyncs, then backs up all devices	6/21/12 8:37 AM

Right click within this portlet to create a New Action Group, Edit an existing, selected one, Execute the selected Action Group, view the Audit trails for the selected group, or Schedule the selected Action Group. You can also share action groups with other users on your system.

The Action Group Editor on page 595 configures new or existing actions. See also Configure an Action Group on page 596 below.

Action Group Editor

This editor lets you create or edit an Action Group.

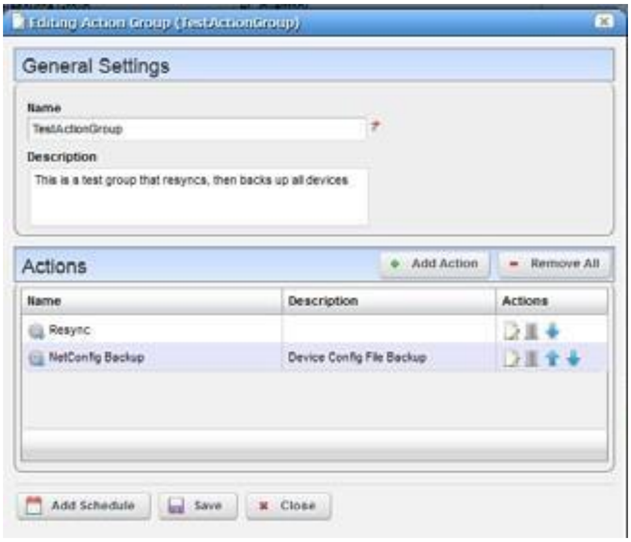

Name: TestActionGroup

Description: This is a test group that resyncs, then backs up all devices

Name	Description	Actions
Resync		
NetConfig Backup	Device Config File Backup	

It has the following fields and buttons:

Name—A unique text identifier for the Action Group

Description —A text description for the Action Group

Actions

This portion of the screen lists all added actions. Use the *Edit this entry* icon (pencil and paper) to edit individual actions, or the *Delete this entry* icon to delete individual actions. The up/down arrows configure the order of execution (top first). Editing an entry opens the editor described in *Add Action* below.

Add Action—This opens an Action Editor where you can select the Action that is to be a member of the group, its Target devices and any Parameters associated with the Action.

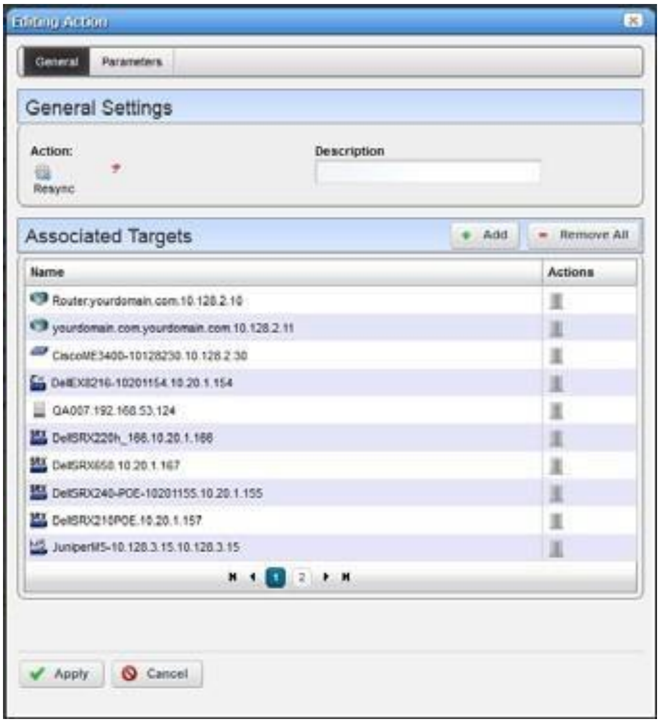

Name	Actions
Router.yourdomain.com.10.128.2.10	
yourdomain.com.yourdomain.com.10.128.2.11	
CiscoME3400-10128230.10.128.2.30	
DellEX8216-10201154.10.20.1.154	
QA007.192.168.53.124	
DellSRX220h_166.10.20.1.166	
DellSRX650.10.20.1.167	
DellSRX240-POE-10201155.10.20.1.155	
DellSRX210POE.10.20.1.157	
JuniperM5-10.128.3.15.10.128.3.15	

Use the *Add* button to add Associated Targets, and the *Delete this entry* icon to delete any added by mistake. Click *Apply* to accept an added (or edited) Action. When you do this the list on the Actions panel of the Action Group Editor changes to reflect the changes you have made.

Remove All — Delete all Actions.

Click *Save* to create the Action Group. Once you have saved the group, you can right-click to *Execute* it manually. Youcan also click *Add Schedule* to schedule its execution. Clicking *Close* ends your editor session without saving any new Action Group or changes you may have made to an existing one.

NOTE:

When you execute an Action Group, the Results view displays a list of targets on the left, and results for the selected target on the right. Click on a different target to see the target's results.

How To:

Configure an Action Group

Follow these steps to configure an Action Group

1 Right-click in the Action Groups portlet and select *New.*

2 Name the Action Group, and optionally type a Description.

3 Click *Add Action.*

4. Select an action (with the green plus at the top of the screen), and optionally add a Description in the field to the right of the selection.
5. To associate devices with this action, click the *Add* button in the Associated Targets panel, and select the device(s).
6. Click *Done*.
7. Review the selected devices and click *Apply* once that list is correct.
8. Review the added Actions, and insure they are in the correct order. Re-order them, if necessary.
9. Click *Save* to preserve your new Action Group.
10. Right-click and select *Execute* to test the Action Group.
11. If appropriate, and execution is correct, right-click to *Schedule* for this Action Group. Configure its occurrence in the Schedule panel of the screen that appears next and *Save* that schedule.

Troubleshooting Adaptive CLI

The following issues can prevent the correct completion of Adaptive CLI execution.

Connectivity—The device can be offline. To detect whet her this is true, right-click the device in the Managed Resources portlet and *Direct Access* > *Ping* it.

Incomplete Discovery —If the device is online and still does not respond to Adaptive CLI, you mayhave onlypartiallydiscoveredit. Right-click the device in the Managed Resourcesportlet and select *Direct Access* > *Terminal.*If that menu option does not exist, it is only partially discovered. Right-click to edit the device, and add a Telnet Management Interface and Authentication in those two tabs of the editor.

Timeouts— Adaptive CLI timeouts may occur because of an unresolved *Continue* prompt. See *Continue Pattern* on page 557 for instructions about how to resolve such things. Depending on the device, you may also configure the device itself not to emit patterns that need a response.

Adaptive CLI Records Aging Policy

You can use Cruz's aging feature to preserve Adaptive CLI information. Click the Redcell > Database Aging Policy (DAP) node of the Control panel, and click the default *Adaptive CLI DAP* and click the edit button on its right.

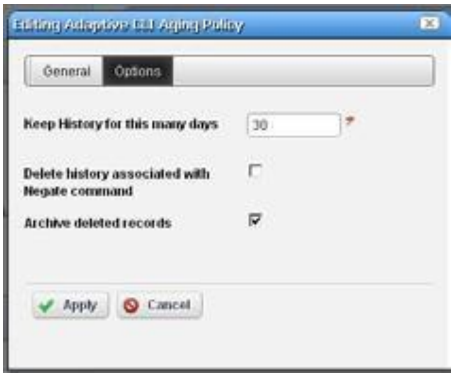

Keep History for this many days: 30

Delete history associated with Negate command ☐

Archive deleted records ☑

After filling in the *General Info* tab, the *Parameters* screen lets you configure the following:

Keep History — Enter the number of days to retain the history in the database.

Delete history associated with Negate command—Check to remove archived records associated with *Negate* (described under General on page 546).

Archive Deleted Records—Check to have deleted archived records saved as a file (configured in the *General Info* parameters too).

Web Service DeploymentFeatures

Right clicking in Actions Portlet supports deploying actionsasweb services, asdescribed in the following section. Web services are typically the concern of administrators, not operators. Administrators using these features are expected to be familiar with the web service technology they configure. Refer to the *Web Services Guide* for more about WSDL.

Deploy Web Service—You can select one or more Actions to deploy as a web service. You must assigneach selected activity a unique Web Service ID, that can ultimately appear in acolumn in the Activities Manager screen. The screen that appears after you select this item lets you assign this ID for each activity connected to the web service. This screen's appearance depends on the web service you select.

The Web Service ID can contain only alphanumeric characters and underscores (_) and must start with either a letter or underscore. This ID represents the input data class (or type) of the activity. It defaults to a valid name reflecting the activity name. You can change the default. Upon successful deployment, the Web Service IDs of the deployed activities appear in the manager. Youcan then export the WSDL file for code generation and web service invocation.

NOTE:

Youcan seedeployed Axis2 webservices listedinthescreen athttp://[application server IP address]:8089/axis2/services/listServices. These may take a little time to appear, so be patient. If you have been patient, and they still do not appear listed, you may have to clear your browser's cache. Clicking the *Activity* link once they appear displays the WSDL.

Undeploy Web Service—Select one or more activities to undeploy from web service. When successful, the Web Service IDs of the undeployed activities are cleared from the manager. Undeployed activities are also no longer accessible for web service requests.

Export WSDL—After deploying and undeploying activiti es, you may want to export the WSDL file for client code generation. The WSDL file contains all the data class types for web service execution.

All activity web service client code shares a common web service method: `TaskSvcExecute`, which takes in `TaskSvcExecuteInParams` as input data and returns `TaskSvcExecuteOutParams` as output data. The following describes the input and output data:

Input Data:

Async—A Boolean type. If true, indicates an asynchronous (not synchronous) request. Asynchronous requests return immediately with a job ID, while synchronous requests await completion of the web service execution before returning job IDs. Either way, you must then use the Job web service to examine the results.

TargetOID —A string representing the ID of the Target Entity Type.

TaskData—This is the base parameters class to be replaced by the activity-specific extended data class (namely the Web Service ID) associated with the activity execution.

Output Data:

JobID—This is a unique ID string used for querying the results of an execution. The job web service is defined in `$OWARE_DEV_CACHE_CLS/ws.war/wsdl/Job.wsdl`. For results, the ParentJobID Status is one of the following values:

0 - Execution in progress (if async request)

1 - Execution succeed

2 - Execution failure

Status is an integer representing one of these values:

0 - Execution succeed

1 - Invalid input data

2 - Execution pending (for async request)

6 - Execution failure

8 - Invalid input target OID

NOTE:

The Status of the TaskSvcExecuteOutParams is different than the status of the job service.

TaskData—The base data class (same as input) returned as output data. Certain activities may produce output data.

RESTful Web Service

This section provides an additional Adaptive CLI type that executes a REST W eb Service. Like other Adaptive CLI behaviors, the application accepts user data and includes that in the REST API payload as JSON scripts. With this type of Adaptive CLI, JSON scripts (which include parameters, not just JSON) can support REST operations: GET , PUT, POST, and so on. You can specify a target URL and authentication. The Cruz application maps any resulting data back to the user data object associated with the Adaptive CLI.

The following topics are covered:

Rest Call: No Authentication/Token

The following example makes a simple REST call without needing an authentication/token:

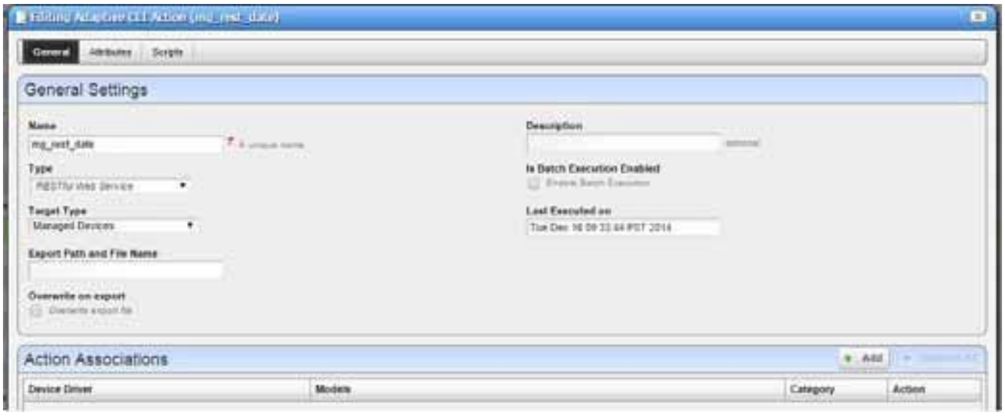

Editing Adaptive CLI Action (mg_rest_date)

General Settings

Name: mg_rest_date * A unique name

Type: RESTful Web Service

Target Type: Managed Devices

Export Path and File Name:

Overwrite on export: ☐ Overwrite export file

Description: optional

Is Batch Execution Enabled: ☐ Enable Batch Execution

Last Executed on: Tue Dec 16 09:33:44 PST 2014

Action Associations

Device Driver	Models	Category	Action

The script in this example:

```
URL:http://date.jsontest.com
requestMethod:GET
```

The result:

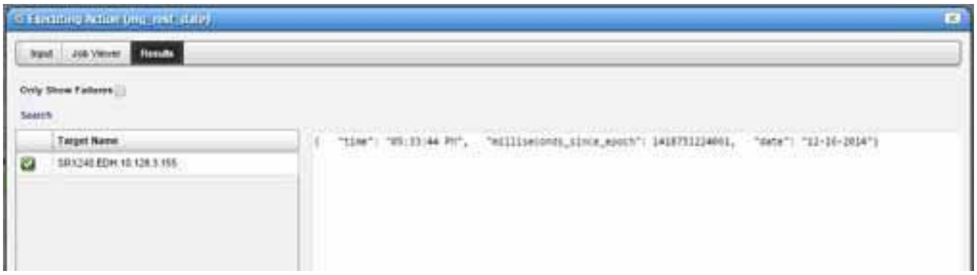

Executing Action (mg_rest_date)

Only Show Failures ☐

	Target Name
✓	SRX240 EDH 10.128.3.155

```
{   "time": "05:33:44 PM",   "milliseconds_since_epoch": 1418751224001,   "date": "12-16-2014"}
```

NOTE:

Sometimes it is useful to specify a target even if none is required so that you can see the Results window. Otherwise, the Job Viewer panel displays the same results. You can also copy and paste output into an online formatter at http://jsonformatter.curiousconcept.com/ to rearrange it so its appearance is more user-friendly.

REST Call Requiring Token Retrieval

The following example is of a REST call requiring Token retrieval before the subsequent request.

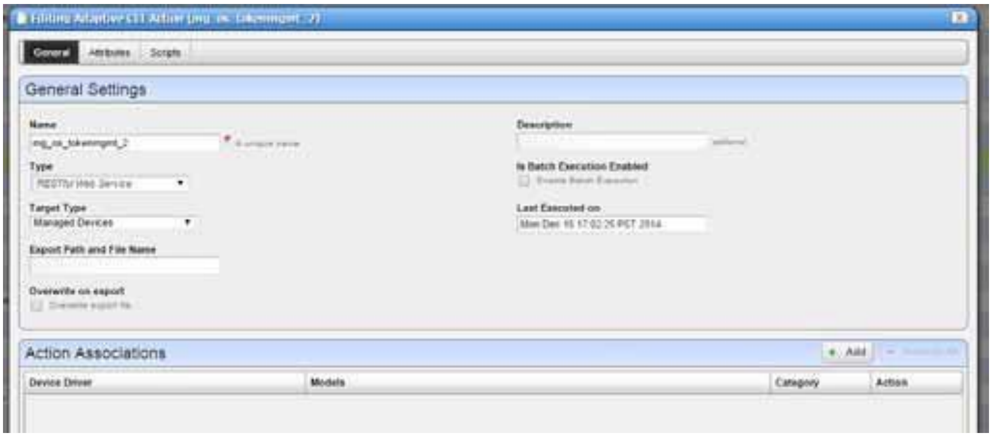

Name: mg_os_tokenmgmt_2 * A unique name

Description: optional

Type: RESTful Web Service

Is Batch Execution Enabled: □ Enable Batch Execution

Target Type: Managed Devices

Last Executed on: Mon Dec 15 17:02:25 PST 2014

Export Path and File Name:

Overwrite on export: □ Overwrite export file

Device Driver	Models	Category	Action

The script contents:

```
TOKEN_BEGIN
TokenProperty: X_Auth_Token
TokenAccessParams: access.token.id
URL:http://10.101.53.2:5000/v2.0/tokens
Method:POST
{
"auth": {"tenantName": "admin","passwordCredentials":
  {"username":"admin","password":"admin"} }
}
TOKEN_END
URL:http://10.101.53.2:5000/v2.0/tenants
Method:GET
```

△ CAUTION:

Do not put lines between TOKEN_END and URL.

Execution Result:

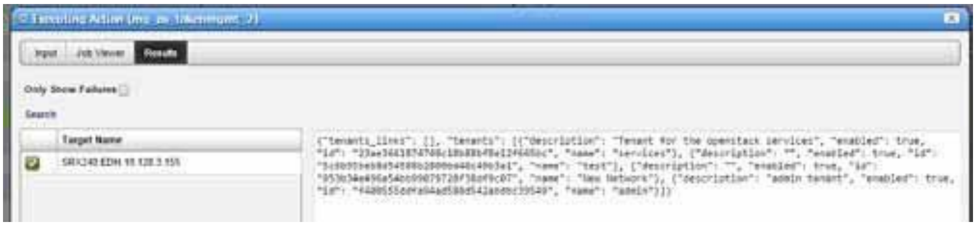

	Target Name
☑	SRX240 EDH 10.128.3.155

```
{"tenants_links": [], "tenants": [{"description": "Tenant for the openstack services", "enabled": true,
"id": "[illegible]", "name": "services"}, {"description": "", "enabled": true, "id":
"[illegible]", "name": "test"}, {"description": "", "enabled": true, "id":
"[illegible]", "name": "New Network"}, {"description": "admin tenant", "enabled": true,
"id": "[illegible]", "name": "admin"}]}
```

NOTE:

The Job Viewer does not display credentials.

If you have already generated a token and saved it in cache, a Token Generation Action that refers to that particular token by name uses the saved token. If any error exists in the Token Generation Action, you may not notice that error in this execution. However, if you execute the same Token Generation Action and the saved token has expired, itmay trigger errors.

REST Call Requiring Base64-Encoded Authentication/ Token Retrieval

The following example script demonstrates a REST call that requires Base64 encoded Authentication and Token retrieval. The subsequent REST call uses the Token retrieved and has a JSON payload.

Example script:

```
TOKEN_BEGIN
TokenProperty:Authorization
TokenPrefix: Bearer
TokenAccessParams: access_token
TokenExpiryTime:1440
TokenAuthenticationType: base64
TokenAuthenticationID:AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
  AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
URL:https://api.sandbox.paypal.com/v1/oauth2/token
RequestMethod:POST
grant_type=client_credentials
TOKEN_END
URL:https://api.sandbox.paypal.com/v1/payments/payment
RequestMethod:POST
{
"intent":"sale",
"payer":{
"payment_method":"credit_card",
"funding_instruments":[
{
"credit_card":{
"number":"5555555555550331",
"type":"visa",
"expire_month":11,
"expire_year":2018,
"cvv2":"555",
"first_name":"Betsy",
"last_name":"Buyer",
"billing_address":{
"line1":"111 First Street",
"city":"Saratoga",
"state":"CA",
"postal_code":"95070",
"country_code":"US"
}
}
}
]
},
"transactions":[
{
"amount":{
"total":"7.47",
"currency":"USD",
"details":{
"subtotal":"7.41",
"tax":"0.03",
"shipping":"0.03"
}
},
"description":"This is the payment transaction description."
}
]
}
```

Result:

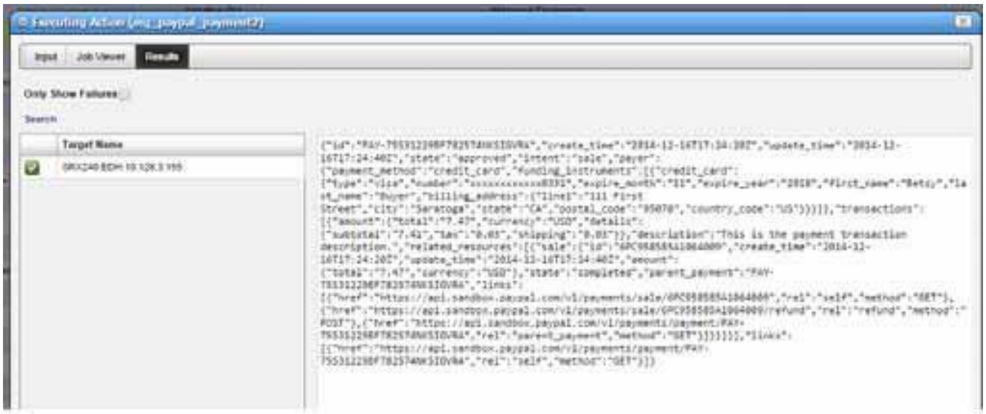

```
{"id":"PAY-75531229BF782574ANKSIOVRA","create_time":"2014-12-16T17:34:20Z","update_time":"2014-12-
16T17:24:40Z","state":"approved","intent":"sale","payer":
{"payment_method":"credit_card","funding_instruments":[{"credit_card":
{"type":"visa","number":"xxxxxxxxxxxx0331","expire_month":"11","expire_year":"2018","first_name":"Betsy","la
st_name":"Buyer","billing_address":{"line1":"111 First
Street","city":"Saratoga","state":"CA","postal_code":"95070","country_code":"US"}}}}]},"transactions":
[{"amount":{"total":"7.47","currency":"USD","details":
{"subtotal":"7.41","tax":"0.03","shipping":"0.03"}},"description":"This is the payment transaction
description.","related_resources":[{"sale":{"id":"6PC95858A1064009","create_time":"2014-12-
16T17:24:20Z","update_time":"2014-12-16T17:34:40Z","amount":
{"total":"7.47","currency":"USD"},"state":"completed","parent_payment":"PAY-
75531229BF782574ANKSIOVRA","links":
[{"href":"https://api.sandbox.paypal.com/v1/payments/sale/6PC95858A1064009","rel":"self","method":"GET"},
{"href":"https://api.sandbox.paypal.com/v1/payments/sale/6PC95858A1064009/refund","rel":"refund","method":"
POST"},{"href":"https://api.sandbox.paypal.com/v1/payments/payment/PAY-
75531229BF782574ANKSIOVRA","rel":"parent_payment","method":"GET"}]}}]}],"links":
[{"href":"https://api.sandbox.paypal.com/v1/payments/payment/PAY-
75531229BF782574ANKSIOVRA","rel":"self","method":"GET"}]}
```

The following format is for the Token portion of the Script. The API used for Token Generation in the REST executions depends on the source queried. Sample1 uses the Paypal API. OpenStack API appears in Sample 2.

- **Sample1**, a base64 authentication with token prefix:

  ```
  TOKEN_BEGIN
  TokenProperty:Authorization
  TokenPrefix: Bearer
  TokenPostfix:
  ```

```
TokenAccessParams: access_token
TokenExpiryTime:1440
TokenAuthenticationType: Base64
TokenAuthenticationID:abc:xyz
URL:https://api.sandbox.paypal.com/v1/oauth2/token
RequestMethod:POST
grant_type=client_credentials
TOKEN_END
```

- **Sample2**, token value extracted from a nested JSON object :

```
TOKEN_BEGIN
TokenProperty: X_Auth_Token
TokenAccessParams: access.token.id
TokenExpiryTime:1440
URL:http://10.101.53.2:5000/v2.0/tokens
Method:POST
RequestPropertyName:abc RequestPropertyValue:def
RequestPropertyName:ghi RequestPropertyValue:jkl
{
"auth": {"tenantName": "admin","passwordCredentials":
  {"username":"admin","password":"admin"} }
}
TOKEN_END
```

The standard Request Properties form-urlencoded are part of the execution code and do not need to be specified:

```
Accept: application/json
Content-Type:application/json;charset=UTF-8
Base64 encoding Content-Type:application/x-www-
```

Callingan Adaptive CLI

The following returns a list of services from the REST call with this JSON script:

```
URL: http://192.168.54.43:8089/rest/application.wadl
Method: GET
PropertyName: Accept PropertyValue: application/xml
```

The following example executes an ACLI action using a REST call:

1 POST http://localhost:8089/rest/auth

```
{   "username" : "admin",   "password" : "secretpassword" }
```

This works like OpenStack’s API, you get a tokenId from that call, add it to the request header `"X-Auth-Token" : tokenId`on subsequent calls.

Alternatively, create a getToken action to reuse:

```
TOKEN_BEGIN
TokenProperty: X-Auth-Token
TokenAccessParams: tokenId
```

```
URL:http://192.168.54.43:8089/rest/auth
Method:POST
{
"username" : "admin",
"password" :"admin"
}
TOKEN_END
```

2 GET http://localhost:8089/rest/actiondefs/

3 Pick a suitable ACLI that you want to call:

```
TokenActionName: getToken
URL: http://192.168.54.43:8089/rest/actiondefs/
Method:GET
```

4 Select this action oid from the result payload:

```
  {
     "name":"Cisco \u0027show hardware\u0027",
     "oid":"com.dorado.redcell.inventory.task.TaskDefinition::ZvfczDh-
 e60G03-sQ-1.313",
     "description":"show hardware",
     "family":"Adaptive CLI",
     "targetType":"EquipmentManager",
     "implementor":"AC_Config"
  },
```

5 GET http://localhost:8089/rest/actiondefs/{oid-from-4}/sample-json

This returns a JSON structure that you can populate, such as:

```
TokenActionName: dp_getToken
URL: http://192.168.54.43:8089/rest/actiondefs/
 com.dorado.redcell.inventory.task.TaskDefinition::ZvfczDh-e60G03-sQ-
 1.313/sample-json
Method: GET
```

The result after execution:

```
{"name":"dp_getDate","description":null,"oid":null,"target":null,"definit
 ionOid":"com.dorado.redcell.inventory.task.TaskDefinition::ZvfczDh-
 e60G03-sQ-
 1.313","sessionId":null,"status":null,"ordinal":0,"groupId":null,"lastE
 xecuted":null,"lastModified":null,"dateCreated":null,"dataValues":[]}
```

6 POST http://localhost:8089/rest/actions/

The JSON structure from the previous step creates a task in Cruz:

```
TokenActionName: dp_getToken
URL: http://192.168.54.43:8089/rest/actions/
Method:POST
{
"name":"Cisco 'show
 hardware'","description":null,"oid":null,"target":"com.dorado.redcell.d
 evicedriver.cisco.ciscorouter.CiscoRouter::WFKyfBmJkHo3G03","definition
 Oid":"com.dorado.redcell.inventory.task.TaskDefinition::ZvfczDh-e60G03-
 sQ-
 1.313","sessionId":null,"status":null,"ordinal":0,"groupId":null,"lastE
 xecuted":null,"lastModified":null,"dateCreated":null,"dataValues":[]}
}
```

7 PUT http://localhost:8089/rest/actions/{oid-from-6}/execute

Executes the following task:

```
TokenActionName: getToken
URL: http://192.168.54.43:8089/rest/actions/
 com.dorado.redcell.inventory.task.TaskDefinition::ZvfczDh-e60G03-sQ-
 1.313/execute
Method:PUT
{
"name":"Cisco 'show
 hardware'","description":null,"oid":null,"target":"com.dorado.redcell.d
 evicedriver.cisco.ciscorouter.CiscoRouter::WFKyfBmJkHo3G03","definition
 Oid":"com.dorado.redcell.inventory.task.TaskDefinition::ZvfczDh-e60G03-
 sQ-
 1.313","sessionId":null,"status":null,"ordinal":0,"groupId":null,"lastE
 xecuted":null,"lastModified":null,"dateCreated":null,"dataValues":[]}
}
```

Calling an Independent Token Generation Action

If a cached token has expired or is invalid, OpenManage Network Manager generates a new token by calling an independent action and saves it in cache. This executes an action requiring authentication.

The format of token generation script (mg_os_generateTokenOnly) is as follows:

```
TOKEN_BEGIN
TokenProperty: X_Auth_Token
TokenAccessParams: access.token.id
URL:http://10.101.53.2:5000/v2.0/tokens
Method:POST
{
"auth": {"tenantName": "admin","passwordCredentials":
 {"username":"admin","password":"admin"} }
}
TOKEN_END
```

The format of a script that uses this token generated by `mg_os_synergy_generateTokenOnly` is:

```
    TokenActionName: mg_os_synergy_generateTokenOnly
    URL: http://localhost:8089/rest/actiondefs/
Method:GET
```

In such a script, TokenActionName is the name of the token generation RESTful Web Services Action.

Supported Script Formats

The following are some sample scripts that confirm to the designed format:

- Sample scripts that do not need authentication
- Sample scripts of a token generation action
- Sample scripts that use a token generation action
- Sample scripts that combine token generation and execution in one action

Sample scripts that do not need authentication

```
URL:http://headers.jsontest.com/
Method:GET

URL:http://10.101.53.2:5000/v2.0/tokens
Method:POST
{
"auth": {"tenantName": "admin","passwordCredentials":
  {"username":"admin","password":"admin"} }
}
```

Sample scripts of a token generation action

```
    TOKEN_BEGIN
    TokenProperty: X-Auth-Token
    TokenAccessParams: tokenId
    TokenExpiryTime: 86400
    URL:http://192.168.53.50:8089/rest/auth/
    Method:POST
    { "username" : "admin", "password" : "admin" }

TOKEN_END
```

The following script contains all available fields for demonstration purposes.

```
TOKEN_BEGIN
TokenProperty: X-Auth-Token
TokenPrefix:Bearer
TokenPostfix:
TokenAccessParams: access.token.tokenId
TokenExpiryTime: 86400
TokenAuthenticationType:Base64
TokenAuthenticationID:abc:xyz
```

```
    URL:http://localhost:8089/rest/auth/
    Method:POST
    RequestPropertyName:abc RequestPropertyValue:def
    RequestPropertyName:ghi RequestPropertyValue:jkl
    { "username" : "admin", "password" : "admin" }

TOKEN_END
```

△ CAUTION:

You must format JSON script headers (example: Accept:application/xml) as follows: `PropertyName:Accept PropertyValue:application/xml.`

Sample scripts that use a token generation action

```
    TokenActionName: mg_synergy_generateTokenOnly
    URL: http://localhost:8089/rest/actiondefs/
    Method:GET

    TokenActionName:mg_pp_tokenGenerationOnly
    URL:https://api.sandbox.paypal.com/v1/payments/payment
RequestMethod:GET
```

Sample scripts that combine token generation and execution in one action

```
    TOKEN_BEGIN
    TokenProperty: X-Auth-Token
    TokenAccessParams: tokenId
    TokenExpiryTime: 86400
    URL:http://192.168.53.50:8089/rest/auth/
    Method:POST
    { "username" : "admin", "password" : "admin" }

    TOKEN_END
    URL: http://localhost:8089/rest/workorderdefs/
Method:GET
```

This is a functioning script with changed values for TokenAuthenticationID

```
    TOKEN_BEGIN
    TokenProperty:Authorization
    TokenPrefix: Bearer
    TokenAccessParams: access_token
    TokenExpiryTime:1440
    TokenAuthenticationType: base64
    TokenAuthenticationID:AQ6pmxAg:14VJ8Y6fA
    URL:https://api.sandbox.paypal.com/v1/oauth2/token
```

```
Method:POST
grant_type=client_credentials
TOKEN_END
URL:https://api.sandbox.paypal.com/v1/payments/payment
Method:POST
{
"intent":"sale",
"payer":{
"payment_method":"credit_card",
"funding_instruments":[
{
"credit_card":{
"number":"55555555550331",
"type":"visa",
"expire_month":11,
"expire_year":2018,
"cvv2":"874",
"first_name":"Betsy",
"last_name":"Buyer",
"billing_address":{ "line1":"111 First Street", "city":"Saratoga",
 "state":"CA", "postal_code":"95070", "country_code":"US" }
}
}
]
},
"transactions":[
{
"amount":{
"total":"7.47",
"currency":"USD",
"details":{ "subtotal":"7.41", "tax":"0.03", "shipping":"0.03" }
},
"description":"This is the payment transaction description."
}
]
}
```

12

Serving Multiple Customer Accounts

Multitenancy lets you manage your customers' tenant networks from a central Cruz while at the same time providing them with secure access to their specific resources. Cruz does this with a two-tiered customer multitenant service provider (MSP) model, assigning each customer a specific security domain for their specific resources. Cruz then restricts access to the domain to the customer or a member of the MSP domain See also Multitenant Batch Imports in the next section.

Multitenancy is a Cruz extension.

Multitenant Batch Imports

You can import Discovery Profiles, Authentications and Contacts to target multitenant domains with a command line importer. The command is `import[item]`, for example `importprofiles`, and these commands are in the `owareapps/redcell/bin` directory. These commands take the import file name as an argument. The required domains should exist in the Cruz before import occurs. Import authentications before importing discovery profiles that refer to them. Example XML files with the `<customer>` tag for domains are in `owareapps/redcell/db`.

Configuring Chat forMultitenancy

Sometimes you may see Colleagues not members of a tenant site in the tenant site's status bar. This is not necessary. To change this so only tenant site personnel appear in the Colleagues panel, do the following:

1. In Control Panel, under *Portal > Users and Organizations* click the link under *Name* to the tenant site.
2. With the *Actions* button in the relevant User's field, select *Edit*.
3. In the *Sites* panel (the link is on the right), remove Cruz, and any other sites present, except the tenant site itself. Only the tenant site remains as the source of chat colleagues who appear in the status bar.
4. Go to the tenant site. You should only be able to see authorized users for the tenant site who are on the tenant site. If you are on the master site, then you can still see all users.

Configuring Multitenancy, Site Management and Access Profiles

Cruz can support multiple client organizations from a single instance. It can also constrain data access for logged on users depending on the client organization to which they belong. To implement multitenancy, follow the How to: Create a Multitenancy Environment on page 613 below.

To understand the data constraints on tenants within a Multitenant Cruz environment, see Constraining Data Access on page 626. See also User Site Access on page 624 for how to give users access to multiple tenant sites.

The following diagram describes the relationship between the pieces of Multitenancy in the order needed to configure a tenant site. See it in narrative form in How to: Create a Multitenancy Environment on page 613

Multitenancy Components

Control Panel

(Multi-Page) Site Template

Page Template

→ Action > View Page
To add/arrange portlets on these templates

Site Management Portlet

Create Tenant Site

(Must assign previously created Site or Page Template while creating the new site for the first time)

Access Profile Templates Portlet

Create Access Profile Templates

(Can configure Cruz capabilities for groups of sites)

Site Management Portlet

Configure Tenant Site Access Profile

(Right-click for AP Editor/AP Template EditoA Assign an Access Profile Template to be logically ed with the site's own AccessProfile.)

Managed Resources Portlet

Right-click to assign devices to Tenant sites.
View and Further Configure Devices for the Site -- page 614.

User Site Access Portlet

You can optionally configure multi-site access for selected users. See Configure User Site Accesson page 615.

Discovery Profiles Portlet

You can optionally assign a discovery profile to a tenant site. See Assigning Sites in Discoveryon page 615.

Provisioning Site-Creating User Permissions

The user(s) allowed to create tenant sites with Site Templates and the other features documented here are typically Administrators. However, if you want to allowanother User like *Tenant Admin* to make tenant sites, you must grant that user's Ro le permissions in two areas within Control Panel:

Redcell > Permission Manager

Click the edit button to the right of the assigned Role, and grant the MSP-related permissions found (*MSP Access Profile Templates* and *MSP Site Management*). Search *All* permissions with the magnifying glass at the bottom of the editor to find these permissions.

Portal > Roles

Grant the selected user's Role permissions within this panel. Click the Role itself, then click the *Define Permissions* panel in the editor. Within the drop-down list, under *Site Content*, grant *Site Template* permissions if you want the user to create templates, and under *Portal*, grant *Sites > Site* and *Sites > Site Pages Variation* permissions configured as appropriate.

How To:

Create a Multitenancy Environment

Creating a Site Template

1 A tenant's website or portal is based on a Site or Page Templates you have created. See How to: Make A Site Template on page 618. Also, see Click "+ Add page" on the grey bar in the middle of the screen, fill in data and click Add page at the bottom on page 620 for instructions about creating them.

2 Right-click in the Site Management Portlet to create a tenant site. Click *Add Site* and the editor appears (see Site Management Editor on page 621).

3 While the Site Management Editor is still open, assign a Site or Page Template you have already created (see Portal > Site Templates on page 618 and Click " + Add page" on the grey bar in the middle of the screen, fill in data and click Add page at the bottom onpage 620). This creates the default website for the customer. Youcan only do this when creating the customer, you cannot do it after the customer is created.

△ CAUTION:

Deleting a Page Template already in use makes the page unusable. Don't do it! Also: A Site Admin can choose to hide some of atenant site's pages from a Site Member through page permissions. The page permissions do not work, however, if the tenant site comes from a site template.

4 *Save* the tenant site.

5 If you want to configure the site further, click *Action > View Pages* in site as it appears in the list of tenant sites. You can add portlets, and they will become part of that site.

NOTE:

The Site Map portlet discloses a site's available pages and the portlets configured to be on each page. Bydefault, theSite Mapportletisavailable fromtheSettingspagebyselecting theSite Mapmenu option. Also: Sites appear in the Control Panel under *Top Level Organizations*, along with the number of assigned users.

Site Map Portlet

Use the Site Map portlet to see where pages, sub-pages, and portlets are located within your installation. Use your browser's search function to find portlet names within this Site Map.

Access this portlet by selecting the Settings > Site Map menu option.

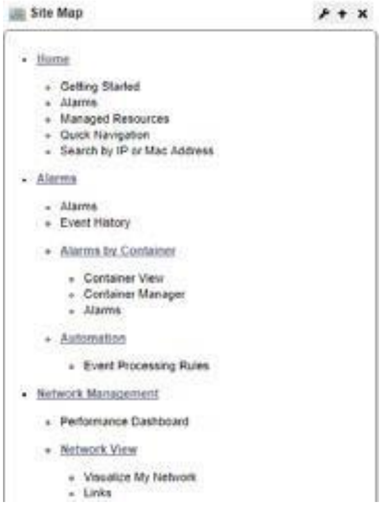

Site Map

- Home
  - Getting Started
  - Alarms
  - Managed Resources
  - Quick Navigation
  - Search by IP or Mac Address
- Alarms
  - Alarms
  - Event History
  - Alarms by Container
    - Container View
    - Container Manager
    - Alarms
  - Automation
    - Event Processing Rules
- Network Management
  - Performance Dashboard
  - Network View
    - Visualize My Network
    - Links

Creating an Access Profile Template

6 Access Profile Templates configure tenant site access to various capabilities. These templates provide an easy way to grant levels of access to resources (for example: *bronze, silver, gold*) to your customer base too.

Optionally, create an Access Profile Template, using the Access Profile Templates portlet described below. Right-click and select *New* in the Access Profile Templates portlet to create a template (see AP Editor/AP Template Editor on page 623).

To create an individual site's Access Profile, right-click the tenant site you created in the Site Management portlet and select *Access Profile* from the menu. Notice you can select from available Access Profile Templates at the top of this screen too.

Every customer automatically has an Access profile which consists of any selected Access Profile Template and any customer-specific Access Profile you configure. Configure the Customer Resource Assignments byclicking *Add* in the Access Profile Editor. Cruz logically ANDs these with any selected template.

View and Further Configure Devices for the Site

7 To see the tenant website, right click the tenant site in Site Management, and select *Go to Site.* The tenant website opens, and you will see all the pages you configured when Creating a Site Template .

8 If you want to confine discovered devices to a particular tenant site, right-click them in the master site's Managed Resources. Right-click a device and select Manage > Domain Access Control to see the menu of available tenant sites.

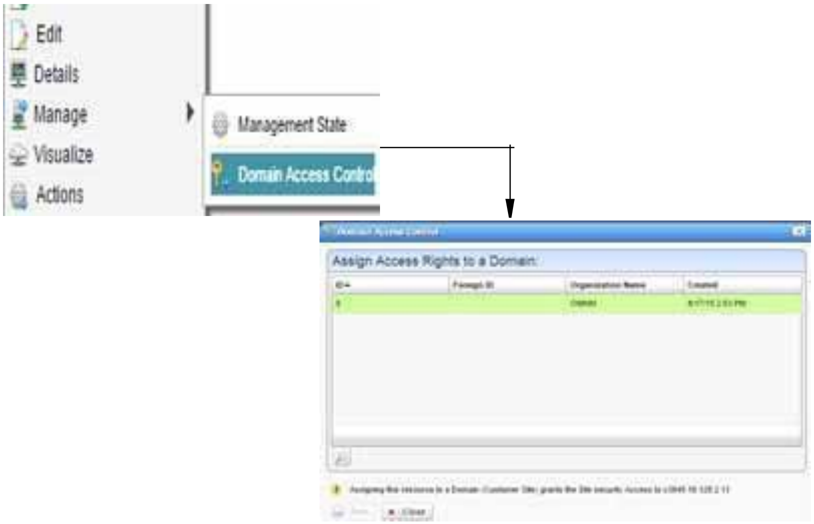

Bydefault, onlydeviceseither discoveredforanassigned site, or thoseexplicitlyassigned with this procedure appear in tenant sites.

Configure User Site Access

9 Youcan optionally assign specific users' access to multiple tenant sites. See How to: Configure User Site Access on page 625 for step-by-step instructions.

Assigning Sites in Discovery

To create discovery profiles for a particular site, select the site in the *New (Specified Site)* action within the Discovery Profile portlet.

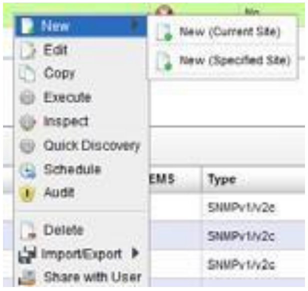

These devices are visible in the master site, but are assigned, by default, to the selected tenant site. Other tenant sites do not see them.

Supported Portlets

The following portlets are Multitenancy-supported. If a portlet does not appear on this list, it does not support Multitenancy:

- Actions
- Alarms
- Audit
- Audit Trails
- Authentication
- Contacts
- Customers
- Dashboard Views
- Discovery Profiles
- Event History
- Event Processing Rules
- Hierarchical View/Hierarchical View Manager
- Links - but only if the portlet is on the same page as resources
- Location
- Map View
- Policies
- Policy Templates
- Report Templates
- Reports
- Resources (including Ports, Interfaces)
- Schedules
- Service Templates
- Services
- Top N Portlets (not the RTFA Portlets though)/Top Problem Nodes
- Topology

NOTE:

For the Dashboard Views portlet, you must make a dashboard for performance monitoring on the tenant site if you want it to appear on the tenant site. Reconfiguring monitors makes dashboards misbehave on tenant sites.

Site Management

This portlet configures customer sites and organizations, including an administrative user for customer sites who can configure additional user accounts within the site. Optionally an organization can have its own website customize-able with non-standard graphics.

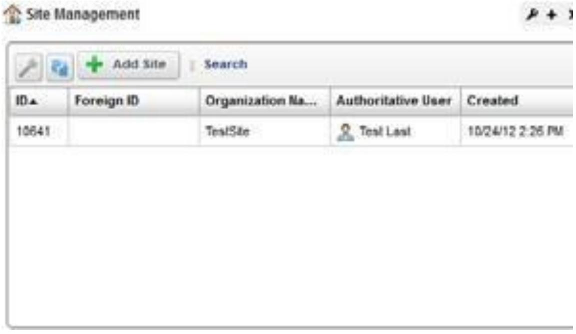

ID▲	Foreign ID	Organization Na...	Authoritative User	Created
10641		TestSite	Test Last	10/24/12 2:26 PM

NOTE:

Rather than the contents management tasks typical for other portlets, the wrench icon in this portlet opens managing restrictions for those logging into the Multitenant environment. See Login Restrictions on page 617 for more about how those work.

The search function can locate a site based on the *Organization Name*, *Foreign Id*, *Authorized User* or the *ID*. This portlet offers the following options:

Add Site — Click the *Add Site* button at the top of the portlet to create a new profile. This opens the Site Management Editor with a few more options than you might see if you right click to *Edit* an existing profile.

After you have created a site, you can see the following in its right-click menu:

Edit — Edit an existing Profile with the Site Management Editor .

Access Profile — Opens the AP Editor, described in AP Editor/AP Template Editor on page 623. Use the selection pick list at the top of the editor to associate an Access Profile Template with the selected site.

Tomake the Access Profilecustomerororganization-specific, youcanaugment the permitted functionality in the selected template by adding shareable resources. Click the *Add* button.

Existing permissions from any selected Access Profile Template do not appear in the permission type totals on the left, nor are they eliminated from the available permissions that appear after you click *Add*.

A Site Admin can choose to hide some of a tenant site's pages from a Site Member through page permissions. The page permissions do not work, however, if the tenant site comes from a site template.

Go to Site [New Window/Tab] —Opens the Customer's site in a new window or tab in your browser.

Delete — Deletes the selected customer.

The expanded and summary portlets are the same. Columns in both include an automatically-provided *ID* for each customer, the *Organization Name*, *Authoritative User* (the administrator for the customer or organization) and *Created* which records the date the customer organization was created. See Portal > Sites/Site Templates in Control Panel below for more about site management capabilities.

Login Restrictions

The site management portlet lets you restrict access to configured network domains. Select the configuration icon (the wrench) which opens the Global Site Settings dialog.

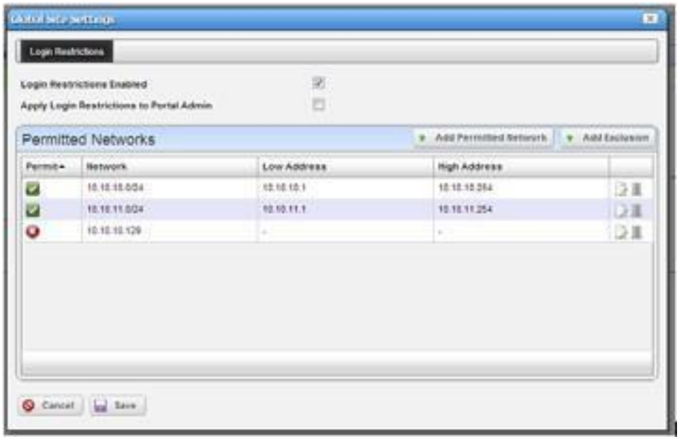

Login Restrictions Enabled ☑

Apply Login Restrictions to Portal Admin □

Permitted Networks

Permit	Network	Low Address	High Address
✓	10.10.10.0/24	10.10.10.1	10.10.10.254
✓	10.10.11.0/24	10.10.11.1	10.10.11.254
✗	10.10.10.129	-	-

Here the administrator can add networks that the primary site’s central domain users can login from, or exclusions of things like a proxy server within one of the permitted networks which allows external access to the web server. When attempting to login from an IP address other than those permitted a message appears saying *Login is restricted from your current IP [IP Address]*.

Notice that you must check *Login Restrictions Enabled* to begin restricting access. When you check that, global portal users can only log in from defined, permitted networks. Youcan also elect to *Apply Login Restrictions to Portal Admin* with that checkbox, too.

Portal > Sites/Site Templates in Control Panel

In addition to the Site Management portlet, authorized users can manage sites from the Control Panel. See the following:

- Portal > Sites
- Portal > Site Templates
- Click " + Add page" on the grey bar in the middle of the screen, fill in data and click Add page at the bottom

△ CAUTION:

Site creation must occur through Cruz’s Site Management portlet. If you bypass this portlet you will not create all the data that let Cruz manage the sites.

Portal > Sites

A list of sites created in Site Management Editor appears in this screen, and clicking the *Actions* button to the right of the listed site lets you manage the page configuration and user membership for tenant sites as well as the main site, if your login is a user authorized to access the main site.

Click the link that names a site to see its logo appear on the Control Panel page, and to select any site template (the template selected for Public Pages also appears for Private Pages). Youcan even check *Enable propagation of changes from the site template* to propagate changes from that template to the site itself.

You can also *Deactivate* or *Delete* a site listed there. These functions supplement the Site Management portlet and Site Management Editor described in this document.

Portal > SiteTemplates

Multitenancy uses Site Templatesconfigured in Control Panel to configure new tenant sites. You can select a Site Template or Page Template when you create a new tenant site. The following describes how to create a new Site Template.

How To:

Make A Site Template

1 Navigate to Control Panel

2 Click the *Portal > Site Templates* link.

3 Click *Add* to create a new site template. Notice that, in addition to configuring the Site Template’s Name and Description, you can check to make a template *Active*, and to *Allow Site Administrators to Modify the Pages Associated with this Site Template* at the bottom of the *New* screen.

4 *Save* the new template.

5 Click the *Actions* button to the right of listed templates to *Edit*, *Manage Pages*, *View Pages*, configure portal *Permissions* or *Delete* a selected template.

The *Permissions* configurable here are for the portal, not Redcell permissions. Use Portal > Roles to configure those. See Access Profile Templates on page 622 for more about configuring them on multitenant types.

6 Click *Manage Pages* to see the tenant site's page setup in tree form on the left, and more editor options on the right and center (for example, alter the look and feel, *Logo* and so on).

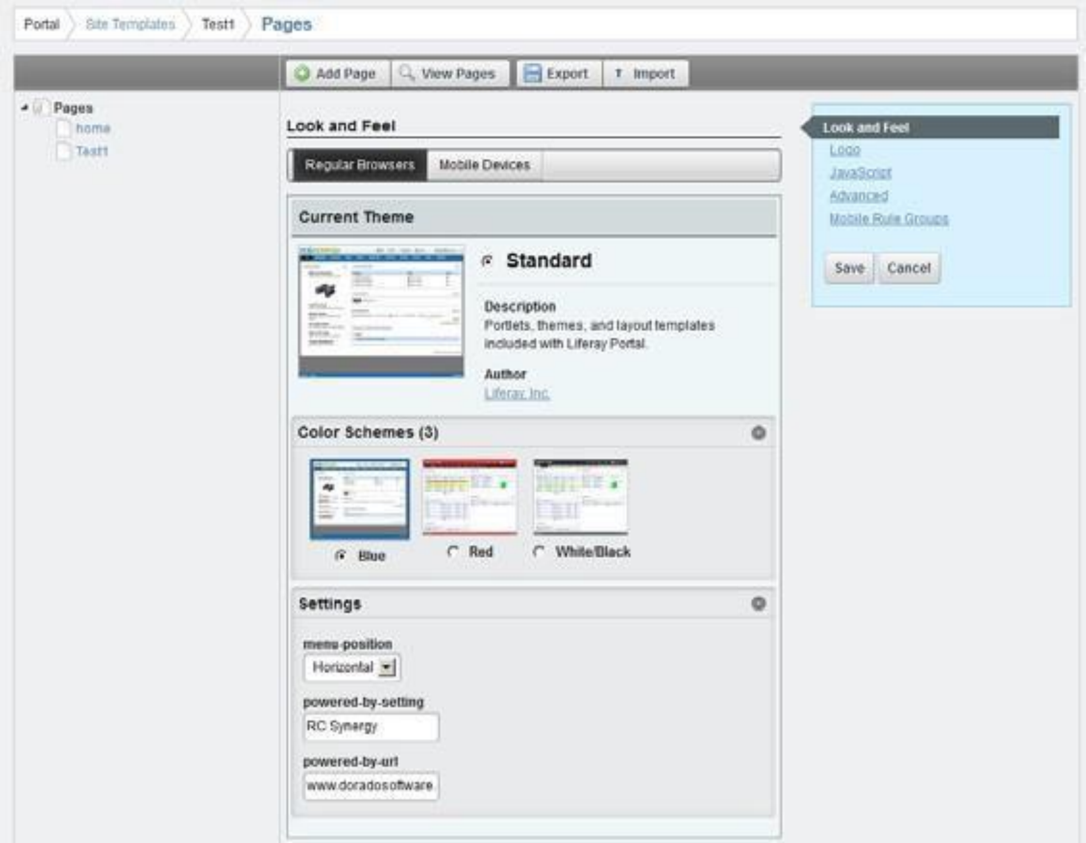

● Standard

Description
Portlets, themes, and layout templates included with Liferay Portal.

Author
Liferay Inc.

menu-position: Horizontal

powered-by-setting: RC Synergy

powered-by-url: www.doradosoftware

7 Click *Add Page* to create a new page. By default, new pages appear below the root node on the tree at the left, but you can drag and drop them to any new location.

8 After you configure the page layout(s) as you like, including the Logo, and color scheme, you can preview them by clicking the *View Pages* button at the top of the editor screen . This opens the new configuration in a new browser tab.

9 To configure portlets on these pages, simply use the *Add* > *Applications* menu item. When you re-open the Site Template, the portlets appear as you have configured them.

10 Notice that you can also *Import* a variety of settings from an exported LAR file with the button at the top right of the screen if you want pages that nearly duplicate each other. Notice also that you can *Export/Import* these Site Templates with the buttons at the top of the editor.

11 Click a page in the tree on the left to expose editing capabilities for that page, including its *Look and Feel* and *Layout* (columns) with links that appear on the right. Notice that you can also *Copy Portlets from Page* to copy the portlet setup from an already-configured page in this Site Template.

How To:

Add a new vertical menu page

1. Go to control panel > Cruz menu section > Site pages
2. Click "+ Add page" on the grey bar in the middle of the screen, fill in data and click Add page at the bottom
3. Page is added at the bottom of the displayed Private Pages hierarchy tree. Drag the menu item to a desired location if necessary.
4. Click save in the details panel on left.

The page is now added in the vertical menu.

Portal > PageTemplates

This lets you edit a single page template as you would several pages in Portal > Site Templates .

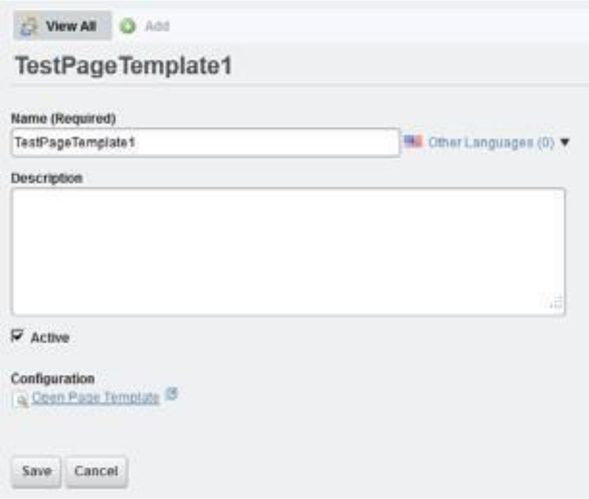

TestPageTemplate1

Name (Required): TestPageTemplate1

Description:

☑ Active

Configuration

Click the *Open Page Template* link at the bottom of this screen to see the page in a separate tab or window and add the desired portlets.

Site ManagementEditor

This editor lets you configure customers, site administrators and their organizations.

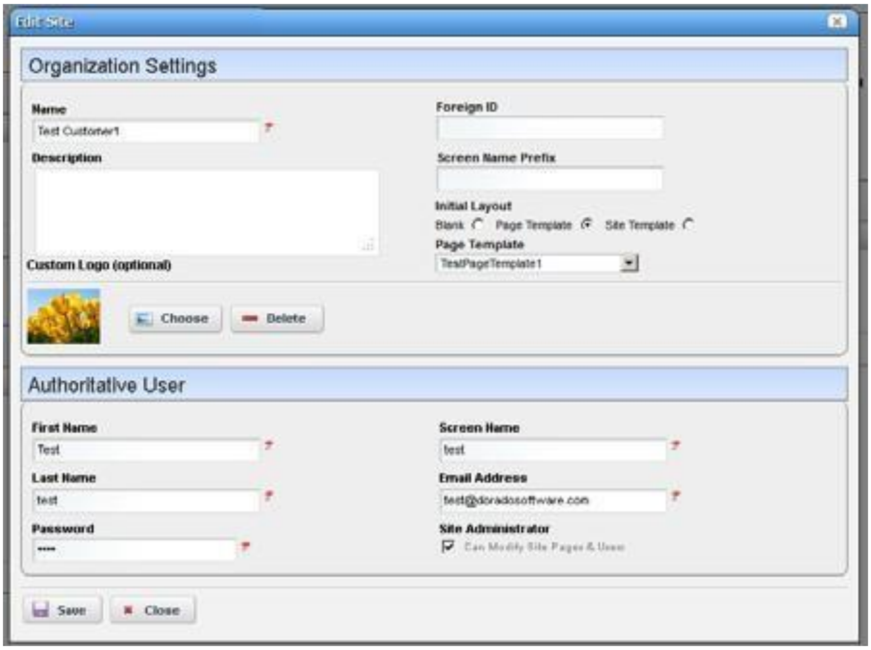

Edit Site

Organization Settings

Name: Test Customer1

Description:

Custom Logo (optional)

Foreign ID:

Screen Name Prefix:

Initial Layout: Blank ○ Page Template ● Site Template ○

Page Template: TestPageTemplate1

Authoritative User

First Name: Test

Last Name: test

Password: ••••

Screen Name: test

Email Address: test@doradosoftware.com

Site Administrator: ☑ Can Modify Site Pages & Users

It has the following fields and labels:

Organization Settings

Name — A unique name for the customer’s organization.

Foreign ID — An optional identifier for the customer’s organization.

Screen Name Prefix—Enter a screen prefix for the customer. The prefix must be unique within the system. The prefix is prepended to all user accounts to insure no accounts conflict across organizations. So, if you create a customer with a prefix of DS, and his screen name is Mark then the user’s login is DS-Mark. This is automated in Site Management Editor, but you must manually enter the prefix for users created in Control Panel.

By default, this field is required. If you do not want to use screen name prefixes, then edit the `oware/synergy/conf/server-overrides.properties` file and uncomment the following property:

```
#site.screen.name.prefix.required=false
```

Description — Any text description for this customer.

Custom Logo (optional)—Select a graphic. The subsequent selection screen lets you select a file, and recommends it be transparent, 50 pixels in height, and proportional in width. Once you have successfully selected and uploaded a logo, a preview appears below this label.

Initial Layout — Select from the radio buttons. This determines the configuration for the initial screen provisioned for the customer’s site. See Portal > Site Templates on page 618 and Click " + Add page" on the grey bar in the middle of the screen, fill in data and click Add page at the bottom on page 620 for instructions about configuring these. To take effect, these must exist before the Customer exists.

Authoritative User

These fields configure the tenant site’s administrator. The user information entered here automatically configures a user in Cruz.

First/Last Name — The name of the administrator.

Screen Name — The screen name of the administrator.

Email Address—The e-mail address of the administrator.

Password— The administrator’s password.

Site Administrator —This grants the user site administration privileges. These privileges let administrators add or remove users within the site and configure pages and layouts.

When an organization's administrator assigns permissions to their users and/or user groups they are constrained by their own access permissions granted by the Multitenant site provider, ensuring that such administrators cannot grant more permissions then they have. See How to: Configure User Site Access on page 625 for information about granting users access to more than one site.

When you try to log in as a multitenant user, even if you do not type it as part of the login ID, Cruz automatically prepends the Screen Name Prefix if you create the user in the Site Management Editor. If you make other users for the tenant site manually in Control Panel, you must manually add the prefix to assign them to the correct site. In either case, you must type that prefix as part of the login ID.

Click *Save* to preserve your edits, or *Close* to abandon them.

Access Profile Templates

Configure data access for Multitenant Service Providers (MSPs) through an Access Profile (AP), configured in these portlets as described in AP Editor/AP Template Editor on page 623, or more accurately, in that AP Template plus whatever customization you configure for Site Management when you assign it an AP.

An AP describes which entities within the system a site can access. When enforcing AP-based access, Cruz identifies applicable entities based upon the Site ID. The fundamental difference between this approach and Site ID filtering is its ab ility to access an entity from multiple sites. For example, multiple customer sites can have access to an AP template.

Thisapproach, along with functional permissions, lets MSPscontrol exactlywhat their customers can do and see.

Forexample, an MSP offers Gold, Silver and Bronze tiers of service. At the Gold level customers have access to 12 pre-defined reports and they can create their own reports from existing report templates. At the Silver level customers can only access the 12 per-defined reports and at Bronze level the customer can only access three reports, a subset of the 12.

To do this within Cruz, the MSP would first create the necessary report templates and report definitions. Assume there are eight templates and 12 report definitions. At the Gold level the MSP would provide access to the eight templates and the 12 report definitions. The MSP could also give gold customers the functional permission to create a report definition. At the Silver level the MSP gives the customer access to the 12 reports but no functional permission to create new reports. Since customer at the silver tier cannot create report definitions, they also cannot access report templates. Finally, at the Bronze level the MSP provides access to the three specific reportsavailable at that service level. In all cases, only target de vices assigned to the respective sites appear in the report(s).

This portlet lets you configure the kind of access available to customers configured in the Site Management portlet.

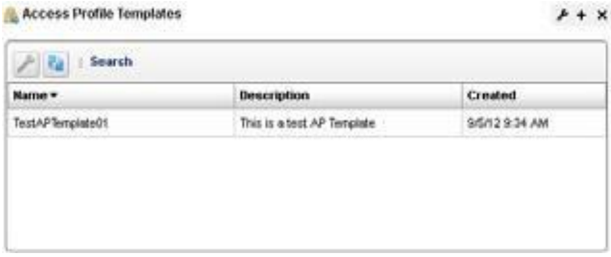

Name ▾	Description	Created
TestAPTemplate01	This is a test AP Template	9/5/12 9:34 AM

This displays a profile’s *Name*, *Description*, and *Created* date by default. Youcan alsoadd a column to display its *Updated* date. The expanded portlet displays the same columns and a Reference Tree snap-in showing connections to the selected Profile.

This portlet’s right-click menu has the following items:

New/Edit — Create a new or edit an existing template with the AP Editor/AP Template Editor . Create Access Profiles when you assign them to sites by right clicking in the Site Management portlet.

Delete — Remove a selected profile.

AP Editor/AP Template Editor

These editors let you create and modify Access Profiles for single sites and Templates that may apply to several sites. To standardize access pro files across multiple customers, create a template. Right-click a site in the Site Management portlet to see the Access Profile Editor, where you can select templates. Right-click in the Access Profile Templates portlet and select *New* to create a template.

These profiles configure access to Cruz resources for customers configured in the Site Management portlet. Use that portlet’s right-click menu to associate Profiles with Customers.

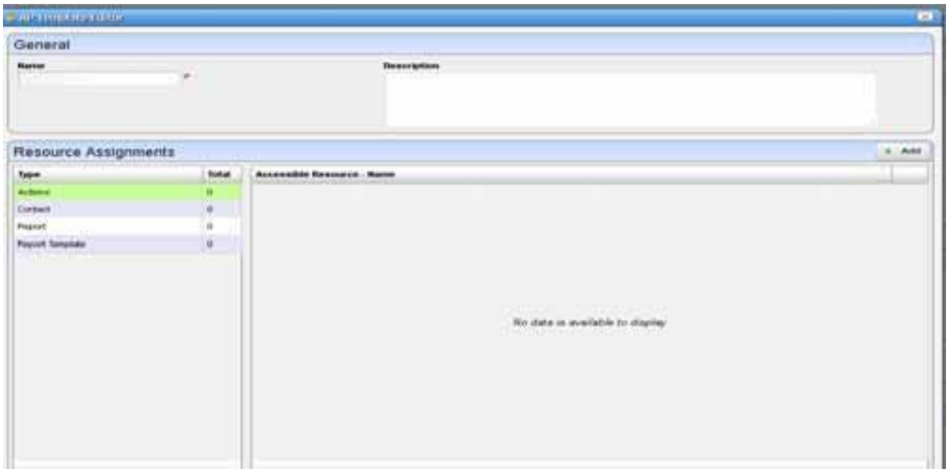

Type	Total
Actions	0
Contact	0
Report	0
Report Template	0

Accessible Resource - Name

No data is available to display

The editor screen contains the following fields:

Template Association

This pick list contains the templates previously configured in Access Profile Templates portlet. This Template Association panel does not appear in the AP Template Editor.

General

Name — An identifier for the Access Profile.

Description—A text description for this Profile.

Resource Assignments

This panel displays the Type of resource on the left. To select resources of that Type, click to select it, then click the *Add* button over the right panel. A selector of available types of access to resources appears. Select the desired ones for the Profile, and they appear listed below the *Add* button. Use the icon to the right of the permission to remove it from a profile. The types of resources constrained by these templates include the following:

- Actions
- Contacts
- Reports
- Report Templates
- Service Policy Templates
- Service Templates

Click *Save* to preserve your edits, or *Close* to abandon them.

The access permissions configured here dictate what users can do and can further be limited by Cruz's functional permissions.

This AP applies to users created in the base domain as well as those created within an organization. Those capabilities include assigning permissions directly to an individual user or to a user group. Any user within the group inherits group-assignedpermissions.

User Site Access

By default, users can access only one Multitenant site. If your network has users who need to see more than one site, you can configure such User Site Access in this portlet.

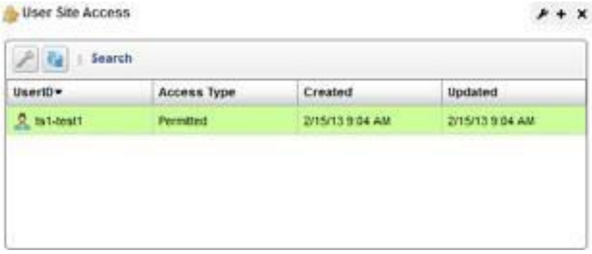

UserID	Access Type	Created	Updated
ts1-test1	Permitted	2/15/13 9:04 AM	2/15/13 9:04 AM

Right clicking in this portlet offers the following menu items:

New — Create a new mapping of a user to a set of sites. When you click this, a screen appears where you can select a user already configured in the system. The User Site Access Policy editor appears after you have selected a user. Configure site access in that editor.

Edit — This opens the User Site Access Policy editor for the selected, already configured UserID. Editing existing policy may take up to 10 minutes to take effect. If user wish to see the effect right away, delete and recreate the policy.

Delete — Delete the user site access configuration selected.

User Site AccessPolicy

This editor assigns Multitenant sites to users.

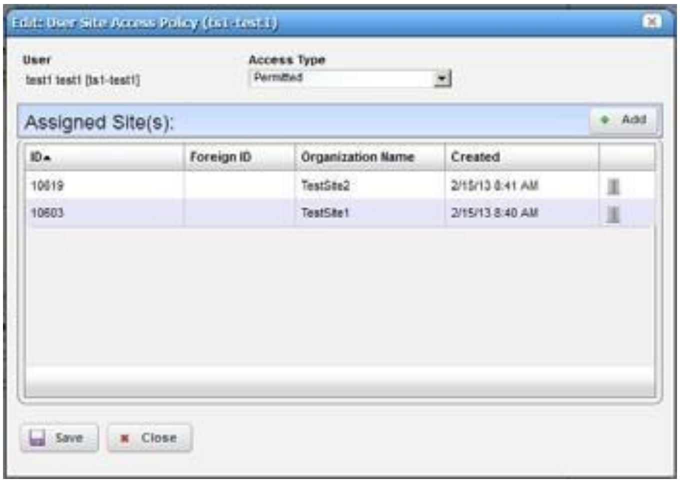

Edit: User Site Access Policy (ts1-test1)

User
test1 test1 [ts1-test1]

Access Type
Permitted

Assigned Site(s):

ID▲	Foreign ID	Organization Name	Created	
10619		TestSite2	2/15/13 8:41 AM	
10603		TestSite1	2/15/13 8:40 AM	

This screen offers the following selections for User Site Access:

Access Type—Select from *Permitted* (allow access to the listed sites) or *Restricted* (deny access to the listed sites).

Assigned Site(s) —Click *Add* to select sites for the selected user's access/restriction. Click the icon on the right of a listed site to delete it.

Click *Save* to preserve a User Site Access Policy, or *Close* to abandon your edits.

How To:

Configure User Site Access

Follow these steps to configure User Site Access:

1 If you have not previously configured Multitenant sites, you must do so before configuring User Site Access. See How to: Create a Multitenancy Environment on page 613 for step-by-step instructions about how to do that.

2 Create any user(s) for whom you want to configure access.

3 Configure those user(s) permissions. See Configure Resource Level Permissions in this guide.

4 Right-click in the User Site Access portlet and select *New.*

5 In the subsequent screen, select the user whose Multitenant site access you want to configure.

6 The User Site Access Policy editor appears.

7 In the *Access Type* pick list, select whether you want to grant the selected user access to the sites you pick (*Permitted*), or you want to deny their access to the sites you select (*Restricted*).

8 Click the *Add* button with the green plus.

9 In the subsequent screen, select the site(s) your user can access or from which your user is to be excluded.

10 Click *Save*. The user you have configured should have access to (or be restricted from) the Multitenant sites you have configured.

Constraining Data Access

Constraining data access for a site involves two primary mechanisms: Site ID Filtering , also referred to as Domain Id filtering, and the kind of filtering provided by Access Profile Templates . The following diagrams how multitenancy can work.

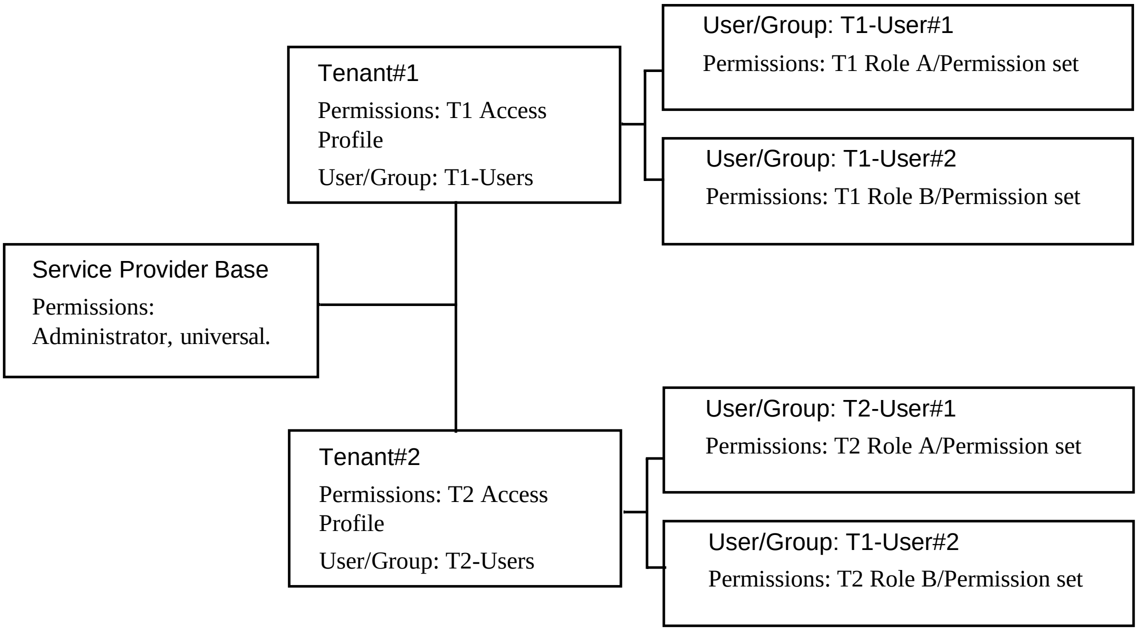
User/Group: T1-User#1
Permissions: T1 Role A/Permission set
Tenant#1
Permissions: T1 Access Profile
User/Group: T1-Users
User/Group: T1-User#2
Permissions: T1 Role B/Permission set
Service Provider Base
Permissions: Administrator, universal.
User/Group: T2-User#1
Permissions: T2 Role A/Permission set
Tenant#2
Permissions: T2 Access Profile
User/Group: T2-Users
User/Group: T1-User#2
Permissions: T2 Role B/Permission set

Manage > Domain Access [Resources]

If you right-click a Managed Resource, you can select Manage > Domain Access to configure different tenant domains’ access to the selected device.

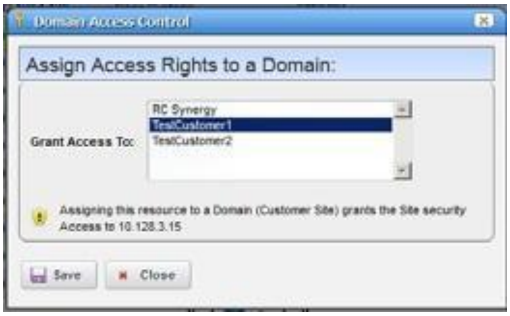

Assign Access Rights to a Domain:

Grant Access To: RC Synergy, TestCustomer1, TestCustomer2

Assigning this resource to a Domain (Customer Site) grants the Site security Access to 10.128.3.15

Click the domain(s) (besides the master domain) where you want to have access to the device and then click *Save*.

Site ID Filtering

Site Management/organization site includeaninternal identifier. Cruzautomaticallyassociates that Site ID with inventory entities created within the context of the site. That Site ID then appears in any database filter whenever users retrieve data within the context of a site, ensuring that the results come only from the site in context.

The one exception to this rule of filtering is that, from the MSP's perspective, no Site ID limits the visibledata set, regardlessofthe organizational site. Inventory entities onlyexist withina singlesite since they hold a single site ID.

The following entities use this approach:

- Resources
- Services
- Policies
- Authentication
- Discovery Profiles
- Locations
- Contacts
- Alarms
- Events
- Event Processing Rules
- Resource Groups
- Audit History
- Schedules
- Performance Dashboards

13

VLAN Visualization

VLAN Visualization providesyou the ability to collect, display, and report on VLAN related data throughout all the managed devices in your network. The collected data can be consumed in various ways, including table, topographical, and report-based formats.

Collecting the Data

The first step in utilizing the VLAN Visualization feature is to collect the data that is needed. This can be done in one of two ways.

Get VLAN Data Executed as an Action

In the first method, the "Get VLAN Data" action can be run from the Actions portlet (which is found, by default, under the "Resources" page:

When executing the "Get VLAN Data" action in this method, you will be presented with a list of managed devices. You can either single select, or multi-select.

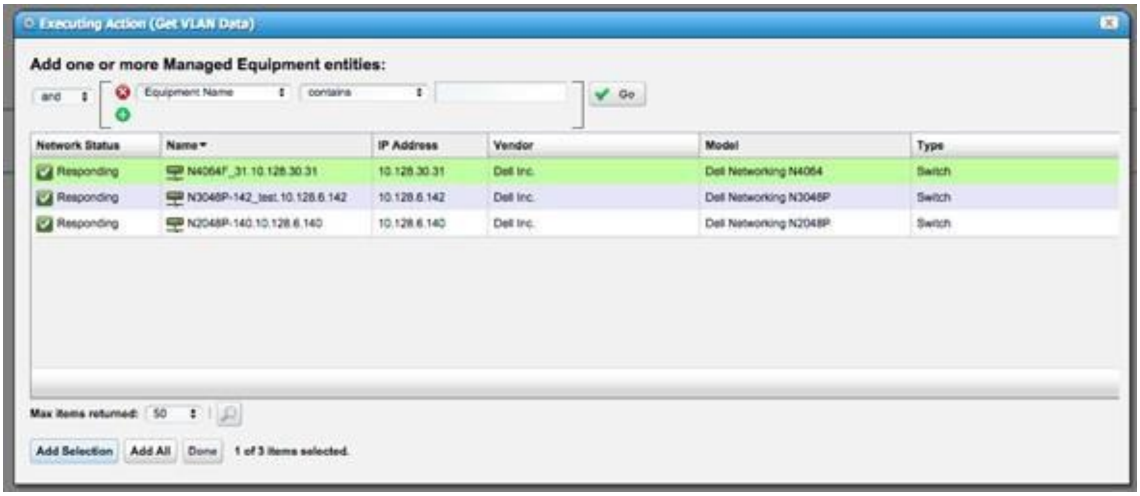

Executing Action (Get VLAN Data)

Add one or more Managed Equipment entities:

and | Equipment Name | contains

Network Status	Name	IP Address	Vendor	Model	Type
☑ Responding	N4064F_31.10.128.30.31	10.128.30.31	Dell Inc.	Dell Networking N4064	Switch
☑ Responding	N3048P-142_test.10.128.6.142	10.128.6.142	Dell Inc.	Dell Networking N3048P	Switch
☑ Responding	N2048P-140.10.128.6.140	10.128.6.140	Dell Inc.	Dell Networking N2048P	Switch

Max items returned: 50

1 of 3 items selected.

After you have finished selecting the devices you wish to run the VLAN data collection against, select the "Execute" option, and the data will be collected.

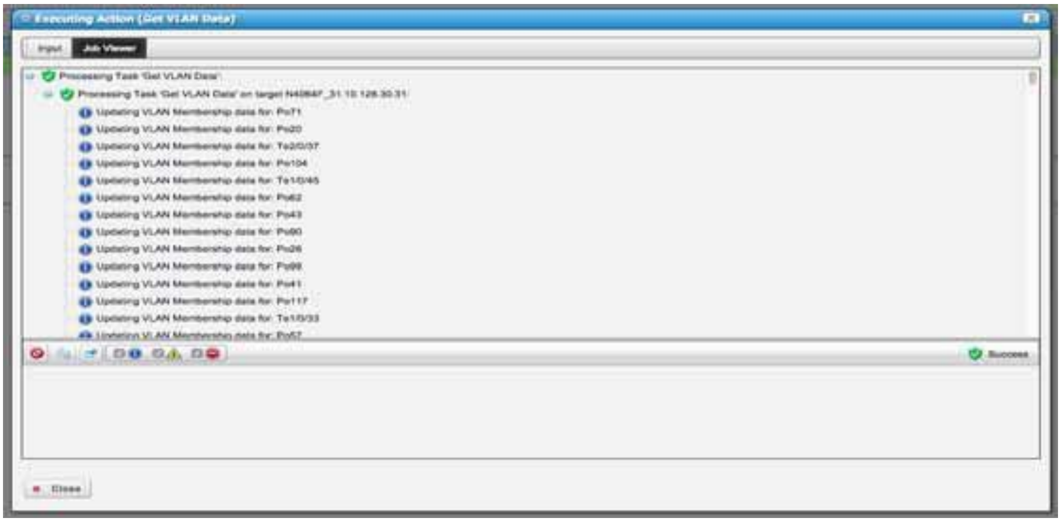

Executing Action (Get VLAN Data)

```
Processing Task 'Get VLAN Data'
  Processing Task 'Get VLAN Data' on target N4064F_31.10.128.30.31
    Updating VLAN Membership data for: Po71
    Updating VLAN Membership data for: Po20
    Updating VLAN Membership data for: Te2/0/37
    Updating VLAN Membership data for: Po104
    Updating VLAN Membership data for: Te1/0/45
    Updating VLAN Membership data for: Po62
    Updating VLAN Membership data for: Po43
    Updating VLAN Membership data for: Po90
    Updating VLAN Membership data for: Po26
    Updating VLAN Membership data for: Po99
    Updating VLAN Membership data for: Po41
    Updating VLAN Membership data for: Po117
    Updating VLAN Membership data for: Te1/0/33
    Updating VLAN Membership data for: Po57
```

Success

Adding "Get VLAN Data" Action to a Discovery Profile

The second way you can have the "Get VLAN Data" action executed against a device, is by including it in the list of actions that are to be executed in your Discovery Profile, when discovering a device.

NOTE:

VLAN Data is collected and stored in the database at the time collection is run. If the VLAN configuration on a box has changed, the "Get VLAN Data" action will need to be run against that device again. This can be accomplished by either manually executing the action, or by scheduling the action to be run at a specified interval, as can be done with all other actions.

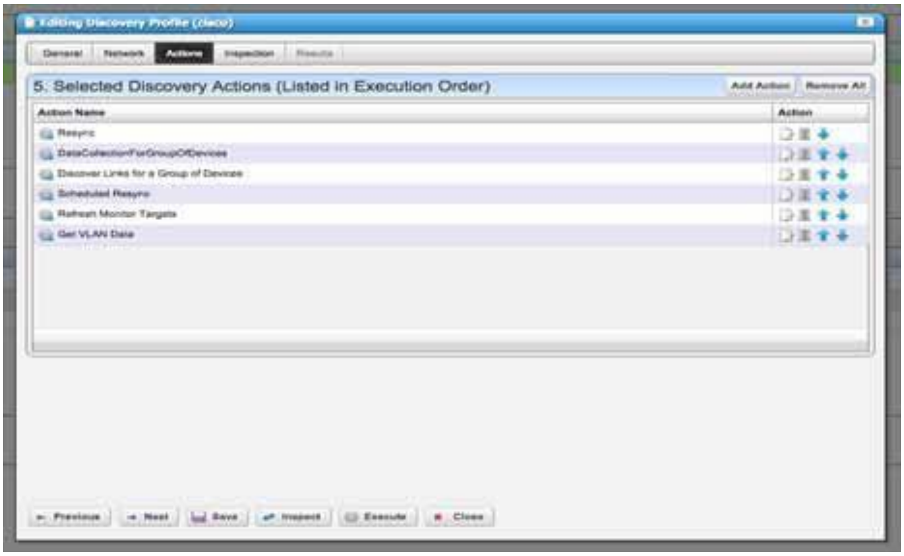

5. Selected Discovery Actions (Listed in Execution Order)

Action Name	Action
Resync	
DataCollectionForGroupOfDevices	
Discover Links for a Group of Devices	
Scheduled Resync	
Refresh Monitor Targets	
Get VLAN Data	

VLAN Visualization Portlets

After having executed VLAN data collection, you are now ready to work with the data that has been collected.

Adding VLAN Related Portlets to Your View(s)

VLAN Visualization consists of four portlets (VLAN Domains, VLAN Domain Assignment, VLANs, and VLAN Memberships) that can be added to your view(s) in any manner that you see fit. You may create a new page or add these to an existing page from Add- > Applications in the menu.

VLANs Portlet

The "VLANs" portlet displays the current list of all VLANs that known on your network. Each entry in the table will have data for:

- "VLAN Domain - The VLAN Domain that this VLAN is considered to be a part of (Please see the "Working with VLAN Domains" section for more information regarding the notion of "VLAN Domains".)
- "VLAN ID - The ID of the VLAN.
- "Name - The name of the VLAN that was captured from the device during data collection.

- "Description - A user settable description that can be used to store information that the user would like to associate with this VLAN entry. This can be edited by right clicking on the entry and selecting "Edit".

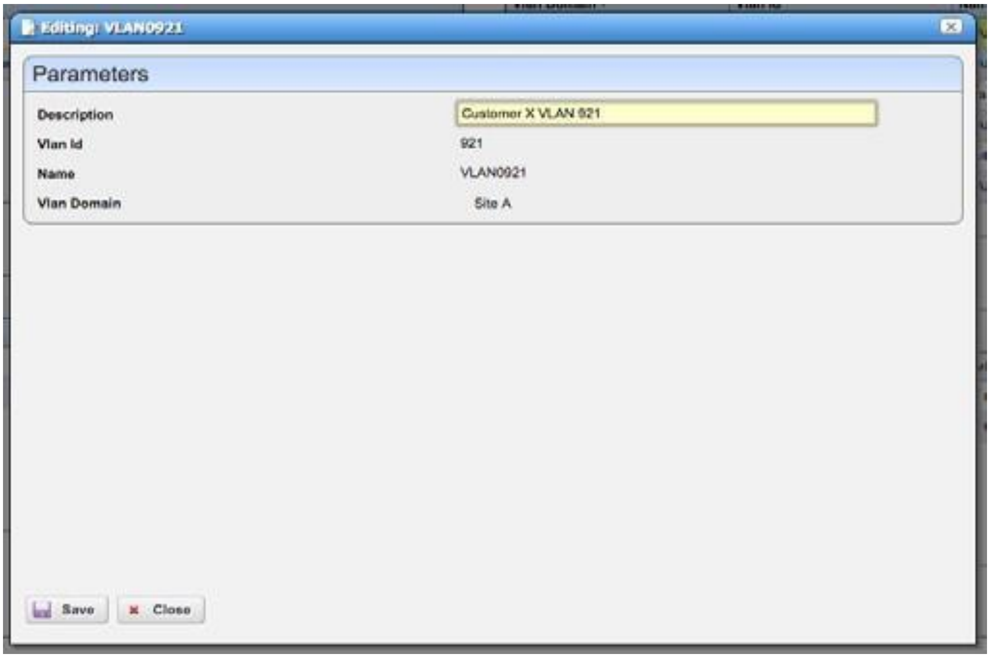

Description	Customer X VLAN 921
Vlan Id	921
Name	VLAN0921
Vlan Domain	Site A

- "Member Count - The current number of members that known (for devices present in inventory).

VLAN Memberships

The "VLAN Memberships" portlet displays a list of all VLAN Memberships that are known on your network. Each entry in the table will have data for:

- "VLAN - the VLAN that this membership is associated with.
- "VLAN DOMAIN - the "VLAN Domain" of the VLAN that this membership is associated with.
- "Device - the device that this membership is from.
- "Endpoint - The port/interface that this membership is from
- "Switchport Mode - Whether this port/interface is set to Trunk or Access.
- "Tagged - Whether this endpoint is Tagged or Untagged in this membership

VLAN Domains Portlet

The "VLAN Domains" portlet displays all data for the "VLAN Domains" that have been created.

A "VLAN Domain" is a concept that we have included in order to allow the user to specify that equipment being managed may be from different labs, different customer sites, etc. that are not in the same network. For example, if you have two groups of managed devices , and they are on entirely different networks, but they both have a VLAN30, you will need a way to inform the software that the VLAN30 on the first group of devices is different from the VLAN30 on the second group of devices. The VLAN Domain concept provides you with the functionality to do this. (NOTE: Please see the "Working with VLAN Domains" section for more information on VLAN Domains.)

Each entry in the VLAN Domains Portlet will contain data for the following:

- Name - The Name of the VLAN Domain. These must be unique.
- Description - Description of the VLAN Domain.

VLAN Domain Assignments

The "VLAN Domain Assignments" portlet will display one entry for each assignment of a port, interface, device, etc. to a VLANDomain. (For moreinformation regard VLANDomain Assignment, pleasesee"Workingwith VLAN Domains").

Working with VLANDomains

As previously mentioned, due to the fact that when collecting VLAN data from a particular device, it is impossible to tell one "VLAN30" apart from a "VLAN30" that might be on a different device - one that is not networked to the first device in any way - the concept of a "VLAN Domain" has been introduced. By making use of these VLAN Domain Assignments, you can inform the software when one "VLAN30" is separate from another "VLAN30".

Creating a VLANDomain

The first step to making use of VLAN Domains, is to create one in "VLAN Domains" portlet. To initiate the creation process, right click anywhere within the table, and select the "New" option (Note: If you have not yet added the "VLAN Domains" portlet to a page, you must do so. Please see "Adding VLAN Related Portlets to Your View(s)" for more information.)

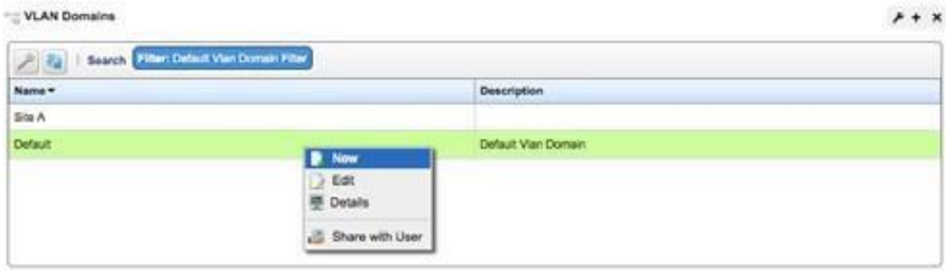

Name	Description
Site A	
Default	Default Vlan Domain

After selecting this option, you will then be presented with a dialog prompting you to enter a name and description for your new VLAN Domain.

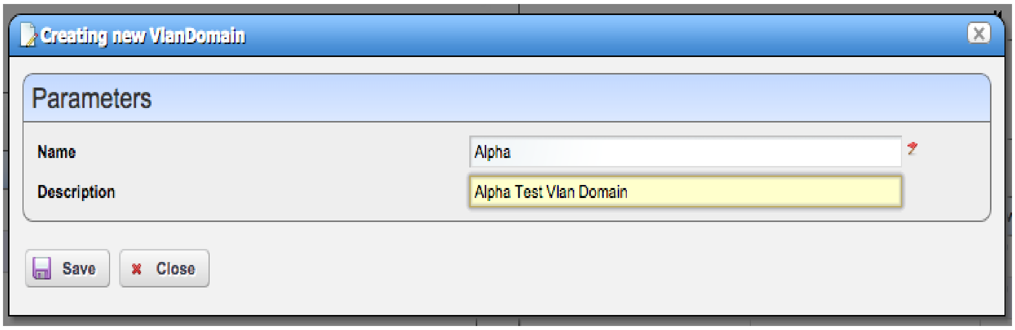

Name: Alpha

Description: Alpha Test Vlan Domain

Afterentering the information, click"Save", and your new VLAN Domain should nowexist within the system, and be ready for use.

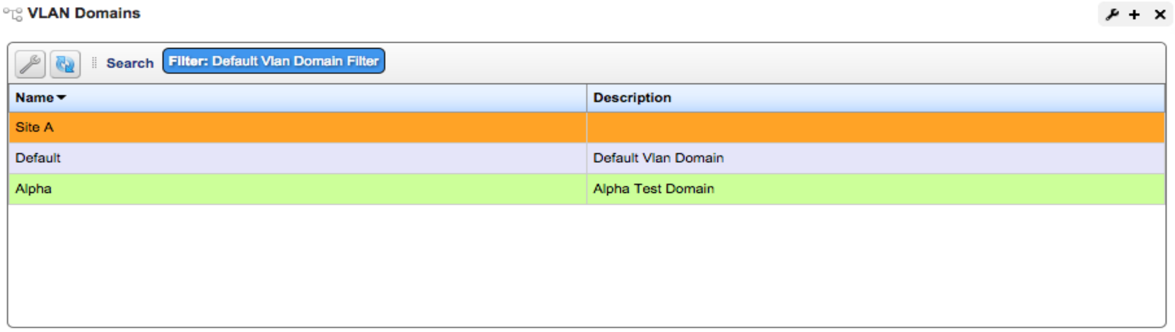

Name	Description
Site A	
Default	Default Vlan Domain
Alpha	Alpha Test Domain

Default VLAN Domain

A"Default VLAN Domain" is seeded automatically. It cannot be deleted, and any piece of equipment that is not specifically assigned to a different VLAN Domain will be assumed to be on this Default Domain.

Assigning to a VLAN Domain

Now that you have a created a VLAN Domain, you can assign equipment to it. Assignment can be performed at any level. Forexample, you can assign a specific port/interface to a VLAN Domain,in addition to being able to assign an entire card, or even an entire device.

Toinitiate the assignment process, you will select "Assign to VLAN Domain" from the actions portlet

Once you have located the action, select the "Execute" option to begin.

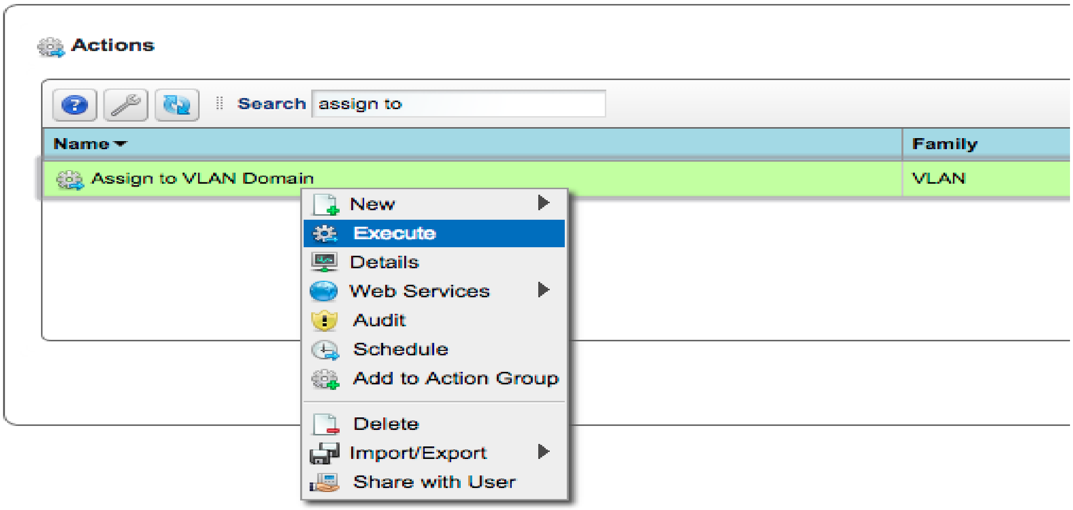

Search: assign to

Name	Family
Assign to VLAN Domain	VLAN

From the first dialog, you will be asked to select what piece of equipment you wish to assign. You can choose what type ofequipment (Card, Equipment and Subcomponents, Interface, Managed Equipment, or Port) byselecting the appropriate option from the upper most combo box. Also note that multi-selection is enabled here, so you can assign multiple objects in one action.

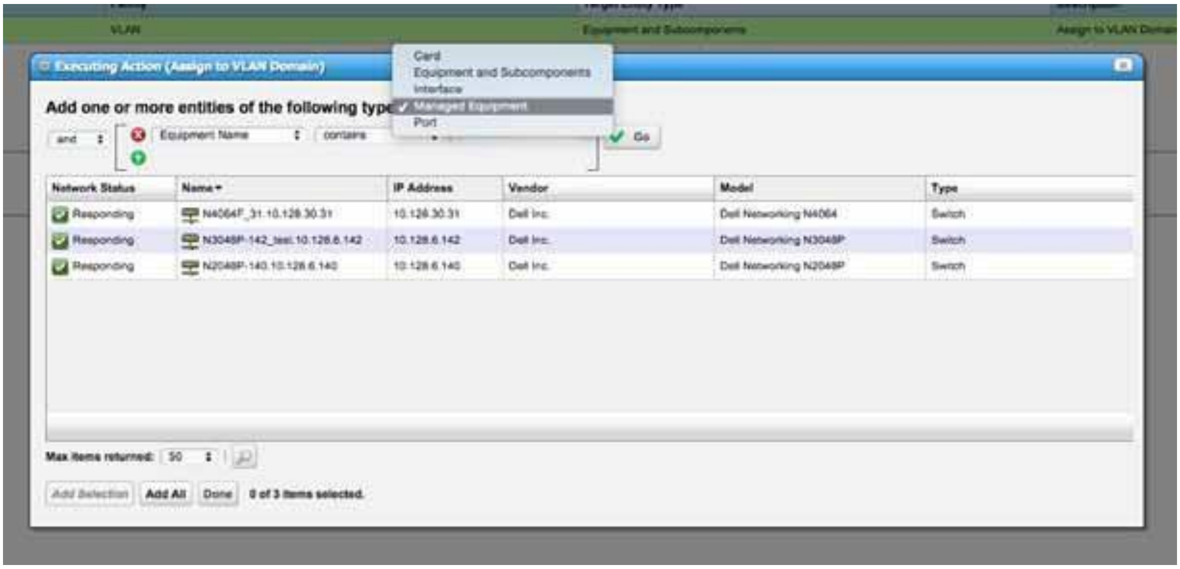

Add one or more entities of the following type

Network Status	Name ▾	IP Address	Vendor	Model	Type
✓ Responding	N4064F_31.10.128.30.31	10.128.30.31	Dell Inc.	Dell Networking N4064	Switch
✓ Responding	N3048P-142_test.10.128.6.142	10.128.6.142	Dell Inc.	Dell Networking N3048P	Switch
✓ Responding	N2048P-140.10.128.6.140	10.128.6.140	Dell Inc.	Dell Networking N2048P	Switch

0 of 3 items selected.

Once you havefinished addingselections, and clicked "Done", you will then bepresented witha dialog that willallow you to select which VLAN Domain you wish to assign your selected equipment to.

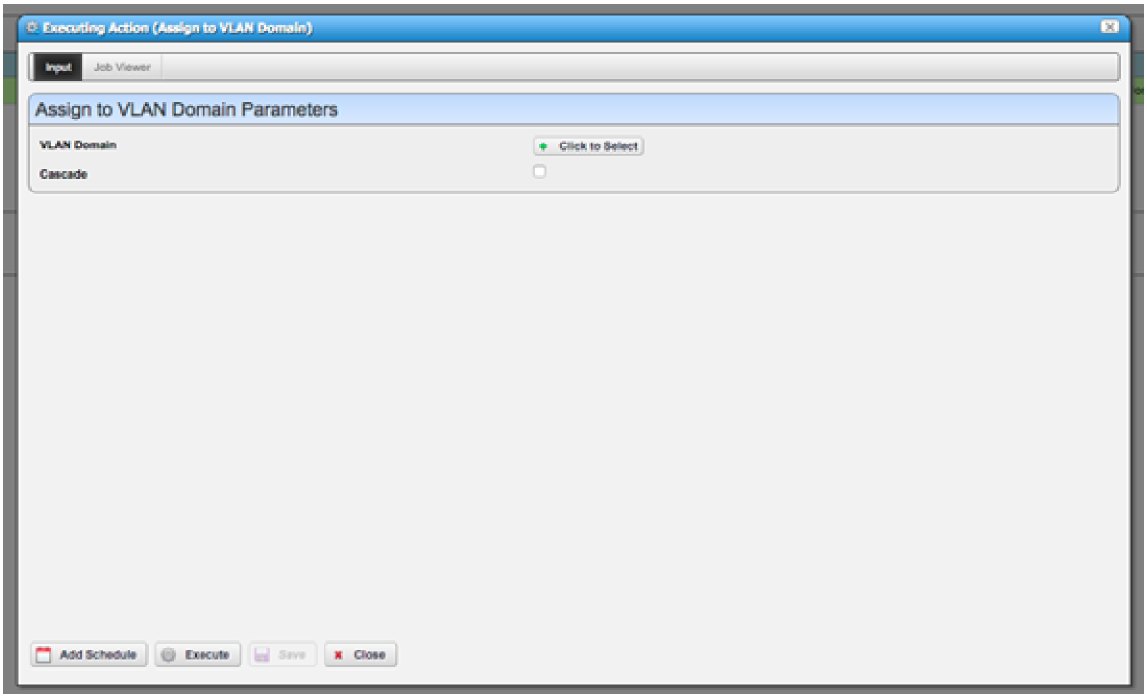

Assign to VLAN Domain Parameters

VLAN Domain: Click to Select

Cascade: □

Clicking the "Click to Select" button will present you with a list of all existing VLAN Domains to choose from.

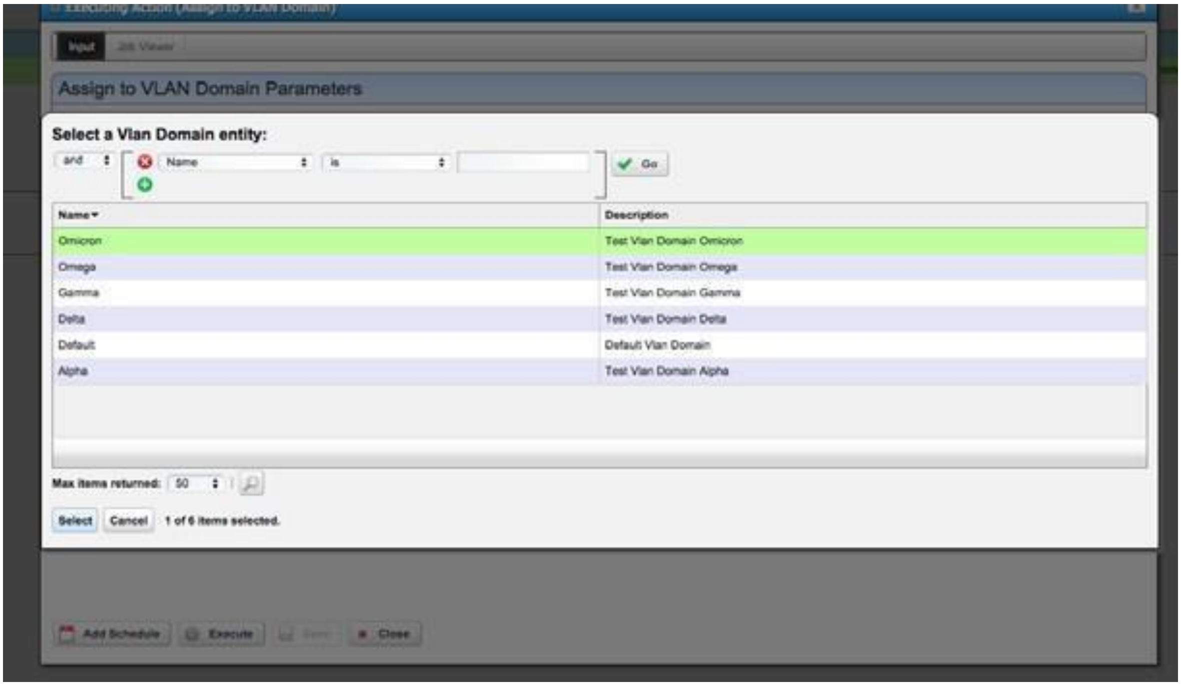

Assign to VLAN Domain Parameters

Select a Vlan Domain entity:

and | Name | is

Name	Description
Omicron	Test Vlan Domain Omicron
Omega	Test Vlan Domain Omega
Gamma	Test Vlan Domain Gamma
Delta	Test Vlan Domain Delta
Default	Default Vlan Domain
Alpha	Test Vlan Domain Alpha

Max items returned: 50

1 of 6 items selected.

After selecting your desired domain, and clicking the "Execute" button, the selected equipment will be assigned to your specified VLAN Domain, and VLAN Membership will bere-calculated.

Toillustrate this, let usconsider the followingexample. Assume you have two devices,"Device1", and "Device2". On both of them, VLANS 10, 20, and 30 exist. If you had run the "Get VLAN Data" action against Device1 and Device2, and ports/interfaces on Device1 and Device2 were tagged/untagged in those VLANs, you would find the following:

- VLANs 10, 20, and 30 would exist and entries for them would be visible in the "VLANs" portlet.
- As there were no VLAN Domain Assignments for anything on either device, each of those VLANs would be assumed to exist on the Default VLAN Domain. Thus, in the "VLANs" portlet, you would see entries for VLAN10, VLAN20, and VLAN30, and each of them would be "assigned" to the "Default" VLAN Domain.
- In the "VLAN Memberships" portlet, you would find membership entries for all VLAN ports/interfaces, on either device, and they would all be associated with the three VLANs mentioned above.

Nowlet's say that we haveassigned Device1 to VLAN Domain"Alpha" (A VLAN Domain that you had previously created). Upon assignment, "Get VLAN Data" will automatically be run against all affected equipment, and you would find the following:

- VLANs 10, 20, and 30 - **on the Default VLAN Domain** - would exist, and entries for them would be visible in the "VLANs" portlet.
- "VLANs 10, 20, and 30 - **on the "Alpha" VLAN Domain** - would exist,and entries for them would also be visible in the "VLANs" portlet.
- All VLANports/interfaceson Device1 - the devicethat wasassigned to the"Alpha" VLANDomain- would have VLAN Membership data visible in the "VLAN Memberships" portlet, and they would all be associated with the 10, 20, and 30 VLANs - **on the "Alpha" VLAN Domain**.
- All VLAN ports/interfaces on Device2 - a device that still has no VLAN Domain Assignments associated with it whatsoever - would have VLAN Membership data placing them in their appropriate VLANs- **on the"Default" VLAN Domain**.

VLAN Domain Inheritance

In order to simplify VLAN Domain assignment, the software incorporates the idea of child components of a device inheriting VLAN Domain assignment from their parent components. For example, you can assign a specific port/interface to a VLAN Domain, and only that port/interface will be considered to be assigned to it. If, however, you assign the card - that several ports are contained in - to a VLAN Domain, all ports on that card will also be assumed to be assigned to the same VLAN Domain that their parent card is assigned to. This inheritance is supported all the way up through the top level - i.e., the device itself.

NOTE:

It is possible to have multiple levels of assignment on any given device; however, please note that when calculating VLAN Membership for a given port/interface, only the first VLAN Assignment found in the parent/child tree will be used!

For example, if you have assigned a the top level device, "Device1", to VLAN Domain "Alpha", but you have also individually assigned a port on Device1, "portA", to VLAN Domain "Omega", when VLAN Membership is calculated for "portA", it will be assumed to be on the "Omega" VLAN Domain, despite the fact that the device it is contained in - Device1 - is assigned to VLAN Domain "Alpha"!

"Cascade" VLAN Domain Assignment

In order to simplify assignment of equipment to VLAN Domains, is the concept of "Cascade" assignment. The "Cascade" option can be selected when executing the "Assign to VLAN Domain" action.

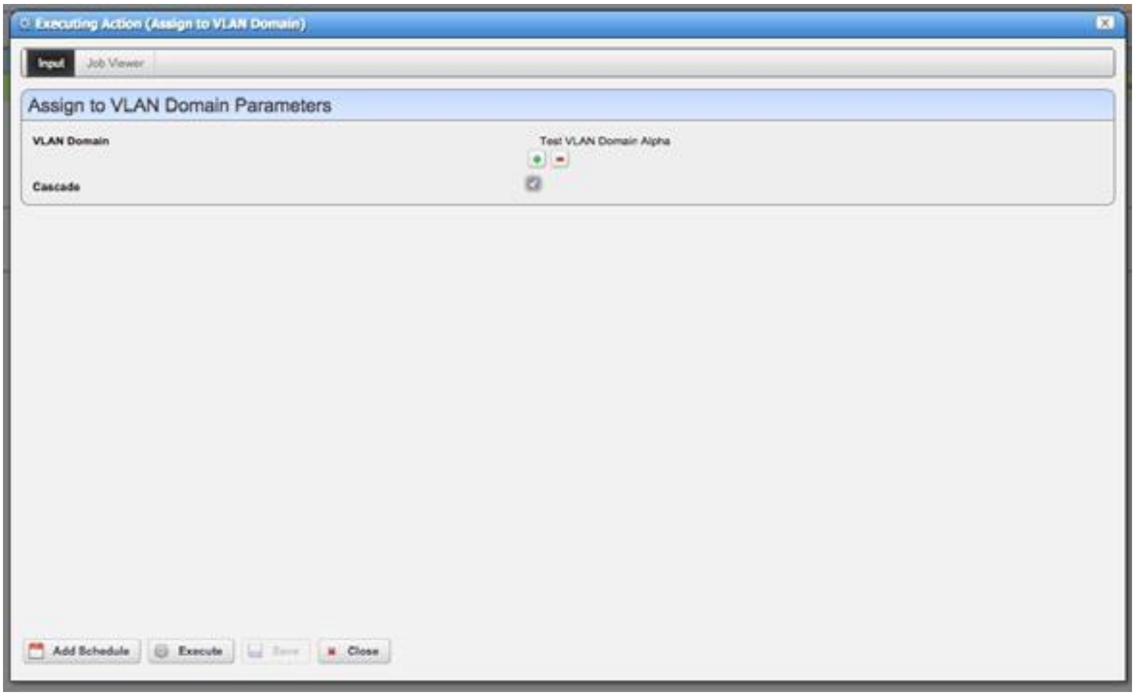

Assign to VLAN Domain Parameters

VLAN Domain: Test VLAN Domain Alpha

Cascade: ☑

When this option is selected, the VLAN Domain Assignment will cascade throughout all other VLAN ports on a given device, and all devices linked to that one. The way it cascades will depend on the initial point (port/interface, or device) that the cascade was started from.

If you have selected a port/interface to initiate the cascade, the following will occur:

- All VLANs that port/interface are members of will be "in play" for the cascade. Forexample, if it is an access port, and it is only on VLAN30, then only VLAN30 will be considered during the cascade. If, however, it is a trunk port, and it is a member of VLANs 10, 20, and 30, all three of those VLANs will be "in play".
- The selected port/interface will be assigned to the specified VLAN Domain.
- All other ports/interfaces, on the same device as the initially selected port/interface - **that share at least one VLAN with the initially selected port/interface** - will also be assigned to specified VLAN Domain.
- The cascade will then look for other devices that are linked to the device that the initially selected port/interface is on, and check to see if VLAN data (for any of the VLANs in play) is going over those links (VLAN data is assumed to be going over a link if both ends of the link have at least one shared VLAN).
- If a linked device is determined to be in the same VL AN, the cascade process will continue that device, and branch out from there.

If you have selected a device to initiate the cascade, the following will occur:

- ALL VLANs that are present on that device will be considered "in play" for the cascade. If, for example, VLANs 10, 20, and 30 are present on the selected device, VLANS 10, 20, and 30 will be considered "in play".
- That top-level device will be assigned to the specified VLAN Domain.
- The cascade will then check to see if there are any other devices that are linked to the initially selected device. If there are, it will check to see if those devices have VLAN traffic - from the VLANS that are considered to be in play - passing over them (VLAN data is assumed to be going over a link if both endpoints have at least one shared VLAN - from the VLANs that are in play - on them).
- If a linked device is determined to be in list of VLANs that are "in play", the cascade process will continue that device, and branch out from there.

Deletinga VLAN Domain

When you no longer need a VLAN Domain, they can be deleted from the "VLAN Domain" portlet. You can initiate this process bygoing to the portlet, selecting the VLAN Do main you wish to delete, and selecting the "Delete" option from the menu.

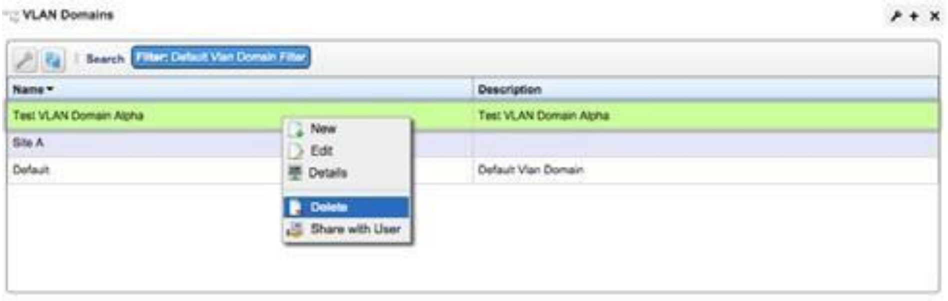

Name ▾	Description
Test VLAN Domain Alpha	Test VLAN Domain Alpha
Site A	
Default	Default Vlan Domain

After confirming the deletion, the VLAN Domain will be deleted, and VLAN Membership data will be recalculated.

NOTE:

The Default VLAN Domain is seeded at installation and cannot be deleted.

Consuming VLAN Data

After having assigned equipment to VLAN Domains and having executed the "Get VLAN Data" for the desired devices, you are then able to view/work with the data that has been collected. There are several places in which the data is available, and they will be discussed here.

VLAN Portlets

The previously covered VLAN Portlets (VLAN Domain s, VLAN Domain Assignments, VLANs, and VLAN Memberships) will be populated with all known/collected da ta. These portlets can be placed on whichever pages you like and will display their relevant data.

In each portlet you can filter the data with simple and "Advanced" filters:

Example "simple" filter, displaying on VLANs with VLAN Id 91:

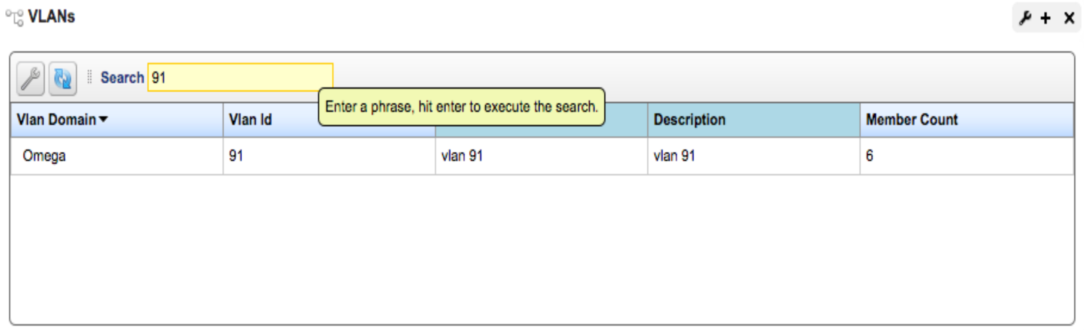

Search 91

Vlan Domain	Vlan Id		Description	Member Count
Omega	91	vlan 91	vlan 91	6

Example "Advanced" filter displaying all VLANs on VLAN Domain "Omega"

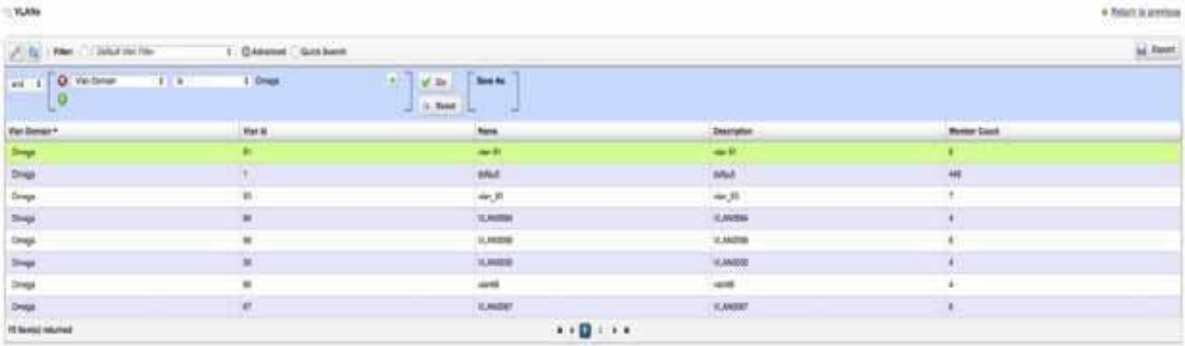

Vlan Domain	Vlan Id	Name	Description	Member Count
Omega	91	vlan 91	vlan 91	6
Omega	1	default	default	446
Omega	93	vlan_93	vlan_93	7
Omega	94	VLAN0094	VLAN0094	4
Omega	98	VLAN0098	VLAN0098	6
Omega	30	VLAN0030	VLAN0030	4
Omega	98	vlan98	vlan98	4
Omega	87	VLAN0087	VLAN0087	6

15 item(s) returned

Details Pages

VLAN Data is present in the "Details" pages in multiple places, and it is context sensitive. For example, when viewing the Details page for a given device, you will find all VLAN Memberships for that device listed under the "Network" tab:

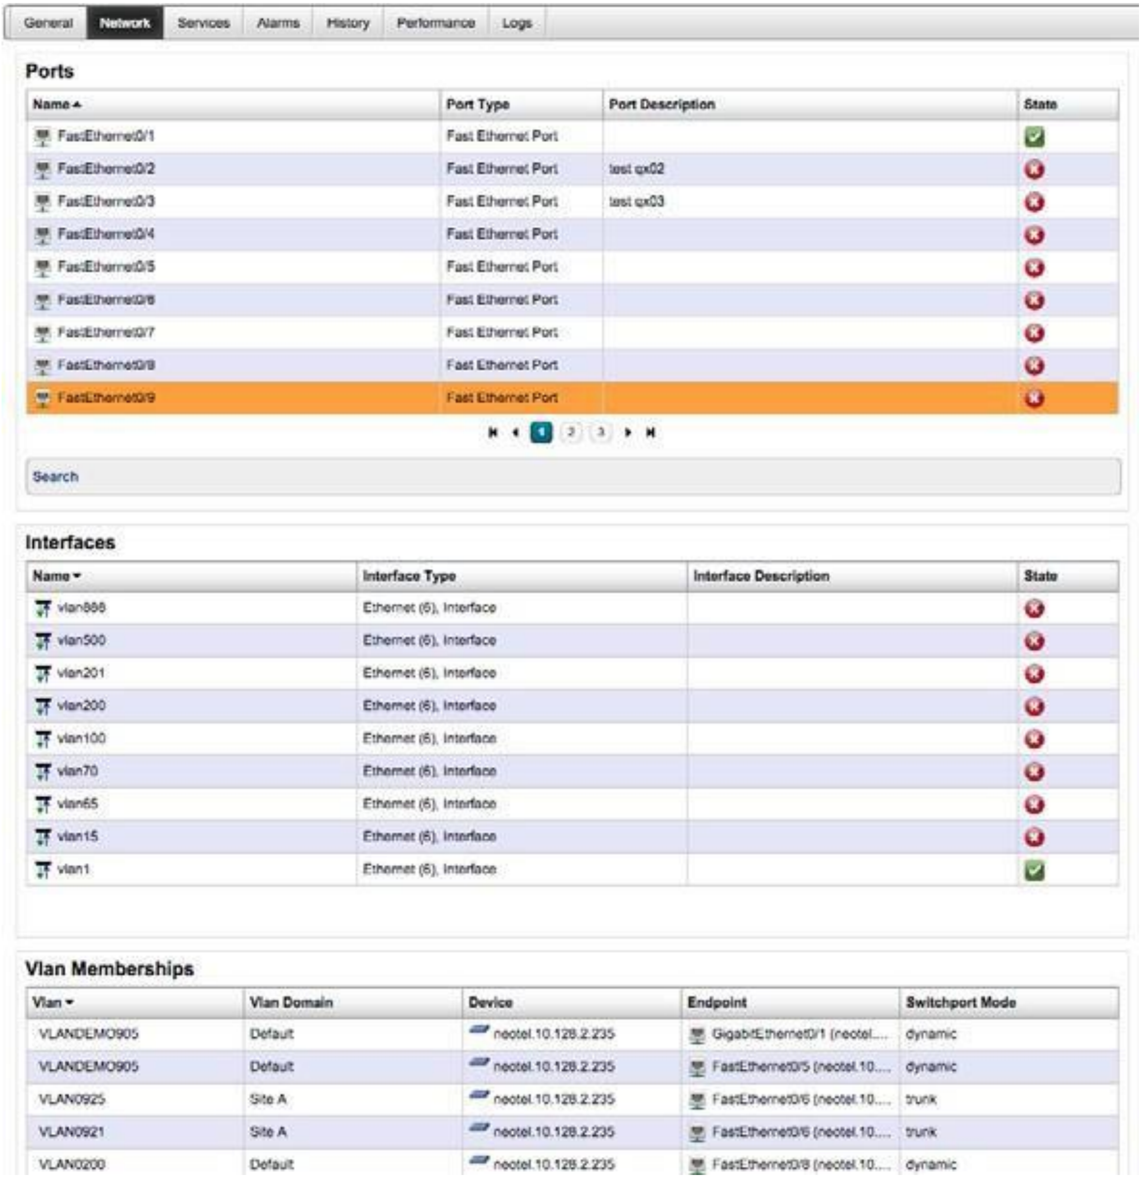

Ports

Name ▲	Port Type	Port Description	State
FastEthernet0/1	Fast Ethernet Port		✓
FastEthernet0/2	Fast Ethernet Port	test qx02	✗
FastEthernet0/3	Fast Ethernet Port	test qx03	✗
FastEthernet0/4	Fast Ethernet Port		✗
FastEthernet0/5	Fast Ethernet Port		✗
FastEthernet0/6	Fast Ethernet Port		✗
FastEthernet0/7	Fast Ethernet Port		✗
FastEthernet0/8	Fast Ethernet Port		✗
FastEthernet0/9	Fast Ethernet Port		✗

Interfaces

Name ▼	Interface Type	Interface Description	State
vlan888	Ethernet (6), Interface		✗
vlan500	Ethernet (6), Interface		✗
vlan201	Ethernet (6), Interface		✗
vlan200	Ethernet (6), Interface		✗
vlan100	Ethernet (6), Interface		✗
vlan70	Ethernet (6), Interface		✗
vlan65	Ethernet (6), Interface		✗
vlan15	Ethernet (6), Interface		✗
vlan1	Ethernet (6), Interface		✓

Vlan Memberships

Vlan ▼	Vlan Domain	Device	Endpoint	Switchport Mode
VLANDEMO905	Default	neotel.10.128.2.235	GigabitEthernet0/1 (neotel....	dynamic
VLANDEMO905	Default	neotel.10.128.2.235	FastEthernet0/5 (neotel.10....	dynamic
VLAN0925	Site A	neotel.10.128.2.235	FastEthernet0/6 (neotel.10....	trunk
VLAN0921	Site A	neotel.10.128.2.235	FastEthernet0/6 (neotel.10....	trunk
VLAN0200	Default	neotel.10.128.2.235	FastEthernet0/8 (neotel.10....	dynamic

You will also find context sensitive VLAN Membership data available in the Details for port/interfaces.

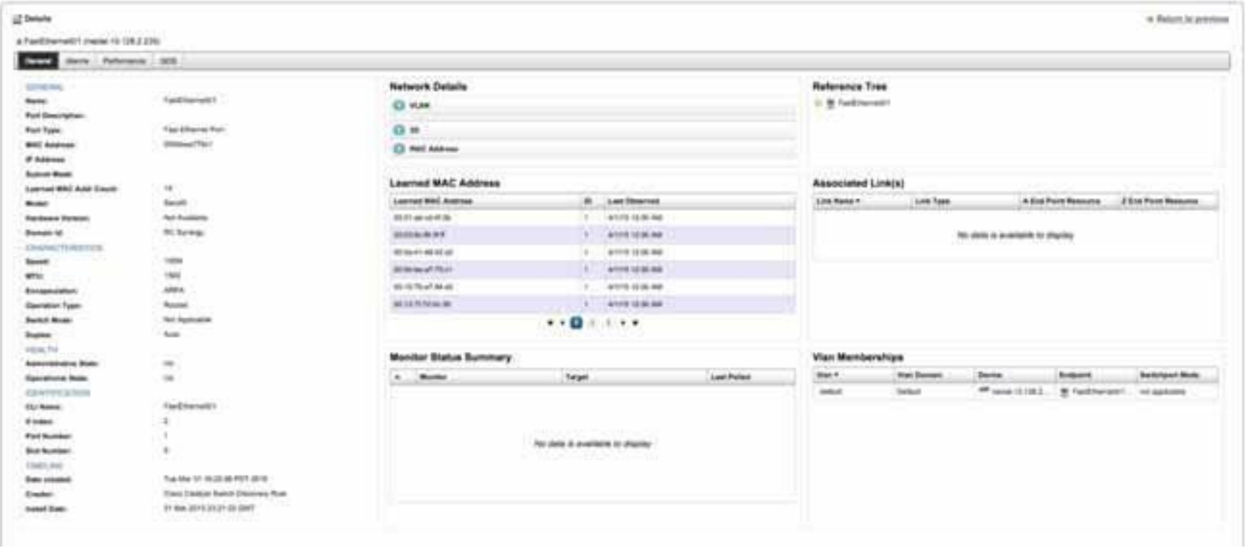

In addition to this, the Details pages for VLAN Domain s, and VLANs also contain context-sensitive listings.

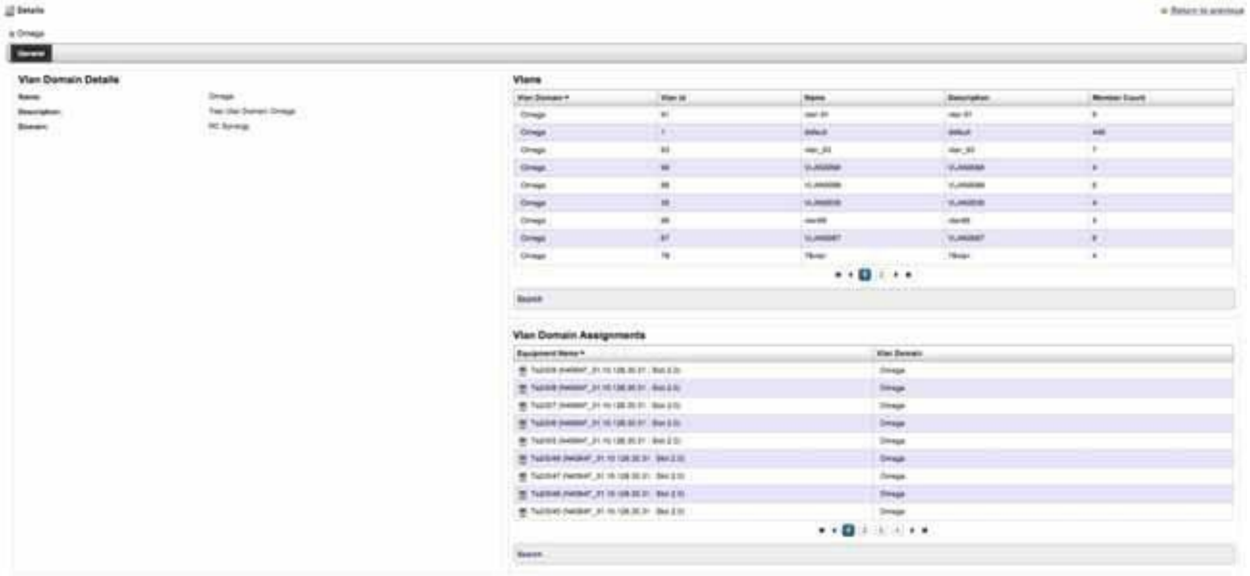

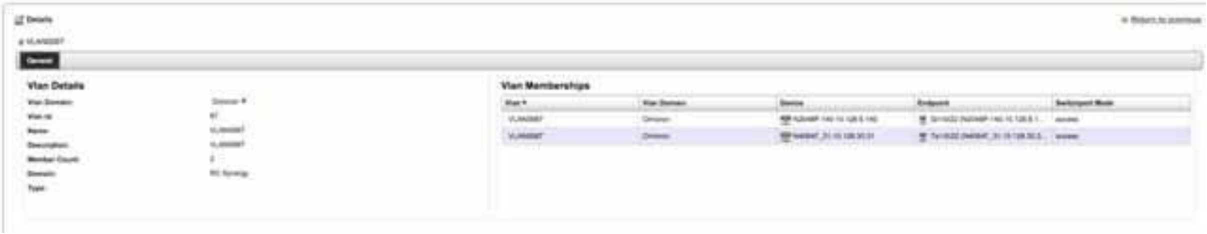

Visualization

In addition to viewing the data it table-based representations, VLANs can be visualized in a topographical representation. To do this, select the VLAN that you wish to Visualize from the "VLANs" portlet, and select the "Visualize" option.

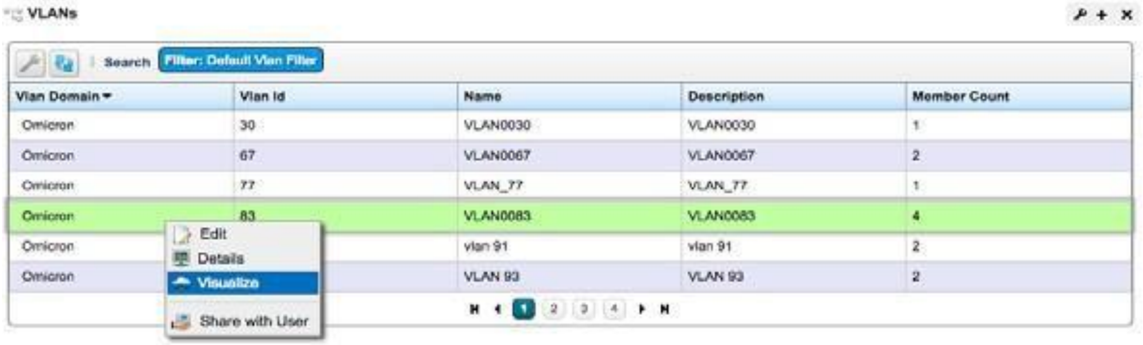

Vlan Domain	Vlan Id	Name	Description	Member Count
Omicron	30	VLAN0030	VLAN0030	1
Omicron	67	VLAN0067	VLAN0067	2
Omicron	77	VLAN_77	VLAN_77	1
Omicron	83	VLAN0083	VLAN0083	4
Omicron		vlan 91	vlan 91	2
Omicron		VLAN 93	VLAN 93	2

Upon doing this, you will be presented with a topographical view of all devices, ports, and interfaces that are members of the selected VLAN. Links that have been discovered that have VLAN traffic - from the selected VLAN - going over them, will also be included in the topographical representation presented to the user.

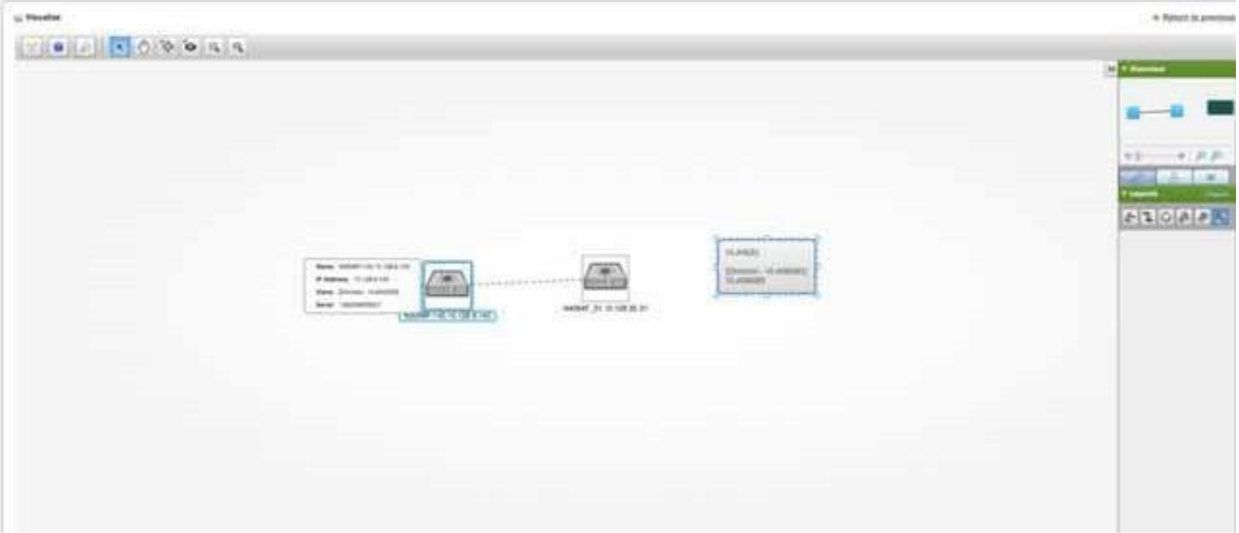

Once you are inside the Visualizer, you will be able to mouse over various node/edges and get additional data. Additionally, ifyouselectdevicelevelnodes, youwillbeable"Expand" thosenodes, to see the individual VLANports/ interfaces on that device.

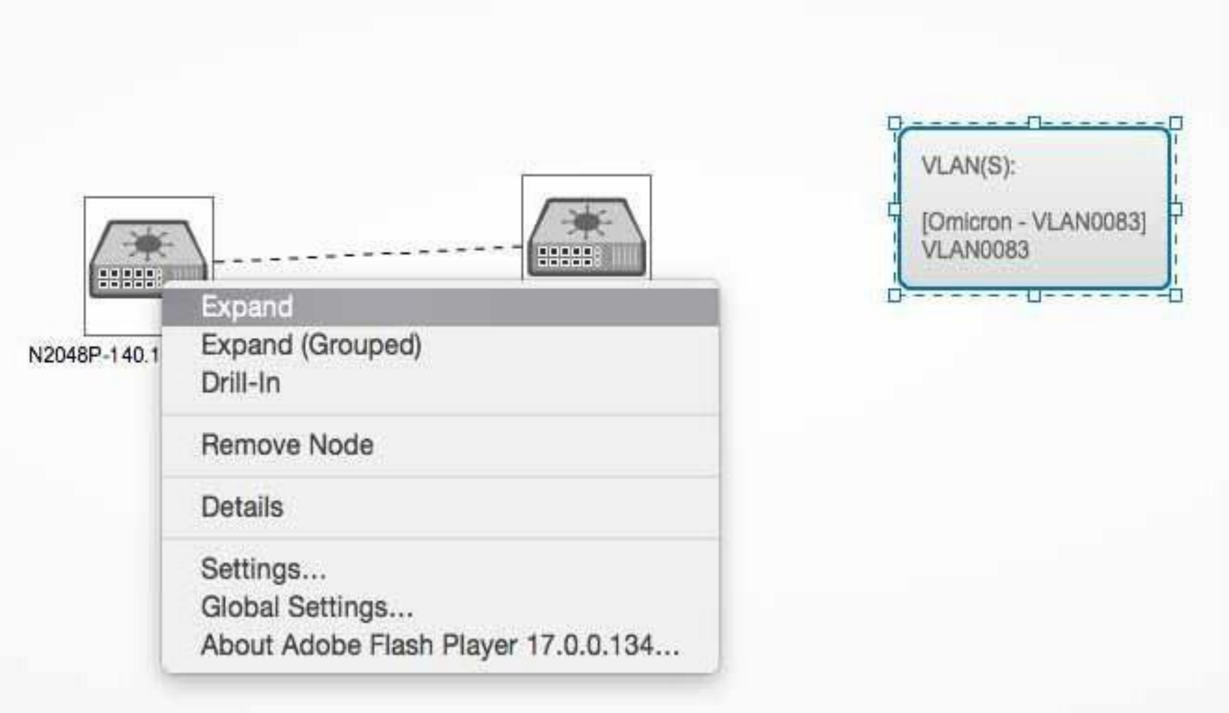

VLAN(S):

[Omicron - VLAN0083]
VLAN0083

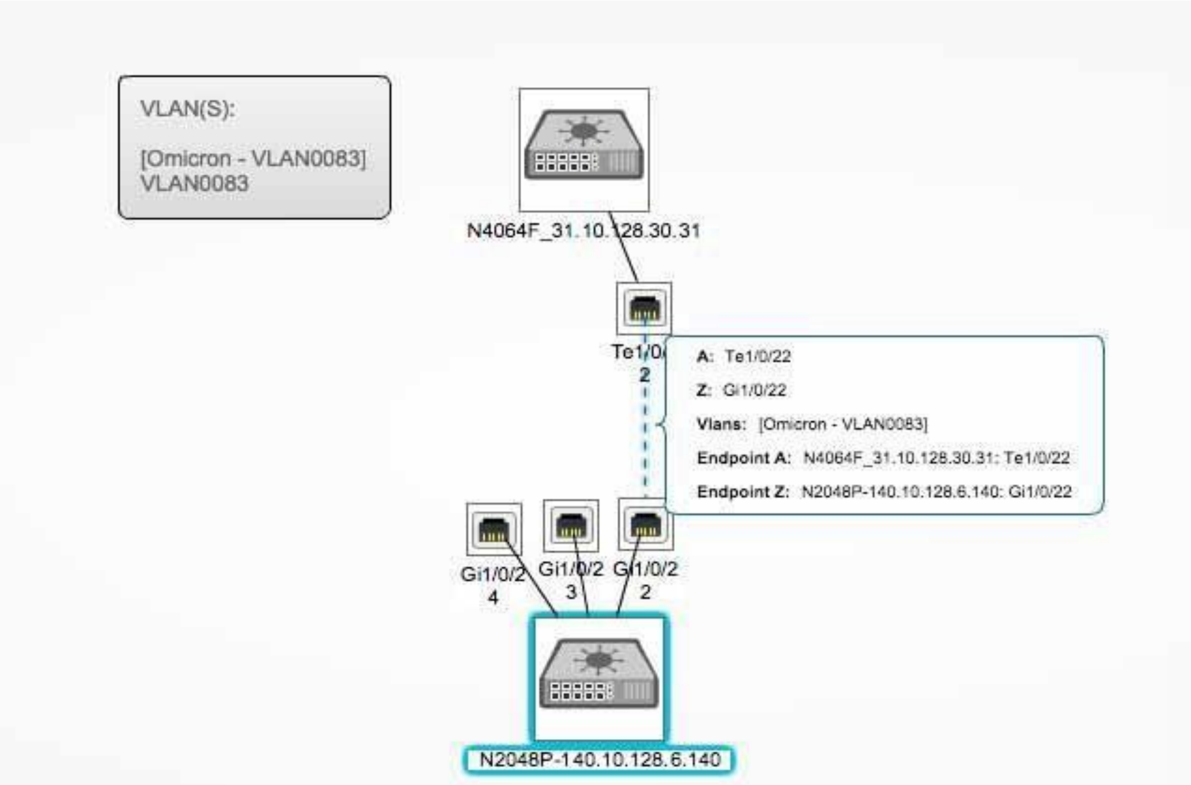

VLAN(S):

[Omicron - VLAN0083]
VLAN0083

A: Te1/0/22

Z: Gi1/0/22

Vlans: [Omicron - VLAN0083]

Endpoint A: N4064F_31.10.128.30.31: Te1/0/22

Endpoint Z: N2048P-140.10.128.6.140: Gi1/0/22

Saving a View

Layout of the nodes/edges is handled on a best-effort basis. It is possible, however, for a user to manually arrange the nodes in manner they desire, and to then save this view for future use. To do this, drag the nodes to their desired location, and then select the "Save this view to the associated context" option from the toolbar at the top of the Visualizer.

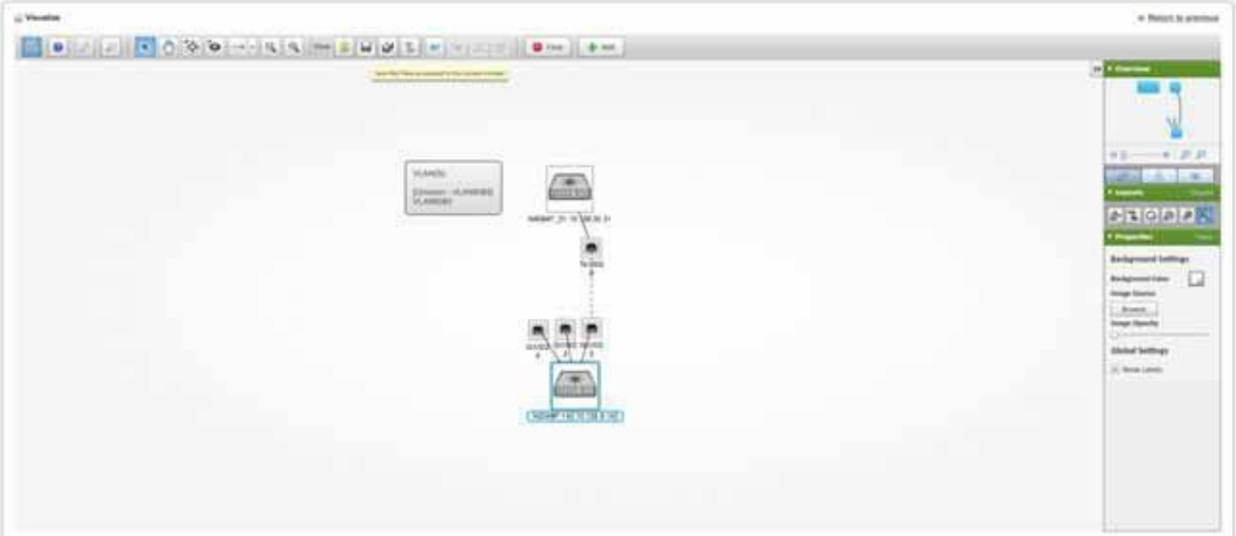

After having saved the view in this way, the saved view will automatically be loaded whenever the same VLAN is selected for visualization.

Multi-Selection of VLANs for Visualization

It is possible to visualize several VLANS at the same time. To do this, you must first maximize the "VLANs" portlet. From this maximized view, you are then able to multi-select VLANs, and then execute the "Visualize" operation.

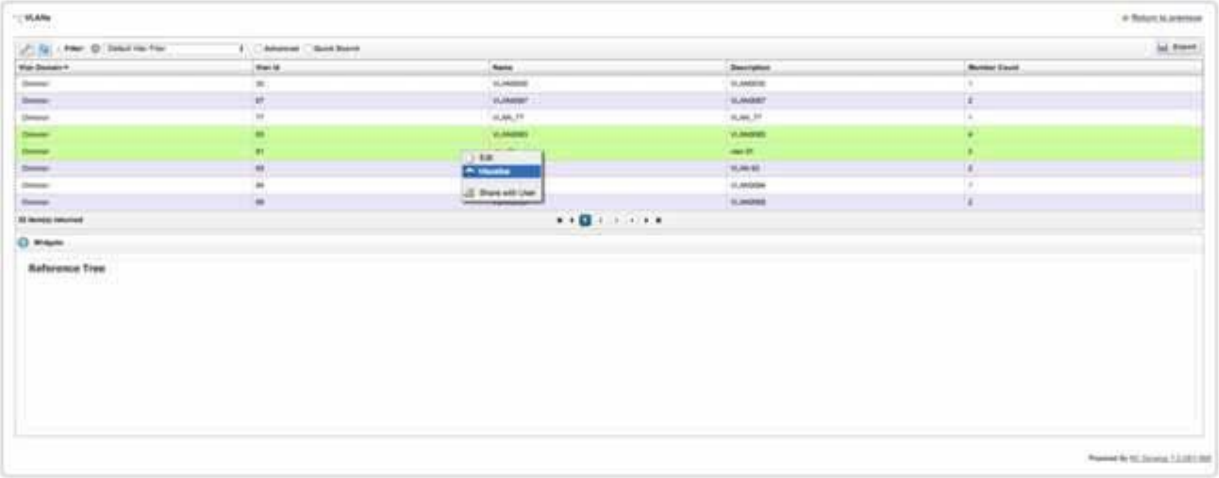

Visualization of multiple VLANs is the same as visualization of a single VLAN, with the exception that all memberships for both VLANs are displayed.

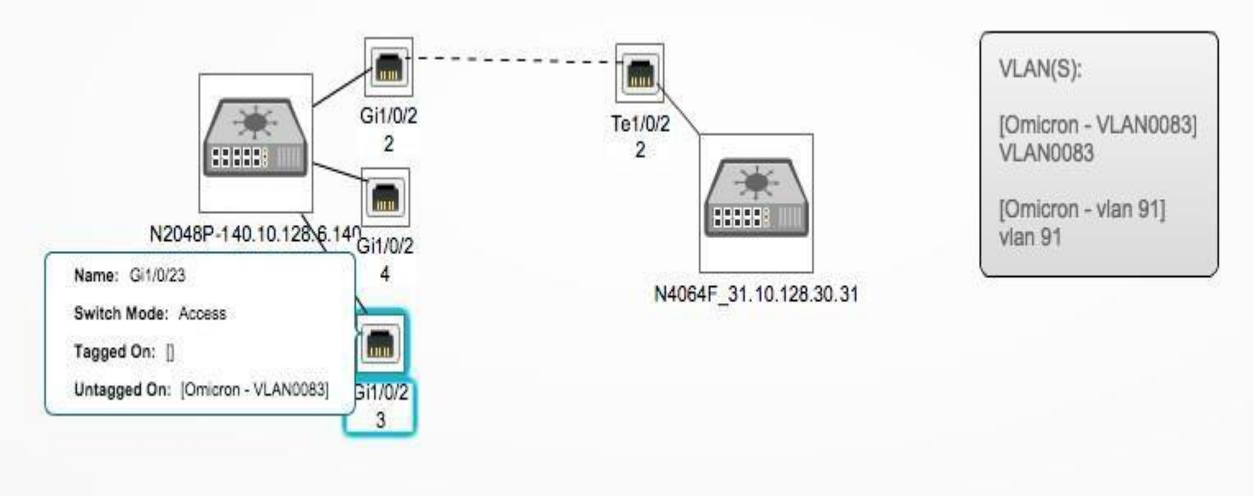

Name: Gi1/0/23

Switch Mode: Access

Tagged On: []

Untagged On: [Omicron - VLAN0083]

Reporting

VLAN, VLAN Membership, VLAN Domain, and VLAN Domain Assignment data is available for use within the standard Reporting features. To make use of it, Report Templates containing the desired VLAN data can be created, Reports using those templates can be constructed, and those Reports can then be executed.

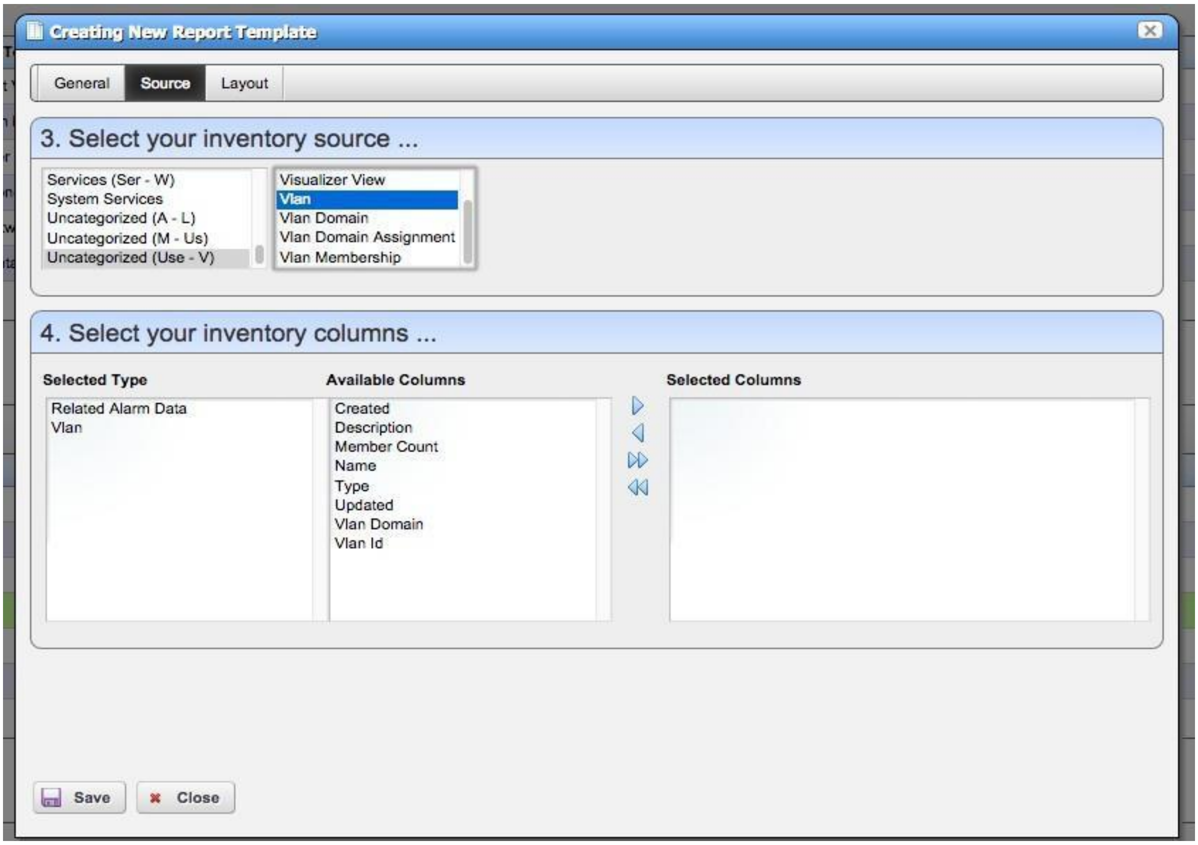

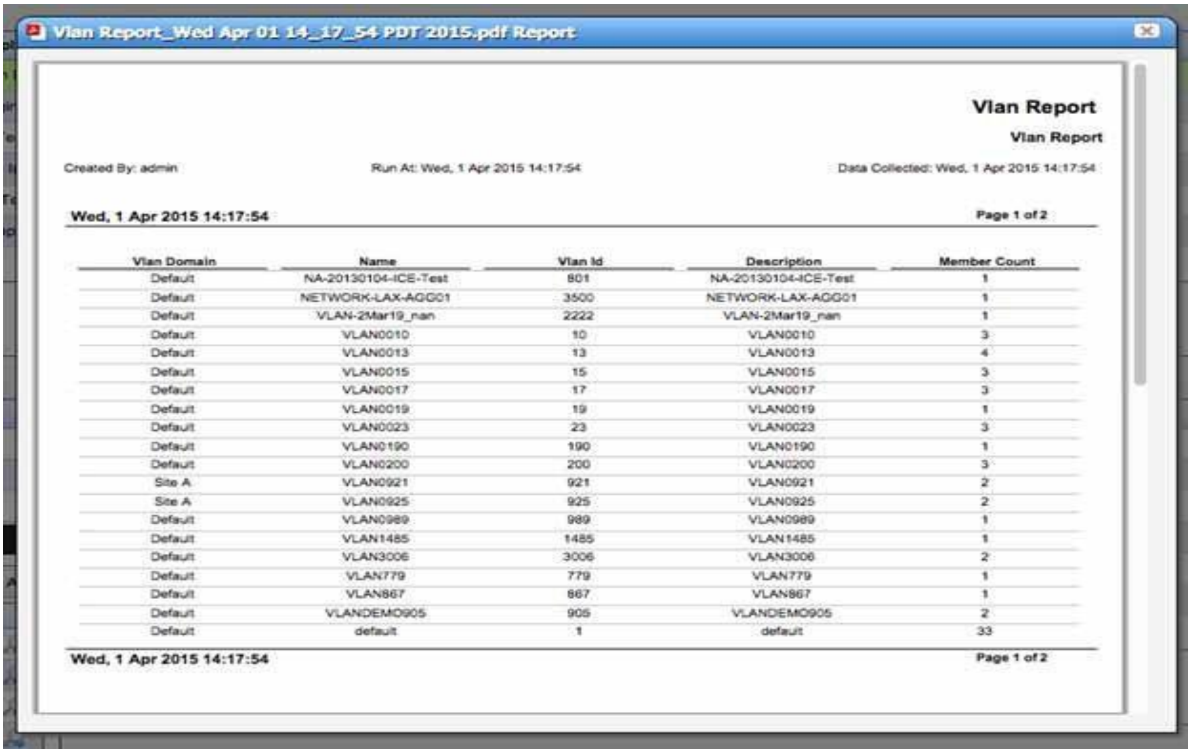

Vlan Report

Vlan Report

Created By: admin Run At: Wed, 1 Apr 2015 14:17:54 Data Collected: Wed, 1 Apr 2015 14:17:54

Wed, 1 Apr 2015 14:17:54 Page 1 of 2

Vlan Domain	Name	Vlan Id	Description	Member Count
Default	NA-20130104-ICE-Test	801	NA-20130104-ICE-Test	1
Default	NETWORK-LAX-AGG01	3500	NETWORK-LAX-AGG01	1
Default	VLAN-2Mar19_nan	2222	VLAN-2Mar19_nan	1
Default	VLAN0010	10	VLAN0010	3
Default	VLAN0013	13	VLAN0013	4
Default	VLAN0015	15	VLAN0015	3
Default	VLAN0017	17	VLAN0017	3
Default	VLAN0019	19	VLAN0019	1
Default	VLAN0023	23	VLAN0023	3
Default	VLAN0190	190	VLAN0190	1
Default	VLAN0200	200	VLAN0200	3
Site A	VLAN0921	921	VLAN0921	2
Site A	VLAN0925	925	VLAN0925	2
Default	VLAN0989	989	VLAN0989	1
Default	VLAN1485	1485	VLAN1485	1
Default	VLAN3006	3006	VLAN3006	2
Default	VLAN779	779	VLAN779	1
Default	VLAN867	867	VLAN867	1
Default	VLANDEMO905	905	VLANDEMO905	2
Default	default	1	default	33

Wed, 1 Apr 2015 14:17:54 Page 1 of 2

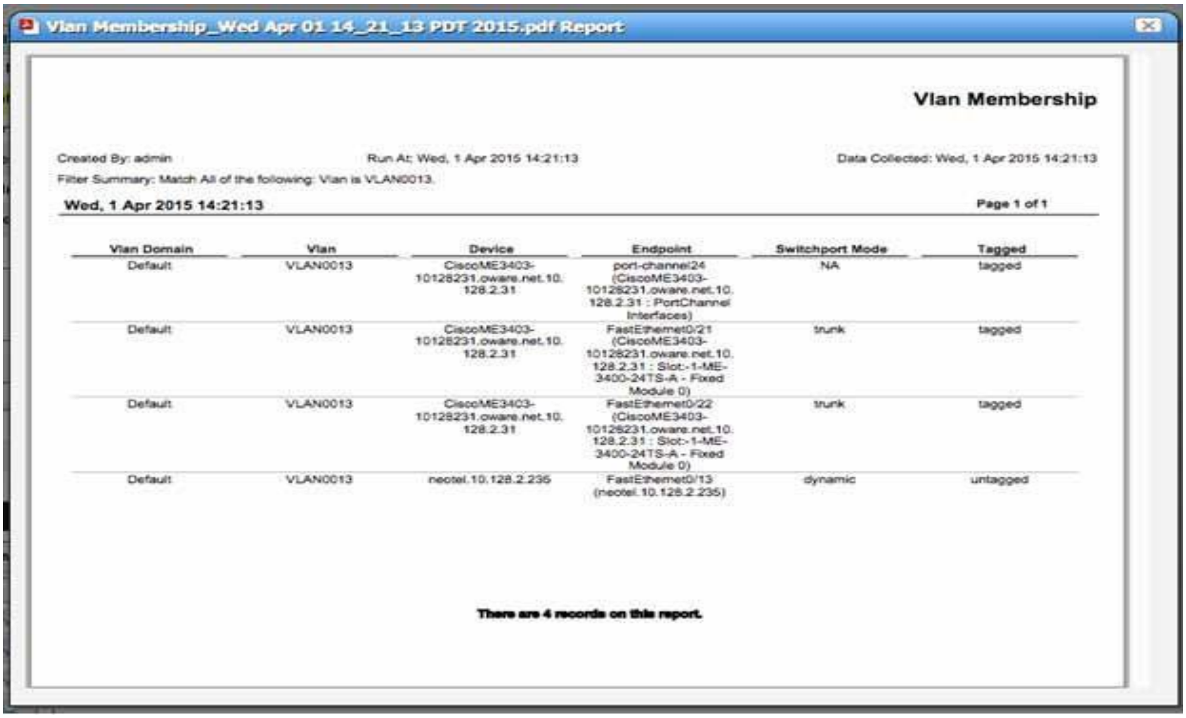

Vlan Membership

Created By: admin Run At: Wed, 1 Apr 2015 14:21:13 Data Collected: Wed, 1 Apr 2015 14:21:13

Filter Summary: Match All of the following: Vlan is VLAN0013.

Wed, 1 Apr 2015 14:21:13 Page 1 of 1

Vlan Domain	Vlan	Device	Endpoint	Switchport Mode	Tagged
Default	VLAN0013	CiscoME3403-10128231.oware.net.10.128.2.31	port-channel24 (CiscoME3403-10128231.oware.net.10.128.2.31 : PortChannel Interfaces)	NA	tagged
Default	VLAN0013	CiscoME3403-10128231.oware.net.10.128.2.31	FastEthernet0/21 (CiscoME3403-10128231.oware.net.10.128.2.31 : Slot:-1-ME-3400-24TS-A - Fixed Module 0)	trunk	tagged
Default	VLAN0013	CiscoME3403-10128231.oware.net.10.128.2.31	FastEthernet0/22 (CiscoME3403-10128231.oware.net.10.128.2.31 : Slot:-1-ME-3400-24TS-A - Fixed Module 0)	trunk	tagged
Default	VLAN0013	neotel.10.128.2.236	FastEthernet0/13 (neotel.10.128.2.235)	dynamic	untagged

There are 4 records on this report.

For more information on working with Reports, pl ease see the associated section in User Guide.

14

Troubleshooting YourApplication

The following describes troubleshooting steps from installation to execution of this application. Installation logs are in the directory `$OWARE_INSTALL_ROOT`. Log files are `setup.log`, `app_setup.log`, and `db_setup.log` (this last log does not appear if you install an Oracle database).

NOTE:

Because this software installs in many different settings, and permits many add-ons and options, not all troubleshooting tips may apply to your installation.

Troubleshooting Prerequisites

Before you begin troubleshooting any serious problem you will need the following:

System details

- Hardware specifics (as applicable)… processor, memory and disk (free space, and so on).
- Operating system and version?
- Cruz version
- Browser and version? Java? Flash?
- Single or distributed installation? Clustered?

Environment details

- How many managed devices?
- Main features used?

Troubleshooting

- Screenshots of errors
- Logs (the `logs.jar` file produced by running `getlogs`)
- A complete description of the problem, and the steps to reproduce, and a complete description of anything already tried that did not work.

Mini Troubleshooting

Suggested mini-troubleshooting steps for a balky application that is already installed and running:

1. Refresh the browser. If that does not work...
2. Clear the browser’s cache (Firefox loves persistent old pages), then refresh.

When you see a difference between direct access behavior between browsers on two different machines, delete temporary internet files. In Windows, open control panel, and open Java. Click the *Settings* button in the *General* panel and click *Delete Files*. Delete for *Applications and Applets* and *Trace and log files.*

If that does not work...

3 Stop and start the browser. If that does not work...

4 Stop and start the web server

For Windows, to start the webserver manager: `oware\synergy\tomcat-X.X.X\bin\startsynergy`. For Linux.

`/etc/init.d/synergy start` or `/etc/init.d/synergy stop`

Worth noting: The tray icon for the web server (▶ is “optimistic” about both when the web server has completely started and completely stopped. Youcannot re-start web server when its Tomcat process still lingers. If you lack patience, kill the (large) Tomcat process then re-start web server. The smaller one is that tray icon.

If that does not work...

5 Stop and start application server. Command lines for this:

`stopappserver` and `startappserver`

If that does not work...

6 Delete the contents of the `oware/temp` directory and restart application server. If that does not work...

7 Reboot the host andre-start the application server, web server and browser.

When troubleshooting (or contacting technical support), you may find pertinent information in logs located in the following directories:

```
..\oware\jboss-5.1\server\oware\log
..\oware\temp\soniqmq.log
..\app_setup.log
..\db_setup.log
```

You can also run `getlogs` from a command line.

NOTE:

If you see errors that say your Linux system has too few threads, make sure you have set the file handles correctly.

You may see: The "<portlet > is temporarily unavailable" in portlets after install or upgrade. Workaround: restart the Web server by stopping and re-starting the synergy service. This can be done from the Windows tray icon by right clicking on the Web server.

Troubleshooting Adobe Flash

Installing the latest Flash version is a part of Cruz’s requested prerequisites. When Flash is not installed on the browser, things like System Topology, selecting a license file, importing a file, selecting an OS image to import, and so on, do not work. Selection dialogs require installed Flash to select files from the local system. When Flash is not installed, such buttons appear to be active but they do not work.

Database Aging Policies (DAP)

DAP policies automatically purge or archive stale data so the database can maintain its capacity. Several pre-defined and pre-seeded DAPscome with Cruz. Youmay need to revise these to fit your system. These start at specific times—see the Schedules portlet for specifics about when.

DAPsamount to preventative maintenance since they help to maintain the database's capacity. Best practice is to do the following regularly:

1 In the Audit Trail Manager, create a Filter for Creation Date = prior Month and Action = DAP Executed.

2 Review the records for Status – Failed. These indicate that a DAP job failed. As long as the following DAP jobs execute, no immediate action is required. If any DAPs are repeatedly failing, then consult the troubles hooting document or Cruz support.

3 Review the DAP jobs entries and compare to the scheduled DAP start times. Confirm that audit records are displaying a corresponding audit record for each scheduled execution.

Installation Issues

If you are having difficulty installi ng the application, refer to the *Cruz Installation Guide,* Preparing for Installation section may help you resolve installation problems.

Installation File Issues

When installing from a compressed file, file corrupt ion can result from an incomplete file copy, or (FTP) transmission. One symptom of corruption: the file will not unzip. Corruption can also appear when you copy unzipped installation files from Windows to Linux. These can pick up erroneous line feed characters. **Workarounds**:

- FTP installation files from Windows to Linux
- Copy the entire ZIP or compressed file first, then uncompress/unzip it.

Installer Failure

- The installer fails, and you are installing on Windows 2012 (this error appears: `Installer User Interface Mode Not Supported.`

  Workaround: Right-click the `win_install.exe` file, and, in *Properties*, select compatibility mode for Windows 7 or Vist a before (re)initiating installation.

- If you created an installation CD from unzipped package:

  -CD formatting limitations can truncate file names to eight characters

  -`Setup.log` complains about the absence of the `owareapps` directory

  -The directory name is truncated to `owareapp` (no "s"), `win_install.exe` becomes `win_inst.exe`

  Workaround: FTP thezipped package. Ifyouareburning an installation CD, use the ISO file, and something like the `mkisofs` utility as the input to the CD-ROM burner.

Additional Windows problems include the following:

- Installer does not provide options to select the desired network interface card /IP address. **Workaround**: Removetheline `InstallMode=Simple`fromthe `setup.ini`filein the installation source directory. Thisenables installation in `standard` mode which provides installation options and detection of a second network interface card (NIC).

- Installer fails immediately with error `Create Process failed ==> %1 is not a valid Win64 application.`

  Solution: Change the value of the `%TEMP%` environment variable. Change the default value of the `%TEMP%` environment variable to another path you have already configured, for example: `C:\Temp`. Use the Windows System Tools menu to do this.
- The installing user must have write access on target directories where the application is installed.
- The same user who installed the software must initiate any uninstallation. Uninstalling may also encounter locked files and directories. Th is may leave files and directories behind since locking prevents their deletion. Forcompleteness' sake, recommended practices are manual deletion or the use ofan unlocker program. Clean directories are important, particularly if you are uninstalling then re-installing,

Other Installation Issues

Youmaysee messageslike `failedtoresetpassword`. Thefollowingdealswith thisand other potential problems:

- Make sure you can resolve the hostname to the correct IP address: `ping -a [IP address]` and `ping [hostname]` and make sure they are in sync.

  If the application becomes unusable after changing the application server IP address (post-install).

  Solution: To change the application server IP address:

  1. Shutdown application/web server
  2. Open command pro mpt/shell and source environment. (for Windows: Type `oware`, for Linux: Type `. [space]/etc/.dsienv`, meaning: `. /etc/.dsienv`)
  3. Modify `oware.appserver.ip` property in file `[installation root]\oware\synergytomcat-x.x.x/webapps/ROOT/WEB-INF/classes/portal-ext.properties`.
  4. Verify the file does not specify the old IP anywhere else. If it does, replace with new IP address and save.
  5. Next, modify the IP address in any shortcut URL properties and click OK. `[installation root]\oware\synergy\tomcat-x.x.xx\bin\portal.url` (Web Document tab)
  6. Type `ipaddresschange -n <new IP address>` in a command prompt/shell.
  7. Restart application and web server.
- Ensure any other of the application's installation(s) have been completely removed, if you have uninstalled a previous version.
- Make sure no other MySQL databases are installed (some Linux packages include them by default).
- You must be user with administrator permissions to install.

Standalone Database Installation Problems

- Errorsappear after installing a standalone database, even though its connection has been verified successful using pingdb utility.

  Solution: After setting the environment in a shell (Windows: `oware`, Linux: `. /etc/.dsienv`) Verify the following commands have been run on the application server.

- `loaddb` (create Cruz database)
- `loaddb -d` instead of `loaddb` if the tablespace has not been created.
- `loaddb -s` (create synergy/portal database)
- `dbpostinstall` (seed components and resolve database schema changes)

After confirming the above, start the application server and synergyportal after application server is up/ready.

Install-From Directory

Installation package files must not be in the installation destination directory.

Windows

Install in `d:\dorado`, unzip package in, for example, `d:\source` or `d:\installsrc`, not `d:\dorado` (the installation default target).

Installer Logs

Installation logs are in `$OWARE_INSTALL_ROOT`. Log files are `setup.log`, `app_setup.log`, and `db_setup.log` (this last log does not appear if you install an Oracle database). An empty or missing `app_setup.log` means no applications were installed. This can result from a truncated `owareapps` directory name. **Solution:** Correct the directory name and attempt the installation again.

Startup Issues

The following are some possible problems with application startup. Remember: after you first install this application, the application server takes longer to start. Be patient the first time you start the application.

The two most common reasons for the inability to start an application server which worked previously are the following:

Application Server Does Not Start

To find application server issues, search the server log (`\oware\jboss-x.x\server\oware\log\server.log`) file forthe word `error`.Reviewthelogforthefirst error or exception. This is typically the item that needs to be resolved and the most relevant for troubleshooting information.

- Installation checks to confirm your hardware is adequate, but if you are installing to a VM, you can reduce hardware allocations after that installation. If you do that, your application server may not start and will provide no logs. This indicates you may have inadequate hardware, or an inadequate portion of your hardware assigned to the VM running Cruz. If this occurs, check the hardware requirements, and reconfigureyour VM, or install on a different machine.

  - Check the total system RAM

  - Check the appserver heap setting in `/owareapps/installprops/lib/installed.properties`.If it exceedsthe totalsystem RAMeither you mustallocate more RAM or reduce the appserver setting

  Application server should never be allocated le ss than 3G ram, and Web server should never be allocated less than1G RAM.

- The following socket errors (or similar) are reported when starting the application.

  ```
  (Feb 14, 2014 4:57:50 PM) [OWProcessMonitor] ERROR! command socket
   failure: Address already in use: Cannot bind
  (Feb 14, 2014 4:57:50 PM) [OWProcessMonitor] re-initializing command
   socket: attempt 1 of 10
  ```

  Thisindicates aportconflict islikelypreventing your systemfromfunctioning properly. This typicallyoccurs when network management or other applications reside on the same server. Determine the source of the port conflict and remove the application. Best practice is to install on a dedicated server platform.

- Port Conflicts, as described in the previous paragraph, can arise after installation. Forexample, if you install other software that uses the SNMP ports (161 and 162), after you have installed Cruz, the installation cannot catch such a conflict. When you try to start it, application server will report it cannot bind to that port and fail. See Ports Used on page 979 for a list of potential conflicts. The `server.log` file lists to such errors when they occur.
- Confirm your license is current and installed. Search the `\oware\jboss-5.1\server\oware\log\server.log` file for `error`, and the license expired error appears.

  To install a license file without a running application server, run the `licenseimporter` script:

  ```
  licenseimporter license.xml
  ```

  Your license may have a different name, but the script syntax is the same.

No valid HA license is found during installation of HA

After installation of thesecondaryappserver, theprimaryservergoesdown (appserver iconchanges from green to red), this most likely happened because incorrect license was applied for HApackage. To check the status why appserver went down, right click on the appserver red icon and select logs/ server option and search for Beginning Server Shutdown. You will see:

2015-08-05 13:02:59,738 3040923 FATAL [com.dorado.mbeans.OWClusterPeerActiveImpl] (Timer-8:) APPLICATION SERVER ERROR

2015-08-05 13:02:59,739 3040924 FATAL [com.dorado.mbeans.OWClusterPeerActiveImpl] (Timer-8:) License Not Found for Clustering Servers in High-Availability Mode

To resolve this issue, contact Support to generate a HA license. Once HA license is received, place it anywhere where appserver is installed then bring up web client and log into Cruz. When at home page, select "License Management" and clickon "Select File" button. Navigate to where the license file is located and click on "Open" button. At this stage, after few seconds, you should receive a popup message stating that license got imported. Restartappserver.

Startup issues for Windows installations

If an application or mediation server goes down ungracefully for any reason the JMS message database may be corrupt in `$OWARE_USER_ROOT\oware\jboss*\ server\oware\data`. When it becomes corrupted the application or mediation server cannot start.

Workaround: Delete thecontentofthe `data`directory. Thisallowsapplication ormediationserver startup.

See Linux Issues on page 681 for additional information about troubleshooting Linux.

Initial Logon after installation fails

If, after installation, your attempt to logon to the application fails with message `Connection to server failed` in the Logon window.

Solution: The most likely cause is that the application server is not running. (Re)start it.

The icon in the Windows System Tray or the presence of the `java.exe` process indicates the status of the application server.

If the icon is red or yellow, no client can connect (although some portlets appear without the benefit of application server). If the icon is red, right-click on it and select *Start* from the menu. The icon turns yellow as the application server starts. Wait until the icon is green and repeat the logon procedure. If this does not work, contact technical support.

Direct Access Fails Because of Java Security Settings

Some Java installations' security settings (in v.1.7 +) may block self-signed websites, interfering with Direct Access. The workaround is to provide a security exception for the application server, as follows:

1. Click Start
2. Type configure java and hit [Enter]
3. Select the Security tab.
4. Click Edit Site List
5. Click Add
6. Typethe Cruz URL(example: `http://192.168.0.51:8080/`. Best practice is to use the IP address of the application server, not `localhost` or `127.0.0.1`)
7. Click OK and Continue.
8. Apply this change, and/or click OK.

Logon Fails with Invalid Logon Message

When you enter your User Id and Password in logon dialog and click Logon, an Invalid Logon message appears.

Solution: Ensure you are entering the correct User Id and Password and click Logon. If you have forgotten the User Id or the Password, or if another user has changed the Password for the User Id, you may have to re-install the software.

Mediation Server on separate machine fails

Distributed mediation server failures to start can occur when multicast isdisabled on your network. The workaround is to this property in
`owareapps\installprops\lib\installed.properties`

```
    oware.application.servers= [application server IPaddress]
```

Correct this on all mediation server machines.

Unsynchronized Clocks in Clustered Installations

All machines in acluster, or distributed system, must havesynchronized system clocks. If this is not true, systems may start, but will not workcorrectly.

Other Failures on Startup

- Another instance ofappserver/medserver maybe running on one host. Common error contents:

```
2005-06-25 08:47:51,968 ERROR [org.jboss.web.WebService] Starting failed
java.net.BindException: Address already in use: JVM_Bind
at java.net.PlainSocketImpl.socketBind(Native Method)
```

Solution: Stop and if necessary, restart application server(s).

Starting and Stopping Servers

Best practice in Windows installations is to start or stop application server with the process monitor icon (installed by the application), when it is installed as a service. Stopping, starting or restarting the service through operating system's Service Manager is not recommended. If the status icon in the tray is green and you rest art procman ("process manager") from Windows' Control Panel services, an error message appears saying the service did not stop properly. The tray icon then turns white. Since the application server is still running, when you try to restart the service again, the icon turns red.

To restore the Process monitor icon to function correctly and show status, stop all Java and WMI processes in the process manager. A system reboot also re-initializes the OWProcMan (process monitor). Note that the service name may be different if your package has been specifically branded. The executable path for the service is `\ ...\oware\bin\owprocman.exe.`

More Failures on Startup

- Failure to connect with a database can occur when...

  –The Oracle instance not running

  –Oracle or separate MySQL lacks connectivity. Use `pingdb -u <user> -p <password>` to check. Default user/password for MySQL: root/dorado

  –Oracle database is not built, or you have not completed its installation

  –MySQL not running (it should start automatically)

  –Your firewall blocks ports the database needs.

  –You see `ERROR. java.lang.OutOfMemoryError` in the log file.

  The following will solve failure to see a login screen in web client:

  1. Shutdown web server
  2. In a command shell type the following: `oware`
  3. Type `mysql -u root --password=dorado` This will log you into mysql database
  4. Copy and paste into mysql shell and hit enter:

     update lportal.layout set typeSettings = `'layout-template-id=1_column\ncolumn-1=58,\n'`

     where (`groupId=10180`)

     AND(`privateLayout=0`)

     AND(`friendlyURL='/login'`);
  5. Restart web service.
  6. Log into Cruz > login form comesup

- Changed properties files

  –Delete the contents of `oware\temp`

- Connection to application server fails

  –Application server has not fully started. Look for `>>>>> Oware Application Server initialization COMPLETE. <<<<<` in the `server.log`

Login Failures

- Invalid Logon

  –Incorrect log in ID or password (the default for web client is User *admin*, password *admin*)

  –User has changed and forgotten password

- Account is inactive

  –Application Security Policy may be configured so passwords or accounts have an expiration date.

  –Use different account.

- Memory errors that indicate too little memory on the application server can also prevent login.

Multi-NIC Host Fails to return to the portal — When you click the *Return to ...* link from Control panel, some unexpected URL appears in the browser. To see the root of this problem, go to *Portal* > *Portal Settings* and compare the *Virtual Host* entry to the application server's IP address. If they are different, then DNS has two different IP addresses associated with the same hostname. Cruz needs an unambiguous IP address and associated hostname for both application server and client.

Workaround(s): 1. Use a local host file and map the IP selected during installation to the hostname. 2. Set the DNS server to only resolve the selected IP address to this machine's hostname.

- Installer fails immediately with error `Create Process failed ==> %1 is not a valid Win64 application.`

  Solution: Change the value of the `%TEMP%` environment variable. Change the default value of the `%TEMP%` environment variable to another path you have already configured, for example: `C:\Temp`. Use the Windows System Tools menu to do this.

Troubleshooting Flow

As part of the troubleshooting process, you can often determine the culprit for problems by a process of elimination. The following questions may help determine what the real issue is:

Discovery/Resync

See Communication Problems on page 658, Preventing Discovery Problems on page 659 and Discovery Issues on page 660

- Can you ping the device you are having difficulty discovering? If you can ping them, and have discovered them, but ping still does not work from within the application, do the devices have an ICMP management interface when you right-click > *Edit* them? If not, add the interface and resync.
- Is your system permitted access to the device (on the Access Control List)?
- Are firewalls blocking access to the device(s)?

NOTE:

The command `service iptables stop` turns off the Linux firewall. Turning it off temporarily is recommended when you first install.

- Is any other software on the application server/mediation server host causing a port conflict? (Uninstall it)
- Is SNMP is configured on the target device and read/trap, write community strings? Is SNMP correctly set up? (check with Network Toolsand MIB browser or a tool like iReasoning's MIB browser)

- Is Telnet or SSH configured on the target device and can you Telnet/SSH to the device through a command line shell or an application like puTTY?

NOTE:

Some devices support only SSH v2. Consult release notes for specifics.

- Areauthentications created in the Authentication portlet with protocolsand passwordsset correctly, with adequate timeout and retries configured for your network's latency?
- Are Discovery Profiles using the created authentications?
- Does device fail to resync after deployment?

  wait up to a minute before resync as SNMP agent is not enabled yet on the device.

Backup/Restore/Deploy

See Backup/Restore/Deploy on page 661

- Is your FTP server installed, up and running?
- Is that FTP server on the same side of the firewall as the devices it addresses?
- Does the device support the type of backup (FTP, SFTP, TFTP) you are attempting?
- Do your authentications grant privileged access? The prompt is typically #, not > at this level of access.
- If the device does not successfully execute the command, then either the authentication you have used does not have permission to do such commands, or the device is configured to prohibit their execution.
- Do FTP and TFTP servers write to the same directory, and have permissions to read/write/execute to that directory?

Alarms/Monitors/Performance

Consult the User Guide's recommendations, particularlyfor Monitoring and for Traffic Flow Analysis. See Alarm/Performance/Retention on page 662.

- Do you have the recommended hardware to handle the number of devices you are managing?
- Are the devices you are monitoring sending only the relevant traps to yoursystem?
- Is your database configured correctly for the expected load? Symptoms of database configuration inadequacies include slowperformance when expanding the Resourcesportlet or right clicking on a port of a device and selecting show performance. This can also occur if your database size has increased significantly sinceimplementation.

  Solution: For MySQL, adjust/increase the innodb_buffer_pool_size to restore performance. Consult the User Guide for more about performance tuning such parameters in MySQL.

NOTE:

The internal event `emsDBCapacity`notifies you about how much of the database you have used.

- Have you tailored your monitoring to the availablecapacity of your hardware? Monitoring or other functionality dependant on writing to the database may stop with error specifying

  ```
  Could not get a database connection
  ```

  One example of an error that appears when an active monitor which is suddenly unable to insert data into the database

  ```
  2014-04-10 11:14:47,376 490736076 ERROR
   [com.dorado.broadscope.polling.PollingResultsDAOImpl]
   (WorkerThread#8[71.192.23.246:58220]:) persistsPollingResults failed.
   Rolling back.
  ```

```
com.rendion.ajl.CheckedExceptionWrapper: Could not get a database
 connection.
```

Solution: If the database can be reached over the network and has been confirmed operational/healthy, the configured pool allocation(s) may be exhausted. Refer to the *Installation Guide* and confirm sufficient pool allocation s have been configured for corepool, jobpool, and userpool.

The *Installation Guide* ’s Clustering chapter contains suggestions for properly sizing pool values based on the number of servers in your environment. Based on this information and your current configuration, increase pool values accordingly.

- Toisolate thesourceofperformance difficulties, doesun-registering Traffic Flowexporters, or turning off monitors have an impact?
- Not receivingflowsin Traffic Flow Analyzer? Makesure that router/switch isconfigured to send sFlows using port 6343 and NetFlow/jFlows using port 9996.

Services

See Service/Policy Troubleshooting FAQs on page 690

- If you cannot discover services, is the FTP server functioning correctly?

Hardware

See Environment/Operating System Issues on page 692.

- Doesyour hardware match the systemrecommendations for the number ofdevices managed, monitoring and concurrent users?
- Have you followed the installation recommendations (particularly important for Linux) in this guide and *Installation Guide* ?

Advanced Troubleshooting

If you remain unable to resolve issues with your system, the following may be helpful.

- When you contact technical support, create a `logs.jar` file with the `getlogs` command, so you can forward it to them.
- You can change the messages your system generates. That may be necessary. See Debug on page 668.

This chapter contains more troubleshooting advice like WMI Troubleshooting Procedures on page 671, and Linux syslog not displaying on page 683 (setting up devices for various vendors).

Upgrade/Data Migration Fails

I. Unexpected Database Behavior after Upgrade: If you observe unexpected behavior after an application upgrade, reviewinstallation logs. Confirmevidence appears that the upgrade executed `dbevolve`. If not...

Solution: Execute the following steps

1. Shut down the application.
2. Opena shell/command prompt on (the primary) application server.
3. Execute command `dbpostinstall`.

   Thisstep resolvespotential database changes between application versions. Youmustrun the command for both MySql and Oracle database environments too.

II. Upgrade Fails with Database Connection Failure: If an upgrade installation fails with th e message with the `app_setup.log`error `Connecting to database...>>>>> ERROR: OWSessionIDRDBMS : Failed to make database connection`, the problem is that the database is not running on the host being upgraded. To cure this problem, manually start th e database, and then re-try the upgrade installation. The following are the startup commands for the embedded database:

Windows:

```
net start mysql
```

You should see the following response in the shell where you execute this command:

```
The MySQL service was started successfully.
```

NOTE:

Ifyousubstitute the word `stop`for `start`in the above, these commandsmanuallystop the database.

Versions

Before you begin more in-depth troubleshooting, you may need to know what versions of various components are installed to ensure they are compatible. To see these before installing, consult the `version.txt` file in the installation source's root directory, or after the application is running, viewits *Manage* > *Show Versions* screen. Differences between `version.txt` and *Show Versions* may occur when you install additional applicationsor patches.

Another way to see the versions for currently installed modules: open a shell (*Start* > *Run* `cmd` in Windows, for example), and type `oware` (`. /etc/.dsienv` in Linux) and [Enter]. Then type `showversions`. The currently installed modules and their versions appear in the shell. Youcan also use the *Manage* > *Show Versions* menu in web client to find this information.

Search Indexes

Sometimes this software maydisplay Control Panelobjectslike Users, Roles, and Organizations inaccurately. This occurs because Search Indexes need to be re-indexed every so often, especially when changes to Roles, Users and Organizations are frequent.

To re-index go to Control Panel > Server Administration and then click on the *Reindex all search indexes*. Reindexing is not instantaneous; expect this to take some time.

Communication Problems

Firewalls mayinterfere with necessarycommunication between elements within or monitored by your system. Best practice when installing is to bring the firewall down, install, then once you have confirmed the installation runs, bring the firewall up with the appropriate ports open. (See Resolving Port Conflicts on page 669, also see Ports Used on page 979 and Ports and Application To Exclude from Firewall on page 990.)Managed devices often have Access Control Lists (ACLs) for management traffic. Best practice is to use a management VLAN or subnet. Note also that in-path devices may filter management traffic creating an obstacle to management messages. Overlapping address spaces may also complicate network management. Identifying such “DMZs” and overlaps is part of network analysis.

Preventing Discovery Problems

Ensure your firewall is not blocking network access to equipment you are trying to discover. The followingdescribes more preventive practices to do when you discovera mixed vendor/mixed class network.

ICMP (Ping)

You can ping devices from a shell or the Network Tools portlet to insure it’s up and online.

Telnet/SSH

1 Manually telnet or SSH connect to a device to verify that you have the correct authentication information (although Discovery Profiles’ *Inspect* function does this too).

NOTE:

Later versions of Windows do not include telnet by default. In addition to free telnet programs you can download and install, like PuTTY, you can open a shell (*Start* > *Run* `cmd`) and type `oware` to get telnet capabilities. Also: Use SSH v2 for Dell devices.

2 If you know the device, look at its configuration file and verify that the SNMP community string is correct.

3 Discover the device.

4 If there are any problems with any devices, then ping them, and/or telnet to problem devices and verify that telnet works/authentication is good.

5 If SNMP problems arise, use this application’s MIB browser tool to troubleshoot them.

Toverify SNMP and WMI connections are working between your system and the devices in the network, use the following tools:

SNMP

1 Open MIB Browser in the web client’s Network Tools portlet, or by right clicking the device.

2 Select RFC1213, system, from the RFC Standard Mibsbranch

3 If necessary, fill out the Authentication tab

4 Select the device tab and information will populate as soon as the query is answered.

WMI

If you are discovering WMI systems on your network, the following may be helpful.

1 Launch the `wmiutil.exe`command line tool from `\owareapps\wmi\bin\"`

2 You need to supply a user anda password along with an IP or hostname

Typing `wmiutil.exe` with no arguments returns launch the WMIUtil User Interface.

```
c:\Dorado\owareapps\wmi\bin\wmiutil.exe -user <user> -password <password> -
host <IP or Hostname>
```

Typing `wmiutil.exe ?` at a command line returns what parameters are available for the command line version.

NOTE:

Even if you do not need a domain to log into your WMI device, the graphic interface for this utility does not work if the domain field is blank. Any content makes it work correctly.

See WMI Troubleshooting Procedures on page 671 and Additional WMI Troubleshooting on page 679 for additional details.

Discovery Issues

Discovery may fail if its authentication or network parameters do not match the configuration of devices discovered. Here, the results panel typically displays a message like `No Devices were detected with selected Discovery Parameters`. Use the *Inspect* function in Discovery Profiles to validate credentials entered.

Some additional sources of Discovery issues, and their solutions:

- Equipment with management IP Addresses in the selected subnet, range, and so on does not exist. Correct the selected range and retry.
- The equipment in the selected range has already been discovered.

  Managed devices can only be discovered once. Those devices that have already been discovered appear in the Discovery Results section of the Discovery Wizard. Update the state ofpreviouslydiscovered devices byselecting *Resync* from the right-click menu. If you want to re-discover these devices, delete them from the Managed Resources portlet
- The SNMP community strings/authentication on the equipment do not match the default values used by this application. Correct the SNMP authentication selected for discovery.
- Discovery finds a device, but features like Direct Access, Backup and Actions do not work. **Solution:** The device likely has a correct SNMP authentication, but an incorrect CLI (Telnet/ SSH) authentication. Either re-run the Discovery Profile after deleting the device and correcting the authentications, or right-click to edit the device and add the CLI authentication and management interface, then right-click to resync the device.

  Note that you must log into the device as a user with enough permissions to accomplish all discovery and other tasks. If you log in as a user with limited permissions, then discovery results reflect those limits.

NOTE:

The *Inspect* feature of the Discovery wizard lets you validate authentications.

Alternative: The device is not supported by your current license or driver set. To request support, use the MIB browser to navigate to RCC1213 ifTable details, and export/save this branch. Navigate to ENTITY-MIB entPhysicalTable details and export/save this branch. Attach the exported files to e-mail to support, or to a trouble ticket.

The following describe additional discovery issues.

HTTP Authentication

Often, an HTTP session with devices thatsupport it exchangesdatawith the deviceafter discovery. This process fails if the HTTP Authentication information is inco rrect. Create HTTP authentications that match your devices' in the Authentications portlet and use it in discovery.

Device O/S Overrides

The device driver installed must support the Operation System version on that device. Verify the equipment'sfirmwareand operatingsystemsareamongthosesupported. Supported firmwareand operating systems appear listed in the release notes, or in *Manage > Show Versions* .

Example: Overridedriver-unsupportedoperatingsystemsfor the Juniper devicesin `/owareapps/juniper/lib/juniper.properties.` Change `com.dorado.juniper.supported.OS.dc.default.max`

Thisrevision does not support new features. Other device drivers have similar override mechanisms.

If devicesappear in Managed Resourcesas Discovered Entities, rather than specific vendors' devices. This can mean the following:

- The equipment's driver is not installed.
- The driver installed but not seeded to database. `Workaround:` Run `ocpinstall -s` on a command line.
- Monitored devices must be configured to connect and send SNMP traps to the element management system.

If your system discovers only top-level equipment, this can mean the following:

- Devices do not have components (interfaces, ports, and so on).
- An incorrect telnet/SSH authentication can have an incorrect password or no enable password. **Workaround:** Youcan right-clickand edit the equipment with thisproblem to add the telnet/SSH authentication. Make sure you also add a management interface, then resync the device.

Backup/Restore/Deploy

Failures of backup/restore capabilities often stem from failures in the external FTP/TFTP server. This means the FTP/TFTP server is offline, bloc ked by a firewall or incorrectly configured. Check in the File Server Manager to correct this. Also, on such servers, FTP and TFTP server must share a directory, and the directory must have all permissions to permit these servers to write, read and delete temporary files.

If deploying firmware fails with the following symptoms:

- Selecting *Deploy* does nothing.
- The FTP/TFTP File Server status is Disabled.

  Workaround: Back up the device first to validate it is connected with the FTP server.

When you use the file backup (NetConfig) option, the internal FTP/TFTP server is provided for testing, not production; do not use it. External FTP servers are essential for performance reasons, and, if necessary, the network equipment using FTP to send/receive configuration files must have FTP enabled.

NOTE:

If you have separate FTP and TFTP servers, they must read/write to the same directory.

If deploying firmware succeeds, but device doesn't reboot:

That might be because of "Console Logging" is enabled. Please disable console logging. The messagesfromconsolelogginginterferewith thecommunicationbetween Cruz and the device(via CLI) and can disrupt supported functionality in Cruz.

Timeout

Timeout can occur when backing up /restoring large config files.

Workaround: Change timeout values in the telnet/SSH authentication object. Right-click > Edit the device and change those values in the Authentication tab. Typically this means doubling the timeout and increasing retries to 2 - 3 times.

NOTE:

Secure FTP connections (scp/sftp) often require SSH services be enabled on the devices addressed. Ensure your system's server and sftp/scp file server can access the devices with SSH too.

GroupFile Management Failure

Ifgroup file management backup operations fail for some devices whileindividual backups to these devices are successful, typically thread pool-related backup errors appear in logs during the related time frame.

Solution: Yoursystem may have insufficient threads available to handle the number ofconcurrent tasks required by the group backup operation. Some threads could have already been in use for other tasks when the group operation began.

To address this, increase the size of the thread pool to handle additional concurrent tasks with the following steps:

1. Shutdown application.
2. Edit `owareapps/installprops/lib/installed.properties`
3. Add the following property...

   ```
   ProvisionThreadPoolMBean.PoolSize=70
   ```

   (Adjust as needed based on current setting and need.)
4. Save `installed.properties`.
5. Start application.
6. Execute the group backup operation.

Byincreasing the Pool Size, theapplication can pe rformadditional concurrent tasks that fall within the scope of this pool.

Alarm/Performance/Retention

If you install your system to monitor alarms, and experience sluggish performance or a rapidly filling database, several remedies are available.

- Configure “chatty” devices emitting many alarms to stop doing so.
- Configure your system’s Suppress Alarms feature to keep performance the database’s capacity at acceptable levels
- Reconfigure yoursystem’sdatabase archivingpolicyfeature toarchivealarms more oftenso the database does not fill up.

Retention Policies

Retention policies tune how long your system retains data. These policiesalso have built-in limits (raw, hourly, daily) that help to avoid potential performance/storage problems. The potential impact when going outside these thresholds is significant and generally not recommended.

Configure retention policies with the following limits in mind:

- Maximum # of days to retain daily data: 180
- Maximum # of days to retain hourly data: 14
- Maximum # of days to retain raw data: 1

Network Monitoring

You can monitor your network’s performance two ways.

- Scheduled polling-based monitoring is more re liable, and specific. It also has a lower impact on network. It may, however, lag network events.

- Event-based monitoring (typically Traffic Flow Analysis, syslog and SNMP) is more up to date, but can be less reliable. It also often does not disclose the root cause of a problem.
- This software does not support Traffic flow analysis on sFlows from devices using sFlow earlier than v5. Typical error content reads “Data array too short” if you have an unsupported sFlow version.
- Traffic Flow Analyzer support in Cruz collects and process flows with source and destination IP address. Switches or devices that only support L2 flow payloads with MAC address as the source and destination payload are not currently supported. Example: Juniper XE devices.

NOTE:

Typical packages come with a default limit to the number of monitored devices. Upgrade your license if you want to exceed the package limit. Because of the performance demands they make, Traffic Flow exporters are licensed separately from the managed resource license count, so a license to manage 50 devices does not necessarily let you have 50 Traffic flow exporters.

Also: The application discards IP v6 flow packets.

NOTE:

Does Traffic flow not appear when it’s expected? Have you made sure the device is registered to display traffic flow (right-click in resources *Traffic Flow Analyzer > Register*)?

- Each monitor should have 10,000 or fewer targets. Use a new monitor to track any targets exceeding that number. The general best practice is to have fewer targets distributed across several monitors.

Using this software’s features, you can create alarms and reports for each. Best practice is to use both polling and event-based protocols. Tune the polling frequencyand event granularity for the specific environment, topology, bandwidth, and notification needs. See the specific Monitor performance recommendations.

NOTE:

Creating a baseline performance measurement report of availability, capacity and performance can provide the basis of capacity planning and proactive network management.

Reachability may vary byprotocol (for example, Telnet works, but not ping), so test multiple protocols. If it is remote, try phoning the affected site and asking for information.

Missing Performance Data/Monitor Stops Polling

This is a problem related to missing performance monitor data accompanied by logs errors like the following:

```
2014-02-12 11:23:45,357 705304239 ERROR
  [com.dorado.core.mediation.base.OWMediationDeploymentHelper]
  (WorkManager(2)-63:) The currently deployed targets for polling
  subscription oware.polling.PollingSubscriptionDO::59x8vSybkeEj6I2 on
  mediation partition MED_PART-medPartition have somehow become out of
  synch with the application server and database.
Actual target count, coming from the database: 184 meditation server
  currently has: 0
We are now resynching this information from the application server to
  the mediation server so that it will once again be accurate.
```

Thiserror indicates that your mediation servers have no current performance monitor subscription targets when 184 targets were expected. This could possibly be due to a mediation server fail-over to another cluster member or a mediation server coming back on-line, etc. This is an expected error when a difference is detected in the expected (database) monitor subscriptions and actual subscriptions in the mediation server. A periodic processexecutes to ensure performance monitor subscriptions remain in sync.

Solution(s):

- Verify connectivity between application servers and mediation servers.
- Verify mediation server state/health and cluster member status (active/inactive).
- Verify polling subscriptions in the Monitor Editor.
- Verify polling skip/miss counts in the Monitor Editor.
- Review number of targets in each monitor, verify each has 10,000 or fewer targets. Monitor any targets over that figure in a new monitor.
- Restart the mediation server process if subscription problems persist.

Reports

- I created a report and didn't specify a location. Where's myreport?

  Solution: The default location for reports is `/oware/temp/reportfw` under the installation root.

Report Missing Data

This software limits reports to 5000 rows by default when saving reports to the database (*Save* checkbox checked). This limit does not apply when not saving and only exporting the report. Increasing this default value is not best practice because of potential performance impact.

Solution: If you must increase the size of reports you save, increase the following report-related property values and restart application server(s).

```
com.dorado.redcell.reports.max.report.size (Increase to save larger
  datasets to database - not recommended)
com.dorado.redcell.reports.max.report.query.size (Increase to include
  more data in exported reports)
```

Follow these steps:

1 Edit `[Installation root]owareapps/installprops/lib/installed.properties` and add/modify the desired properties.

```
com.dorado.redcell.reports.max.report.size=<new value>
com.dorado.redcell.reports.max.report.query.size=<new value>
```

2 Restart application server(s).

Web Portal

Web portal problems can occur as described in the following section:

I. Web portal performance is slow or login page inaccessible.

Solution: Check/verify portal memory heap settings and increase as needed.

To manually change the web portal heap settings, verify sufficient system memory exists then modify the setenv.sh (Linux) or setenv.bat (Windows) file:

```
set "PORTAL_PERMGEN=768m"
set "PORTAL_MAX_MEM=4096m"
set "PORTAL_INIT_MEM=4096m"
set "PORTAL_32BIT_MAX_MEM=768m"
```

These files are in the `[Installation root]\oware\synergy\tomcat-x.x.xx\bin directory.`

For Linux, restart the portal service to apply new memory settings. In Windows, besides updating `setenv.bat` you must run `service.bat update` in that same directory.

II. WebPortal Displaying Errors: The application web portal displays errors immediately after starting application processes.

Solution: Allow the application/web server more time to fully start before attempting to access the web portal.

III. Web Portal Down, Cannot Access/Display Login Page: The application web portal displays errors or cannot be accessed.

Solution: Verify the Portal Oracle database password has not expired. By default, `netview` is the default user to connect to database. This appears in `/opt/dorado/oware/synergy/tomcat-[version]/webapps/ROOT/WEB-INF/classes/portal-ext.properties`

```
jdbc.default.username=netview
jdbc.default.password=dorado
```

Connect using SQL*Plus to set new password, you can even use the same password you had earlier.

```
$ sqlplus
SQL*Plus: Release 11.2.0.1.0 Production on Mon Dec 13 01:12:07 2010

Copyright (c) 1982, 2009, Oracle. All rights reserved.

Enter user-name: netview
Enter password:
ERROR:
ORA-28001: the password has expired

Changing password for netview
New password:
Retype new password:
Password changed

Connected to:
Oracle Database 11g Release 11.2.0.1.0 - 64bit Production
```

After resetting the password, best practice is to set profile/policy expiration to LIFETIME to prevent this from expiring again.

NOTE:

Users `netview` and `synadmin` need to have same password.

MySQL Database Issues

See Optimizing Your System on page 112 for preventive `my.cnf` performance tuning tips.

I. Slow Performance: Ifyoursystem'sperformanceslows to theextentapplication is unusable, and its log contains the error below or similar entries:

```
com.mysql.jdbc.CommunicationsException: Communications link failure due
 to underlying exception:
** BEGIN NESTED EXCEPTION **
java.net.SocketException
MESSAGE: Software caused connection abort: recv failed
STACKTRACE:
java.net.SocketException: Software caused connection abort: recv failed
```

Solution: Follow these steps:

1 Review disk space, verify adequate space is available onpartition.

2 Shutdown MySQL database.

3 Edit the `[installation root]\oware3rd\mysql\[version number]\my.cnf` file. Review database size configuration and add another data file to extend size as needed. Save file.

For example: Change:

```
innodb_data_file_path =/ibdata/ibdata1:1024M:autoextend:max:10500M
```

To:

```
innodb_data_file_path = /ibdata/ibdata1:1024M;/disk2/
 ibdata2:1024M:autoextend
```

4 Restart MySQL. Youcan also refer to the following link additional detail: dev.mysql.com/doc/refman/5.1/en/innodb-data- log-reconfiguration.html

II. TablespaceFull/Applicationcrashes. Logentriesindicating tablesor tablespacearefull. Likely accompanied by application crashes. Examples of log entries (which may reflect any table).

```
The table 'rc_notification_hist' is full
```

or

```
InnoDB: Warning: Cannot create table 'owbusdb/pm_dtl_7879' because
 tablespace full
```

Solution: Follow these steps:

1 Review disk space, verify adequate space is available onpartition.

2 Shutdown MySQL database.

3 Open `[Installation root]\oware3rd\mysql\[version]\my.cnf` file. Review database size configuration and add another data file to extend size asneeded.

For example, change

```
innodb_data_file_path =/ibdata/ibdata1:1024M:autoextend:max:10500M
```

To:

```
innodb_data_file_path = /ibdata/ibdata1:1024M;/disk2/
 ibdata2:1024M:autoextend
```

4 Save the file.

5 Restart MySQL. Youcan also refer to the following link additional detail: dev.mysql.com/doc/refman/5.1/en/innodb-data- log-reconfiguration.html

III. MySQL Connection Exceptions: The following error, or similar errors, appear in log:

```
Caused by: com.mysql.jdbc.CommunicationsException: Communications link
  failure due to underlying exception:
BEGIN NESTED EXCEPTION **
java.net.NoRouteToHostException
MESSAGE: No route to host
STACKTRACE:
java.net.NoRouteToHostException: No route to host at
  java.net.PlainSocketImpl.socketConnect(Native Method)
```

Solution: This indicates a connectivity issue between your application server and the database. Discover the root cause of this communication issue and correct it. Here are some things to try:

- Check with `ps -ef | grep MySQL` in Linux or in Windows' Services utility to make sure your database is running. If not, re-install (or uninstall/reinstall) until this daemon/service starts without problems.
- Execute pingdb to test database connectivity.
- Check network interfaces and connectivity between application server and database.
- Try connecting with MySQL Workbench or another tool.
- Verify database is up and healthy.
- Verify database login/password has not changed or expired (default user/password: root/dorado)

IV "Too Many Connections" Error. Ironically, this error may indicate the `max_connections` parameter in your `my.cnf` files is too *small*. To use more connections, change the setting in the `my.cnf` file. For example, in `\oware3rd\mysql\X_X_X-64\my.cnf`, under the `[mysqld]` section add:

```
max_connections = 500
```

You can login to mysql to check current settings:

```
mysql -u root --password=dorado
mysql> show variables like 'max_connections';
```

To check open connections

```
mysql> SHOW STATUS WHERE `variable_name` = 'Threads_connected';
```

Oracle Database Issues

Best practice for Oracle database users is to have a database administrator configures their Oracle application before installing Cruz to use an Oracle database. This ensures correct configuration, best performance and connection with the database.

Patches

You can patch OCPs or add them new (for new ap plications, drivers, or features). If you are installing onlya singlecomponent, rather than a re-installation with the full installation wizard, installation is a three-step process:

1. Extract (`ocpinstall -x [filename(s)]`)
2. Optionally, update/verify the database schema (on database servers only). (`ocpinstall -l [filename(s)]`)

3 Seed the database (on database servers only). (`ocpinstall -s [filename(s)]`)

NOTE:

Whenusinghostnameasoracleconnection URL, sometimeinstaller mayappend colontotheURL after theupgrade.Workaround: remove thecolonfrom the URLintheinstalled.properties andportal-ext.properties.

Refer to the *Installation Guide* for more about Oracle databases.

If you are adding jars for patches, turn applicatio ns off, rename and move the old jar before you copy in a new jar. Then delete the contents of `oware/temp`.

Debug

When an error appears in logs (see Logs on page 670) it indicates for which Java class you need to increase the level of logging if you want debugging information. Youcan change logging levels to get additional (debug) information.

NOTE:

Best practice is to now alter *Log Categories* in the Application Server Statistics portlet by clicking that button. This alteration simplifies editing `log4j.xml` files since it provides a graphical interface, and if you have more than one server, it alters the log levels for all servers at once.

Flipdebug

You can easily turn debug on or off with the `flipdebug` script. Run this while the application server is running and remember to wait a few moments for the application server to pick up your changes. Here is the usage (just type the script name to see these options):

```
Usage: flipdebug [-d] [-t] [-r] [-h|-?] [-p] product[,product]


-d Turn on debug for all packages
-t Turn on trace for all packages
-p Turn on trace/debug only for product[,product] (no spaces between
  products)
-r Revert to original log4j settings
-h|-? Display usage
```

The *product* name in -p matches the directory under `owareapps`. The debug and trace write to the log and to stdout.

Fine-Tuning Debug

To fine tune debug further, create a file, whose name ends in `log4j.xml` in the `owareapps\installprops\server\conf` directory with the categories you want changed. If the class does not exist within the `log4j` file, add it and set it to debug. Changes preserved in such a file remain in place until you change them again and are not overwritten when software upgrades occur.

To increase logging levels to DEBUG, change `WARN` or `INFO` to `DEBUG` in categories like the following:

```
- <category name="com.dorado.redcell">
-     <priority value="WARN"/>
-   </category>
```

```
...
-  <category name="redcell">
-    <priority value="WARN"/>
-  </category>

-  <category name="RedCell">
-    <priority value="INFO"/>
-  </category>
```

To see what categories are available, look in `\oware\jboss-x.x.x\server\oware\conf\log4j.xml`. Thisfileconcatenatesallloggingcategories, but is generated, and should *not* be changed.

When application server starts, it detects logging levels in these categories and concatenates them into theserver's`log4j.xml`from`*log4j.xml`filesinthe `server\conf`directoriesofinstalled components under `owareapps`.

When it starts, application server processes loggi ng for components in order of their dependency and overrides any detected settings from a file whose name ends in `log4j.xml` in the `installprops` directory.

This application applies detected changes once a minute. The log4j file scanner can then detect any subsequent changes up to a minute after making them. The `server.log` is not truncated when this occurs.

The following are additional categories that allow logging level changes:

For Mediation Server registration with App Server, add the following category:

```
<category name="com.dorado.mbeans.OWMedServerTrackerMBean">
  <priority value="WARN"/>
</category>
```

For SNMP and Syslog, change INFO to DEBUG in

```
<category name="com.dorado.core.mediation">
  <priority value="INFO"/>
</category>
```

Toviewdebug output:

Server

Debug does not appear in real-time in the applicat ion server shell (if you have one). View real-time and historical logs in the `oware\jboss-x.x.x\server\oware\log` directory.

NOTE:

Typing `oware` inaWindowsshellletsyouuseLinuxcommandslike`tail -f server.log`.Tailing it lets you watch the log file as it is generated.

Resolving Port Conflicts

Installation scansfor port conflicts, but these may arise after you install too. Ifyour applicationruns with others, conflicts related to those other applications' ports are possible. For example: the application can have trouble communicating with the built-in TFTP server for backups. Port contention of TFTP on UDP port 69 with other applications can cause this. Try rebooting the system to clear any unused ports and verify that UDP port 69 is not in use before starting the application.

Finding Port Conflicts

You can find ports in use with the following command line:

```
netstat -a -b -o | findstr [port number]
```

Use this command to track down port conflicts if in stallation reports one. Best practice is to run this software on its own machine to avoid such co nflicts. Freeware port conflict finding programs like Active Ports are also available.

Logs

You can execute the `getlogs` script to package relevant logs if you need technical support. Run this script in a command shell where you have sourced the Oware environment (in Windows *Start > Run* `cmd,` then run `oware`, or in Linux `. /etc/.dsienv`, and then invoke the `getlogs` script). This script creates a `logs.jar` file in the root installation directory, and moves any existing copy of `logs.jar` to `oware\temp`. This jar compresses all logs necessary for troubleshooting. Read the jar yourself, or forward this file to technical support to help troubleshoot problems.

NOTE:

Server logs are in oware\jboss-x.x.x\server\oware\log. Also: If you install with an Oracle database, because the Oracle installation is outside Cruz's installer, Cruz does not create db_setup.log.

The `getlogs` script also gathers the tomcat web server logs for the web portal.

If `getlogs` halts, and does not produce a `logs.jar`file, it may be because someone installed this software, or a previous version as root, so `getlogs` does not have access to directories and groups owned byroot. Change the permissions and/ orownership of those groups and directories to make `getlogs` work.

Increasing Startup Logging

Forapplications based on Oware 6.2.1 and later (see your `versions.txt` file or *Manage > Show Versions* for your version), you can add the following line to `oware\lib\pmstartup.dat`. If you add this line, this software logs all output from the startappserver script to a file:

```
server.out.filename=/opt/dorado/pmserver.log
```

The destination you specify can be any valid path and file name. This helps when the server never starts or errors occur during deployment that would not be in the usual server.log.

Tuning Log Retention

The following properties are in `appserver.properties` and `redcell.properties` file for purging log files. As with all other properties, best practice is to override them in `owareapps/installprops/installed.properties.`

```
owappserver.properties
# This property defines how many days to retain server log files. All log
  files
# older than the specified retention days are purged. Back up older log
  files if
# you want to retain them. Set the property to -1 to disable this option.
  The
# default is 7.
oware.server.log.files.retention.days=7
```

```
redcell.properties
# This property defines where redcell client log files are stored The
# following is also the default:
redcell.log.files.location=oware.user.root/owareapps/redcell/logs

# This property defines how many days to retain the client's
# log files. Files older than the specified age are purged.
# Setting the property to a negative value disables log file deletion.
# The default is 7.
redcell.log.files.retention.days=7
```

Log Generation Fails with “Build Failed” Error (Linux)

Log generation and the build processfails when attempting to generate logsand isaccompanied by an error like this:

```
BUILD FAILED
/opt/dorado/oware/conf/owrtbuild.xml:46: Problem creating jar: /opt/
  dorado/logs/ocpinstall_14332.log (Permission denied)
```

To successfully build logs, included files must be owned by the installing user (example: MyUser).

Solution: Locate and change the ownership of file(s) breaking the build process. To repair these, follow these steps:

1 Open a shell and source environment by typing `. /etc/.dsienv`.

2 Type `getlogs`.

3 Navigate to the file location(s) that appear in any error.

4 Type `ls -l` and review the owners for the files in this directory.

5 Type `su root` and enter root password.

6 Change the file ownership from user `root` (or other) to the installing user.

Type, forexample: chown `MyUser:MyUser ocpinstall_*` to changeownership ofall files beginning with `ocpinstall_` .

7 Type `ls -l` to confirm new file ownership.

8 Type `exit` to return to the installing user prompt.

9 Repeat this process until getlogs builds a `logs.jar` successfully without error. Youmay need to correct file ownership in several lo cations before a successful build can occur.

To avoid a reoccurrence, do not perform anyapplication-related command line operations while logged in as `root`. Such tasks must always be executed by the installing user.

WMI Troubleshooting Procedures

Thefollowingsections describe troubleshooting common WMI problems. To monitor with WMI, the following must be true:

- WMI must be enabled on the remote, monitored server and functioning properly.
- The remote server must be accessible through a Remote Procedure Call (RPC) connection to run WMI queries.

If your system does not meet these conditions WMI displays an *Unknown* status.

Examples of what may prevent the above can include the following:

- Not having local Administrator rights on the monitored host.
- Firewalls blocking the WMI traffic.
- An operating system not configured for WMI.
- An error in the credential password.

To help diagnose these issues, test the WMI services, the remote WMI connections, and you system's component configuration.

The following topics provide troubleshooting assistance:

- WMI Troubleshooting on the local host.
- Testing Remote WMI Connectivity
- Verify Administrator Credentials
- Enable Remote Procedure Call (RPC)
- Configure Distributed Component Object Model (DCOM) and User Account Control (UAC)
- Add a Windows Firewall Exception for Remote WMI Connections
- Checking the Authentication portlet to ensure correct credentials exist.

Finally, if these troubleshooting tips are not enough, see Additional WMI Troubleshooting on page 679.

WMI and Operating Systems

Best practice is to avoid using Windows Vistaand Windows 2008 for network monitoring of WMI. WMI works well Windows 7, even for larger networks. But with Vista and Window 2008 this is not true. Some tests even indicate that Windows 7 is up to 70 times faster than Windows 2008 or Vista. In these tests, hardware (CPU, memory) does not strongly affect WMI monitoring performance, nor does virtualization.

WMI Troubleshooting

To troubleshoot WMI, do the following:

1. To test WMI locally, click *Start* > *Run*, then enter `wbemtest.exe` and then click *OK*. The `wbemtest.exe` program comes with Windows that supports WMI.
2. Click *Connect* on the Windows Management Instrumentation Tester window.

3 Enter `\root\cimv2` in the field at the top of the dialog box next to the *Connect*button.

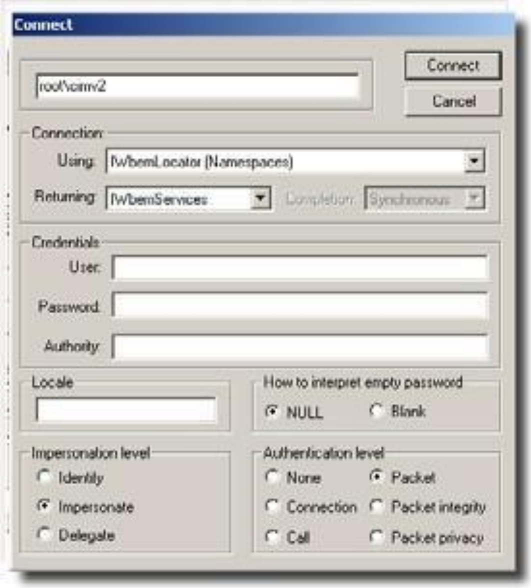

Connect

root\cimv2

Connection

Using: IWbemLocator (Namespaces)

Returning: IWbemServices

Completion: Synchronous

Credentials

User:

Password:

Authority:

Locale

How to interpret empty password: ● NULL ○ Blank

Impersonation level: ○ Identify ● Impersonate ○ Delegate

Authentication level: ○ None ● Packet ○ Connection ○ Packet integrity ○ Call ○ Packet privacy

4 Click *Connect*.

5 Click the *Enum Classes* button.

6 Select the *Recursive* radio button. Leave the superclass name blank, and then click *OK*.

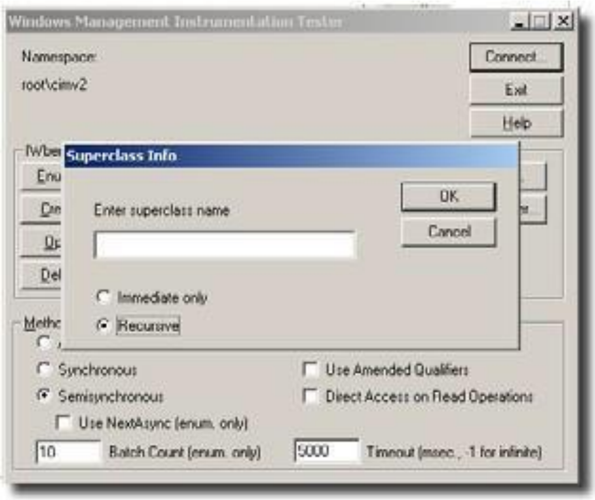

Windows Management Instrumentation Tester

Namespace:
root\cimv2

Superclass Info

Enter superclass name

○ Immediate only
● Recursive

Method
○ Synchronous
● Semisynchronous
☐ Use NextAsync (enum. only)
10 Batch Count (enum. only)
☐ Use Amended Qualifiers
☐ Direct Access on Read Operations
5000 Timeout (msec., -1 for infinite)

7. If the WMI classes you are querying appear in this list, local WMI services are functioning correctly.

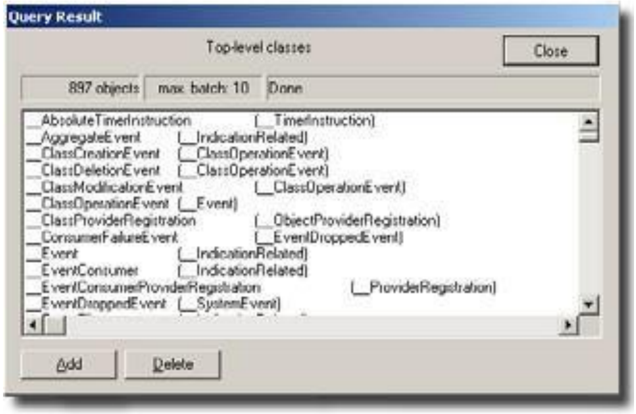

```
__AbsoluteTimerInstruction          (__TimerInstruction)
__AggregateEvent       (__IndicationRelated)
__ClassCreationEvent   (__ClassOperationEvent)
__ClassDeletionEvent   (__ClassOperationEvent)
__ClassModificationEvent            (__ClassOperationEvent)
__ClassOperationEvent  (__Event)
__ClassProviderRegistration         (__ObjectProviderRegistration)
__ConsumerFailureEvent              (__EventDroppedEvent)
__Event                (__IndicationRelated)
__EventConsumer        (__IndicationRelated)
__EventConsumerProviderRegistration            (__ProviderRegistration)
__EventDroppedEvent    (__SystemEvent)
```

8. If the list does not appear or does not contain the desired WMI class, WMI is not functioning correctly. Continue reading this section for guidance on repairing WMI services on the target server.
9. Click *Exit*.
10. If you did not see the desired classes, Reset the WMI Counters , and re-test until those classes appear.

Reset the WMICounters

At times, the WMI performance counters may not get transferred to WMI because services were delayed or started out of order (see support.microsoft.com/kb/820847).

To manually reset the WMI counters:

1. Stopthe WindowsManagement Instrumentation (WMI)service.
2. Open a shell (Click *Start* > *Run*, type `cmd`, and then click *OK*).
3. At the command prompt, type `winmgmt /resyncperf`, and then press [Enter].
4. At the command prompt, type `wmiadap.exe /f`, and then press [Enter].
5. Type `exit`, and then press [Enter] to close the command shell.
6. Restart the WMI service.

After resetting the WMI counters, retest WMI. If resetting the WMI counters did not solve your problem, see “WMI is Still Not Working, Now What?” on page 12.

Testing Remote WMI Connectivity

Ensure WMI is running on the remote, monitored host. This is similar to WMI Troubleshooting on the local host described above.

1. Click `Start > Run`, enter `wbemtest.exe` and then click *OK*.
2. Click *Connect* on the WMI Tester window.
3. Enter `\\Target_Primary_IP_Address\root\cimv2` in the field at the top of the dialog box. Replace *Target_Primary_IP_Address* in this entry with the actual host name or primary IP address of the target server.

4 Enter the appropriate administrator user name in the User field, the password in the Password field, and `NTLMDOMAIN:NameOfDomain`inthe Authorityfield. Replace *NameOfDomain* with the domain of the user account specified in the User field.

5 Click *Connect*.

6 Click *Enum Classes.*

7 Select the Recursive radio button without entering a superclass name, and then click *OK*.

8 If the WMI class list appears, remote WMI is functioning correctly. Skip to the next topic and validate your credentials.

9 If the list does not appear, remote WMI is not functioning correctly. Continue reading this topic for guidance on restoring remote WMI connections on the target server, and retest remote WMI after completing each troubleshooting step.

11. Click the *Close* button, and then click *Exit*.

Verify Administrator Credentials

Onlyacredential that hasadministrator rights on the target server has the necessary permissions to access the target host's WMI services. Make sure that the username and password you are using belongs to an administrator on the target host.

If the administrator credential is a domain member, be sure to specify both the username and the domain in standard Microsoft syntax. For example: `DOMAIN\Administrator`.

Enable Remote Procedure Call (RPC)

Remote WMI connections use RPC as a communications interface. If the RPC service is disabled on the target server, remote WMI connections cannot beestablished.

These steps show how to enable the RPC service:

1 Log on to the target host as an administrator.

2 Click *Start* > *Run*, then type `services.msc`, and then press [Enter].

3 Right-click Remote Procedure Call (RPC), and then click *Start* on the shortcut menu.

Configure Distributed Component Object Model (DCOM) and User Account Control (UAC)

If the target computer is running Windows Vista or Windows Server 2008, you may be required to make settings changes to allow remote WMI requests (See msdn.microsoft.com/enus/library/aa822854(VS.85).aspx).

DCOM — Edit Default and Limits permissions to allow the following actions:

- Local launch (default permission)
- Remote launch (default permission)
- Local activation (limits permission)
- Remote activation (limits permission)

For more information, see Enabling DCOM on page 676.

WMI Namespaces—Modify the CIMV2 security to enable and remote enable the account used to access the server or workstation through WMI. You must ensure the security change applies to the current namespace and subnamespaces. For more information, see Enabling Account Privileges in WMI on page 677.

User Account Control—Remote UAC access token filtering must be disabled when monitoring within a workgroup environment. For more information, see Disabling Remote User Account Control for Workgroups on page 678.

Enabling DCOM

WMI uses DCOM to communicate with monitored target computers. Therefore, for Application Performance Monitorto use WMI, you musthave DCOM enabledandproperlyconfigured. Follow these steps to enable DCOM permissions for your Application Performance Monitor credentials:

1 Log on to the target host as an administrator.

2 Navigate to *Start* > *Control Panel* > *Administrative Tools* > *Component Services.* (Only the Classic view of the Control Panel has this navigation path). Youcan also launch this console by double-clicking `comexp.msc` in the `/windows/system32` directory.

3 Expand *Component Services* > *Computers.*

4 Right-click *My Computer*, and then select *Properties*.

5 Select the COM Security tab, and then click *Edit Limits* in the Access Permissions grouping.

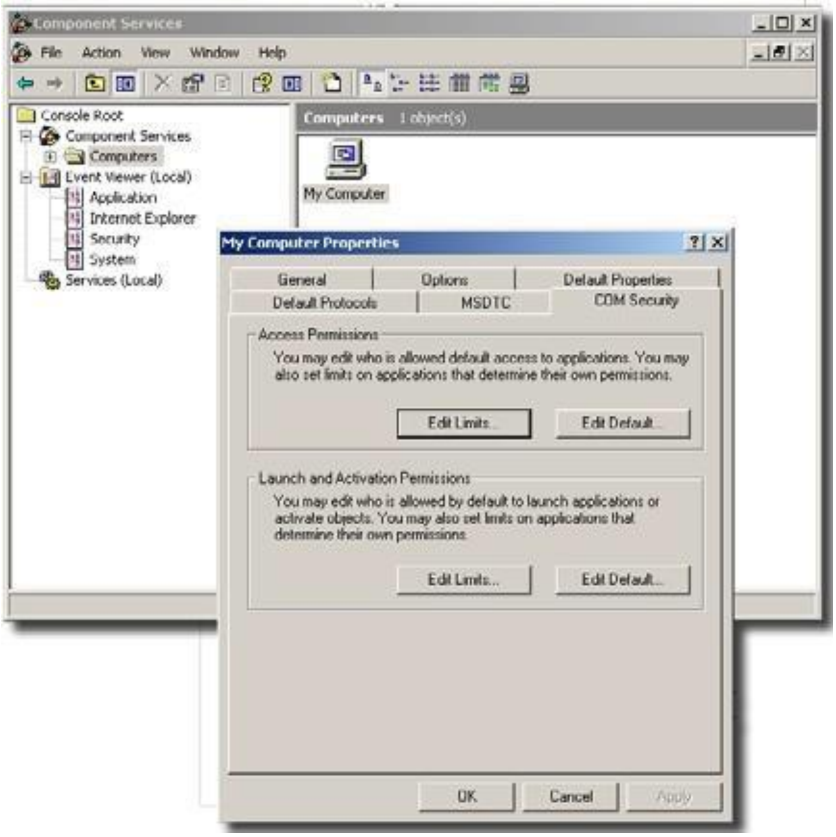

You may edit who is allowed default access to applications. You may also set limits on applications that determine their own permissions.

You may edit who is allowed by default to launch applications or activate objects. You may also set limits on applications that determine their own permissions.

6 Ensure the user account collecting WMI statistics has *Local Access* and *Remote Access*, and then click *OK*.

7 Click *Edit Default*,and then confirm the user account collecting WMI statistics has *Local Access* and *Remote Access*, then click *OK*.

8 Click *Edit Limits* in the Launch and Activation Permissions grouping.

9 Ensure the user account collecting WMI statistics has *Local Launch*, *Remote Launch*, *Local Activation*, and *Remote Activation* enabled, and then click *OK*.

10 Click *Edit Default*,and then ensure the user account collecting WMI statistics has *Local Launch*, *Remote Launch*, *Local Activation*, and *Remote Activation* enabled.

11 Click *OK*.

Enabling Account Privileges in WMI

The account you specify for authentication must possess security access to the namespace and subnamespaces of any monitored target hosts. To enable these privileges, complete the following procedure.

To enable namespace and subnamespaces privileges:

1 Log on to the host you are monitoring as an administrator.

2 Navigate to *Start* > *Control Panel* > *Administrative Tools* > *Computer Management* > *Services and Applications*. (Classic View of the Control Panel this navigation path).

3 Click WMI Control, and then right-click and select Properties.

4 Select the Security tab, and then expand Root and click CIMV2.

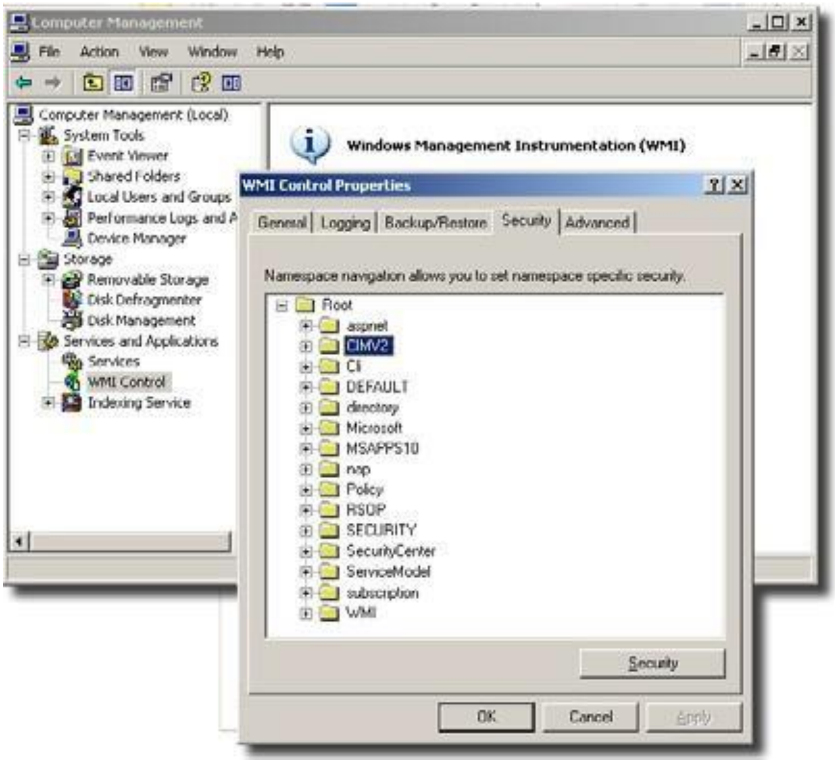

Namespace navigation allows you to set namespace specific security.

5 Click Security and then select the user account used to access this computer and grant the following permissions:

–Enable Account

–Remote Enable

6 Click *Advanced*, and then select the user account that accesses this computer.

7 Click *Edit*, select *This namespace and subnamespaces* in the *Apply to* field, and then click *OK*.

8 Click *OK* on the *Advanced Security Settings for CIMV2* window.

9 Click *OK* on the *Security for Root\CIMV2* window.

10 Click *Services* in the left navigation pane of *Computer Management*.

11 Select *Windows Management Instrumentation* in the *Services* result pane, and then click *Restart*.

Disabling Remote User Account Control for Workgroups

If you are monitoring a target in a workgroup, you must disable remote User Account Control (UAC). Althoughthis notrecommended, it is necessarywhenmonitoringaworkgroupcomputer. Disabling remote user account control does not disable local user account control functionality.

△ CAUTION:

The following modifies or creates aregistry key. Changing the registry can have adverse effects on your computerand mayresult in anunbootablesystem. Bestpracticeis to backupyourregistry before making these changes.

To disable remote UAC for a workgroup computer:

1 Log on to the host you want to monitor as an administrator.

2 Click Start > Run and enter `regedit`.

3 Expand `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System`.

4 Locate or create a DWORD entry named LocalAccountTokenFilterPolicyand provide a DWORD value of 1.

NOTE:

To re-enable remote UAC, change the DWORD value to 0.

Add a Windows Firewall Exception for Remote WMI Connections

If the target computer has Windows Firewall enabled, it must have a Remote WMI exception to allow remote WMI traffic through (See /msdn.microsoft.com/en-us/library/aa389286(VS.85).aspx). Follow these steps to add this exception:

1 Open a command shell (Click *Start* > *Run*, type `cmd`and then press [Enter]).

2 At the command prompt, type

```
netsh firewall set service RemoteAdmin enable
```

Press [Enter]

3 Type `exit` then press [Enter].

If adding the firewall exception did not solve your problem, see Additional WMI Troubleshooting below.

WMI Authentication

If the above troubleshooting has been done correctly, ensure the authentication credentials for the WMI device in *Resources* match those for an administrator. Select the device in the Resources screen and click *action* > *Open* (or right-click and select *Open*). Go to the *Authentication* node of the tree and confirm that the correct authentication objects appear there.

Additional WMI Troubleshooting

The above discusses the most common errors behind WMI failures. After trying these, if you are unable to get WMI services working, consult the following articles about this subject on Microsoft's Technet and Developer Networks:

- *WMI Isn't Working!: Troubleshooting Problems with WMI Scriptsand the WMI Service.* (See www.microsoft.com/technet/scriptcenter/topics/help/wmi.mspx)
- *WMI Diagnosis Utility: A New Utility for Diagnosing and Repairing Problems with the WMI Service* (See www.microsoft.com/technet/scriptcenter/topics/help/wmidiag.mspx)
- *WMI Troubleshooting* (See msdn.microsoft.com/enus/library/aa394603.aspx)

NOTE:

While the above URLs are believed correct, they may change.

jstack Debugging in Windows 7

Technical assistance sometimes uses the jstack stack trace tool to debug problems in this software. When you install application server to run as a service (autostart), the user SYSTEM owns the application server process. Youalso cannot log into Windows 7 as user SYSTEM. For security's sake, no other user can access a service running as the SYSTEM user in later Windows 7 kernels. The jstack tool therefore does not work if you run it as the (non-SYSTEM) user logged in to Windows 7.

Workaround: Toview a jstack output for application server, orany service (and its subprocesses) running as the user SYSTEM, you must run jstack (and jps) as user SYSTEM. Windows 7 provides no direct way to log in as user SYSTEM, so the following sidesteps this prohibition:

1. Open a command shell (Click the *Start* icon and type `cmd` in the *Search Programs and Files* field.)
2. At the command prompt, type:

   ```
   sc create testsvc binpath= "cmd /K start" type= own type= interact
   ```
3. Then type:

   ```
   sc start testsvc
   ```

   The `sc start` command immediately creates a newcommand shell owned by SYSTEM, even if the original command window failed to start with error 1053 (this is expected since `cmd.exe`does not have any service-related code in it).
4. Open an oware shell inside the SYSTEM-owned command prompt created in Step 3. (Type `oware` at the command line.)
5. In that oware shell, run `jps` to see the process ID (PID) of the application server's Java process (`OWLaunchV2`).
6. Then run `jstack [PID]` in the SYSTEM shell.

7 To delete the testsvc when you are finished, type this on a command line:

```
sc delete testsvc
```

FAQs about Monitoring Mediation Servers

After makinga UDP-based JGroups discoveryrequest and receivinga response froman application server in the cluster, each mediation server makes an RMI (TCP) call to an application server every 30 seconds. This RMI call results in a “call on cluster” on the application server cluster, using JGroups (UDP by default), to call the agentHeartbeat method of the OWMedServerTrackerMBean on each application server in the cluster. The primary application server updates the timestamp for the medserver in question, and the others ignore the call. Every five seconds, the primary application server checks to see if it has not received a call from a mediation server in the last 52 seconds. If it has not, it attempts to verify down status by pinging the suspected mediation server. Then it issues an RMI call on that mediation server. It considers the meditation server down if the ping or the final RMI call fails. This avoids false meditation server down notifications when a network cable is pulled from an application server.

- Does the application server wait 15 seconds after receiving the mediation server's response? Or does it monitor mediation server every 15 seconds regardless of the mediation server's response?

  The receipt of the mediation server's RMI call is on a different thread than the monitoring code. The monitoring code should run every 5 seconds, regardless of the frequency of mediation server calls. However, after investigating the scheduling mechanism used (the JBoss scheduler - http://community.jboss.org/wiki/scheduler), it is possible that other tasks using this scheduler could impact the schedule because of a change in the JDK timer implementation after JDK 1.4.

- What kind of functionality (JMS?) does application server use to send and receive Cruz messages?

  The application server does not actively monito r the mediation servers unless it fails to get a call from one for 52 seconds. If it does try to verify a downed mediation server, it uses an RMI call.

  The RMI calls use TCP sockets. It may use multiple ports: 1103/1123 (UDP - JGroups Discovery), 4445/4446 (TCP - RMI Object), 1098/1099 (TCP - JNDI), or 3100/3200 (TCP - HAJNDI), 8093 (UIL2).

- What kind of problem or bug would it make application server to falsely detect a mediation server down? Forexample, would failing to allocate memory cause application server to think a mediation server is down (dead)?

  An out of memory error on an application server could result in a false detection of a downed medserver.

- If such memory depletion occurs as described in the previous answer, would the record appears in the log? If it doesn't appear in the log, would it possibly appear if the log-level is changed?

  An out of memory error usuallyappears in the log without modifying logging configuration, since it is logged at ERROR level.

- The log shows that a mediation server was detached from the cluster configuration, but what kind of logic is used to decide the detachment from the cluster? For instance, would it detach application servers if they detect the mediation server down?

JBoss (JGroups) has a somewhat complex mechanism for detecting a slowserver in a cluster, which can result in a server being “shunned.” This logic remains, even though we have never observed the shunning of a server resulting in a workable cluster. This is the only mechanism which automates removing servers from the cluster. The configuration for this service is located in `$OWARE_USER_ROOT/oware/jboss-x.x.x/owareconf/cluster-service.xml`. Shunning can be disabled byreplacing all `shun='true'` instances with `shun="false"`. A flow control option also exists which regulates the rate of cluster communication to compensate for one server being slower in processing cluster requests than another. The detection of a mediation server being down with the heartbeat mechanism described here does not attempt to remove the medserver from its cluster.

- Why does Mediation server not appear in the control panel?
  Make sure you have followed the instructions in the *Installation Guide.*

Linux Issues

The following are issues with Linux installations:

- Install in `/opt/dorado`, unzip package in, for example, `/opt/installs`, not `/opt/dorado`.
- Linux (executed as the root user) uses this command:

  ```
  /etc/init.d/owaredb start
  ```

  You should see the following response in the shell where you execute this command:

  ```
  Starting MySQL[ OK ]
  ```

- If you experience problems with discovery, and see errors on startup similar to the following:

  ```
  [com.dorado.core.mediation.snmp.SRSnmpEventReportDispatcher] (Thread-36
   RecvTrap Exception :
  com.dorado.core.mediation.snmp.SRSnmpException:
  at com.dorado.core.mediation.snmp.SRSnmpSession.nRecvTrap(Native Method)
  at
   com.dorado.core.mediation.snmp.SRSnmpSession.recvTrap(SRSnmpSession.jav
   a:733)
  at
   com.dorado.core.mediation.snmp.SRSnmpEventReportDispatcher.run(SRSnmpEv
   entReportDispatcher.java:96)
  at java.lang.Thread.run(Thread.java:662)
  ```

  or

  ```
  ERROR [com.dorado.core.mediation.syslog.OWSysLogListener]
   (OWSysLog.Listener Received a null SysLog message. SysLog port may be in
   use. Shutting down SysLog listener.
  ```

  You may be able to solve this issue by increasing the available memory on the entire system or lowering the heap memory used by your system. The former option is best practice.

Application Server Memory (Linux and Windows)

I. Linux application server appears to have low memory.

Solution: Memory statistics using TOP can be deceiving. Linux may have borrowed some free memory for disk caching. To determine if this is the case:

1 Open a shell and execute command: `free -m`

Thisreturns theamount of(true) free/available memoryforapplication usein megabytes. See cache value in line `-/+ buffers/cache: 26441 37973` below…

```
[redcell@AppRedcell01 ~]$ free -m
total used free shared buffers cached
Mem: 64414 63823 590 0 364 37018
-/+ buffers/cache: 26441 37973
Swap: 65535 11 65524
[redcell@AppRedcell01 ~]$
```

Here, 37,973M is still free for application use.

See www.linuxatemyram.com/ for more detail on this topic.

Alternatively,

2. If TOP reveals an excessive and abnormally high memory usage for the application java process, you may need to restart your application server and evaluate installed/available memory with regard to your sizing and application usage requirements. Install more server memory as needed.

II. Genuine memory issues appear if logs contain an error like `java.lang.OutOfMemoryError: GC overhead limit exceeded`, and application server performance is slow, possiblypreventing logintoweb portal. Youmayalsosee many other errors, for example, from performance monitoring:

```
WARN [com.dorado.broadscope.polling.PollingResultsDAOImpl]
  (WorkManager(2)-99:) Low on memory. Discarding this batch.
```

Solution: These errors indicate memory resources are low or have been depleted.

To address this, first review any potential causes for an increase in memory usage. For example, has there been a significant increase in performance monitor load, perhaps from reducing polling times or an increase in targets/attributes? Have there been any other changes?

Assumingenoughserver memory is available, increase heap size. Adjust memoryvaluesbelow according to your environment and configuration needs:

1. Shut down the application.
2. Open `owareapps/installprops/lib/installed.properties` file for editing.
3. Modify the `oware.server.max.heap.size`property

   ```
   oware.server.max.heap.size=3072m
   ```

   In this example, a recommended increase would be 25% to 4096m.
4. Increase `oware.server.min.heap.size` to match (4096m)
5. Save changes to `installed.properties`.
6. Restart the application.

NOTE:

Heap adjustments workfor Windows too.

III. Out of Memory errors like `Out of Memory: unable to create new native thread`inlogsforserver (applicationor mediation) mayindicate memoryresourcesareenough, but threads are not.

Solution: The operating system may be limiting the number of available threads for use by the application. Check/modify the `ulimits` settings with thesesteps:

1 Open shell/CLI and type `ulimits -a`.

Open files and User Processes should not be set to typical defaults (1024). Change these with the next steps.

2 Open `/etc/security/ limits.conf` for editing.

3 Add the following lines and save.

```
<installing user> soft nofile 65536
<installing user> hard nofile 65536
<installing user> soft nproc 65536
<installing user> hard nproc 65536
```

4 Restart the application processes.

Refer to the *Cruz Installation Guide* for Linux installation and upgrade instructions and best practices.

Linux syslog notdisplaying

Application does not display syslog messages.

On Linux based platforms, under certain circumstances, a race condition at application startup may impact syslog event/messaging functionality. If syslog messages are not displaying as expected, please apply the following **workaround** to restore functionality.

1) Shutdown the webserver.

2) Restart the appserver.

3) Start the webserver after the appserver'sstatus shows 'ready'.

This process may need to be repeated if the server is restarted.

HA installation issue onLinux

Servers within HA system will fail on startup caused by the following exception:

```
ERROR [org.jgroups.protocols.UDP] (Timer-1,192.168.54.41:37377:) failed sending message to null (69 bytes)
```

The recommended workaround is to modify a line with the run.sh script found in install_path/ oware/jboss-5.1/bin

From

```
JAVA_OPTS ="$JAVA_OPTS -Djava.net.preferIPv4Stack = false"
```

To

```
JAVA_OPTS ="$JAVA_OPTS -Djava.net.preferIPv4Stack = true"
```

Servers restart is required.

Linux HA does not support IPv6 as default.

While IPv6 is supported on Windows HA, Linux HA does not support IPv6 as default. To acquire IPv6 on Linux HA, it's suggested that users must follow these steps enable unicast within the Mediation cluster. Apply the configuration changes to all Mediation servers.

1 Add the property oware.unicast=true to installed.properties file located in .../dorado/ owareapps/installprops/lib directory.

2. Locate .../oware/jboss/server/oware/deploy/ cluster/jgroupschannelfactory.sar/META-INF/ jgroups-channelfactory-stacks.xml.
3. In the TCP section (you can search by < stack name ="tcp"), comment this portion:

```
<!--Alternative 1: multicast-based automatic discovery. -- >
< MPING timeout ="3000"
num_initial_members = "3"
mcast_addr ="${jboss.partition.udpGroup:230.11.11.11}"
mcast_port ="${jgroups.tcp.mping_mcast_port:45700}"
ip_ttl ="${jgroups.udp.ip_ttl:2}"/ >
```

4. And Uncomment

```
<!-- Alternative 2: non multicast-basedreplacementfor MPING. Requires
a static configuration of a ll possible cluster members. >
<TCPPING timeout ="3000"
initial_hosts ="${jgroups.tcpping.initial_hosts:localhost[7600],localhost[7601]}"
port_range = "1"
num_initial_members = "3"/-- >
```

△ CAUTION:

Make sure you modify stack "tcp" section, not "tcp-sync" section

Example:

```
<!--Alternative 1: multicast-based automatic discovery.
< MPING timeout ="3000"
num_initial_members = "3"
mcast_addr ="${jboss.partition.udpGroup:230.11.11.11}"
mcast_port ="${jgroups.tcp.mping_mcast_port:45700}"
ip_ttl ="${jgroups.udp.ip_ttl:2}"/ >
-- >
Alternative 2: nonmulticast-basedreplacementfor MPING. Requires
a static configuration of a ll possible cluster members. >
<TCPPING timeout ="3000"
initial_hosts ="${jgroups.tcpping.initial_hosts:10.35.35.200[7600],10.35.35.201[7601]}"
port_range = "1"
num_initial_members = "3"/
```

Where 10.35.35.200 and 10.35.35.201 are IPaddresses of Mediation servers.

5. Restart Mediation servers. (# service oware stop/start)

Device Prerequisites

Often, devicesrequirepre-configuration before theyaremanage-able bythissoftware. Forexample, the management system application server must have access to the device, and often must be listed on the access control list for the managed device.

Common Device Prerequisites

The following are common prerequisites:

Credentials—WBEM credentials have a role in discovering the device. Yoursystem must have access to the computer using Administrative only credentials. These are the same credentials as the user installing WBEM on the device.

Telnet/SSH credentials are necessary for other supported applications.

Forfull functionality, this WBEM device driver requiresadministrative (root) access. Many devices may only allow root logins ona local console.

Insuch cases, configure the Telnet/SSHauthentication for these devices to loginasa non-root user—and, in Authentication Manager, enter `su` in the *Enable User ID* field and enter the root user's password in *Enable User Password* in that same authentication. This enables full device management functionality with root access.

NOTE:

CredentialsforTelnet/SSHshouldhaveaprivilegelevelsufficienttostopservicesandtorestartthe computer system.

Firewall—Some firewallsinstalledon thecomputer mayblock Web-Based Enterprise Management requests. Allow those you want to manage.

License—Make sure you have the correct WBEM driver license installed. Licenses come in the following types:

- Major Vendor by Name - Such as Dell, Compaq, HP, Gateway.
- Server/Desktop individual license support.
- Generic computers - non-major vendors.
- ALL - this gives the driver all capabilities for any computer system.

Aruba Devices

By default, only SSH interactions work on these devices. If you want to use telnet you have to configure the device through the console or through an SSH session and turn iton.

Cut through or direct access sessions are only supported for SSHv2. You must create an SSHv2 management interface for the device and use it when attempting direct access. If you use SSHv1 the session does not connect (the ArubaOS does not support SSHv1), and if you select telnet, the driver cannot log into the device automatically, and must login manually.

SNMP v2c only supports read operations, not write. SNMP v3 supports both read and write, but not SNMP v3 informs. To manage Aruba devices you must use SNMP v2 or v3. Up to Aruba OS3.1.1.2, SNMPv2c (read-only) and v3 (read-write) are recommended. SNMP v1 does not work correctly.

NOTE:

Although SNMP handles the bulk of the communication with this device, and you must supply the correct SNMP authentication information, some information comes through telnet interaction, so you must supply telnet/SSH authentication too for all device interactions to work correctly.

Backup and Restore

You can backup or restore text files that reflect the *Startup Config and Running Config* as well as backup/restore of a binary *Flash* memory for the selected device. Youcan compare, store, view and version the text files, and can restore either text or flash memory to the device if you have the File Management option installed.

Because the Aruba mobility controller's flash memory backup is a compressed, binary `.tar.gz` file, the displayed Current Config is not always textual. The flash file is binary, so you cannot view it as text. Nevertheless, you can restore it as long as the backup file has the extension of `.tar.gz` Using this application to back up the flash automatically creates the file with this extension.

Avaya Device Prerequisites

Youmust do the following for the device driver to function correctly with Avaya devices. Whether you use RTCP or not, do the bullet point setup steps outlined in Setting Up RTCP on page 686, below.

Cut-Through

Avaya Communication Manager uses a special port for telnet sessions. To use the Telnet Cut-thru feature in this software, you must modify the port. Youcan do this during the discovery process by creating a Telnet authentication object and entering port 5023. Alternatively, once you install Communication Manager, open it from the Resources screen, select the Authentication tab and create a new Telnet management interface with port 5023. When you initiate a Telnet cut-thru session, Customer Manager asks for an emulation type. Select type 4410.

Setting Up RTCP

The following describes setting up RTCP with Avaya devices. Do this on the Avaya Media Server with the Media Server's Management Web Interface using a browser (for example, Internet Explorer) to access the IP address of the media server. For an S8300, S8400, S8500 use the server's IP address. For the S87xx, use the active server IP address—not server A or server B but the active server address. When accessing the web manager, after logging in, navigate to the Maintenance Web Page to see the following menu choice:

NOTE:

When changing SNMP settings Avaya recommends stopping and restarting the agent.

- Enable SNMP v2 on the device (under *Alarms* > *SNMP Traps)*
- Ensure the SNMP community strings match on theauthentications you haveconfigured and on the device (Alarms > *SNMP Agents*).

- Mediation Server (or Application Server, if enabled as a Mediation Server) must be allowed access on the device, either through its specific IP address(es) or by checking *Any IP address* under *Alarms* > *SNMP* Agents.
- *Security* > *Firewall* must have *SNMP, HTTP* and/or *HTTPS*, and *Telnet* and/or *SSH* and *RTCP* enabled both as *Input* and *Output* from the Server. If your system collects SNMP Traps and syslog messages, select those for Output from Serveronly.
- You must confirm, or start, the SNMP agent (Master Agent) on the device under *Alarms* > *Agent Status.*

After modifying the alarm and security areas of the web manager, you need to access the *Communication Manager* command line interface (CLI) using telnet, the *Native Configuration Manager* (from the main web management page) or use Avaya's *Site Administration* software. If telnetting to the CLI, choose the `w2ktt` terminal type upon login. You must have a Communication Manager login which may or may not be the same as the web manager login and password.

Execute and change the following:

1 Change **system-parameter IP-options** and specify under RTCP Monitor Server the:

- Default server IP address = the Application server IP address
- Default port = Must match what you put in your software
- Default RTCP Report Period = you can leave at default
- Enter/submit the changes

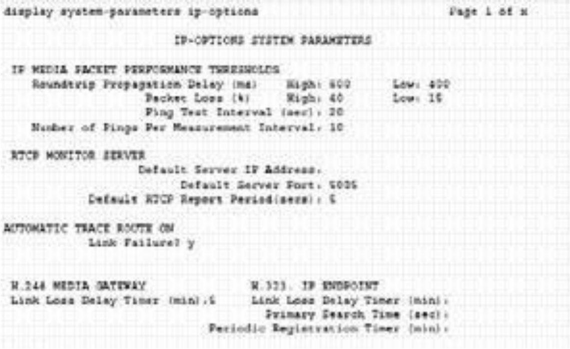

```
display system-parameters ip-options                          Page 1 of x

                          IP-OPTIONS SYSTEM PARAMETERS

 IP MEDIA PACKET PERFORMANCE THRESHOLDS
    Roundtrip Propagation Delay (ms)    High: 800       Low: 400
                       Packet Loss (%)    High: 40        Low: 15
                       Ping Test Interval (sec): 20
    Number of Pings Per Measurement Interval: 10

 RTCP MONITOR SERVER
                      Default Server IP Address:
                             Default Server Port: 5005
              Default RTCP Report Period(secs): 5

AUTOMATIC TRACE ROUTE ON
              Link Failure? y


 H.248 MEDIA GATEWAY                    H.323  IP ENDPOINT
 Link Loss Delay Timer (min):5          Link Loss Delay Timer (min):
                                           Primary Search Time (sec):
                                   Periodic Registration Timer (min):
```

2 Change **IP-network-region x** (for every network region that pulls RTCP information)

- On page one, make sure that RTCP Reporting is enabled.
- Under RTCP monitor server Parameters,

  –Use default server parameters (this uses the parameters you set up on the `system-parameter ip-options` form)

  –If you want to specify a separate server IP address, specify it under server IP address and server port

–Enter/submit the changes.

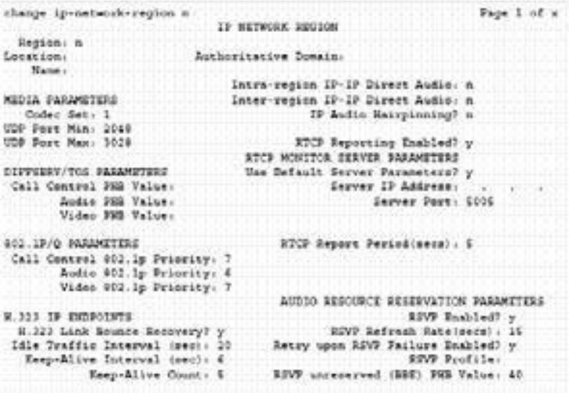

```
change ip-network-region n                                             Page 1 of x
                                  IP NETWORK REGION
  Region: n
Location:                       Authoritative Domain:
    Name:
                                     Intra-region IP-IP Direct Audio: n
MEDIA PARAMETERS                     Inter-region IP-IP Direct Audio: n
   Codec Set: 1                                  IP Audio Hairpinning? n
UDP Port Min: 2048
UDP Port Max: 3028                             RTCP Reporting Enabled? y
                                       RTCP MONITOR SERVER PARAMETERS
DIFFSERV/TOS PARAMETERS                Use Default Server Parameters? y
 Call Control PHB Value:                             Server IP Address:   .   .   .
        Audio PHB Value:                                  Server Port: 5005
        Video PHB Value:

802.1P/Q PARAMETERS                         RTCP Report Period(secs): 5
 Call Control 802.1p Priority: 7
        Audio 802.1p Priority: 4
        Video 802.1p Priority: 7
                                            AUDIO RESOURCE RESERVATION PARAMETERS
H.323 IP ENDPOINTS                                           RSVP Enabled? y
  H.323 Link Bounce Recovery? y                  RSVP Refresh Rate(secs): 15
 Idle Traffic Interval (sec): 20        Retry upon RSVP Failure Enabled? y
   Keep-Alive Interval (sec): 6                              RSVP Profile:
            Keep-Alive Count: 5         RSVP unreserved (BBE) PHB Value: 40
```

Brocade Devices

This software will not telnet connect to some devices if they use the factory default password. You must set the password to something other than that default. This software does not recognize the additional prompt asking that default password be changed each time login occurs.

Follow these steps and update firmware on the non-RX devices:

1. Download the firmware update (.zip file) from www.brocade.com
2. Extract that zip file to the download directory of the External FTP file server your system uses.
3. Create an empty file named `release.plist` and load into the OS Image manager portlet.
4. Deploy that image, selecting the device and `release.plist` file loaded in your system. To deploy the image, select the device(s) and select the `release.plist` file in OS Manager.
5. Remove any remaining files before attempting the next Brocade firmware update.

NOTE:

You cannot deploy these updates using this software’s internal FTP server.

For RX devices, download and deploy firmware updates as you ordinarily would, registering the OS image in the OS Images manager (see the *OS Images* section of this guide), and deploying it to either a device or group with the *action* menu.

BIG-IP F5

Cruz supports the BIG-IP F5 load-balancing appliance and software. It supports the following capabilities, and requires the listed device configurations:

- SNMP— This requires adding the IP address from which you manage the F5 to its Client Allow List under System -> SNMP -> Agent -> Configuration). Supports SNMP-based default device/resync using the ifTable - creates interface sub-components.
- Event Management—SNMP Trap Definitions (Tip: search for event definitions beginning with “big”).
- Reports—Run any default Cruz reports against F5 inventory.

Features not supported

- Any CLI-based functionality (NetRe store, CLI Cut-Thru, and so on)

- Link Discovery
- Port creation

Cisco Devices

The followingsections discuss prerequisites and limitations of Cisco device management capabilities.

Setup Prerequisites

You must include the *Enable* user ID. Omitted enable user IDs may interfere with correct device management.

You must create Authentication objects for Cisco Switches with an Enable user and password. Otherwise authentication for Interface level DC login sessions fail. The Cisco Router authentication rules automatically send the Enable user/password at the interface level.

For IOS devices, access to various system comma nd modes on the device may be defined by specifying an access privilege level for a user accoun t. Access privilege level 15 is required to access Enable (privileged) mode. This software requires the user account for authenticating with and managing a device has a privilege level of 15.

Ensure that user accounts associated with CLI Authentication objects in your system are configured on the device with privilege level 15. Consult your device's manuals for additional information about configuring this privilege level.

△ CAUTION:

If you upgrade IOS when you have not copied running-config to Startup-config, provisioning services may fail because of the unexpected device errormessage below. Thismessage causes "writememory" command to fail. In the application the user may see a message like "failure to terminate telnet session" Device message: *Warning: Attempting tooverwrite an NVRAM configuration previously written by a different version of the system image.* Also: Currently unsupported: Chassis View, Discrete configuration.

Saving Running-Config to Startup Config

Tosave the running-config to startup-config wh enever the router's configuration is updated uncomment the following in cisco.prop erties to enable this feature.

```
#cisco.ios.save.running.config=true
```

This copying behavior should not be fatal to th e configuration that was updated but a job message displays the failure or success.

NOTE:

Nexus configuration restore to start-up is not supported. When you attempt this, an error including `This command is deprecated`appears on 3000 series devices. On 5000 series, the error is `sysmgr_copy_nvram_dest_action: src uri_type = 1 is not supported yet.` (25242)

Copying to the Device Rather than Your System

Youcan have backup copy running configurations to the device's disk rather than the system's database. This requires no intervention of FTP since everything occurs on the device itself. The option appears when you change the following properties in the cisco.properties file in `owareapps/cisco/lib:`

```
#flags that enable direct copy run start
cisco.rmc.save.config=false
cisco.backup.save.config=false
```

When these are *true* you can back up running-config to startup-config for Cisco devices. You can select this option in addition to standard backup in the backup configuration screen (choose *running-config* from the File System pick list, and *startup-config* from the File Server Protocol pick list.) You can also trigger it from the *System* > *General* screen's *Save Config* button.

Service/Policy TroubleshootingFAQs

The following sections answer frequently as ked questions about Services and Policies.

Service Integrity Check

- When you right-click in the Service or Policies portlet, Maintenance > Integrity Check checks the service targets to see if they are still in the database. Are these targets just physical resources like ports or subinterfaces?

  Whatever the target is for the service, most typically this is the router itself.

- Will Cruz do this check when you invoke it? For example, if an equipment discovery finds a port is gone, will that flag the service as compromised or is it only flagged when Cruz runs the integrity check on the service?

  Youmust run the service integrity check; Cruz does not connect general device resync with any consequence to services.

About Services Discovery

- Are altered services flagged ascompromised if service rediscoverydetects services in the network have been changed manually is some way (for example, a different VLAN assignment), or does the Cruz service just get aligned to what is in the network? How would a user know it was different?

  Services in the database get updated to what is found on network. Here are the potential outcomes:

  –If Cruz finds a “provisioned” service in the database, but not on network, the service is flagged as “Not Found.”

  –If data found on network does not pass validation, the service is flagged as “Error”

  –If the data on network is different than what appears in Cruz’s database, Cruz saves the new data and versions the previous configuration, updating the Last Modified field. Users can compare/diff versions and revert to previous version through the application. Service Discovery Job/Audit trails log the creation of a new version and that the service was modified.

  Depending on what changes, Cruz may create a new service. If a key field changes, then what is discovered is a different service. Cruz cannot distinguish that VRF1 was modified to VRF2. These are therefore two different services. In such a case, Cruz marks VRF1 “Not Found” and creates a new service,VRF2.

- Howdo Idiscover pooled resourcesinto Cruz pools? Withservice discovery? Withequipment discovery? With something else?

  Servicediscovery, throughthecreation/updateofservicesmaintainsthepoolallocationsused/required by those services. Cruz requires no separate pooldiscovery.

Adaptive CLI FAQs

- Why share an existing schema from another Adaptive CLI (ACLI) versus creating a new one with each ACLI?

  One reason to use the same schema is to accommodate a complementary ACLIs. For example, one ACLI creates an entity and you want a script to remove the same entity. For such examples, the valid values, labels, and so on, for the attributes are always going to be the same when you create and delete ACLIs. Therefore, it is safest to use the same single referenced version of the Schema. You can share the same schema, and your delete script can mark the attributes it does not use as *Not applicable*.

- What is the best practice for exporting ACLIs to import later into another system?

  If you have ACLIs that you need to export so that you can import them into a production system, then the recommended practice is to create a separate file for each ACLI and export them one at a time.

  You can group select multiple ACLIs in the Adaptive CLI Manager and export them to a single file, but this can be difficult to maintain if changes are being made frequently to the ACLIs. Best practice is that only ACLIs directly sharing a common Schema (example a Create ACLI and its complimentary Delete ACLI) be exported to the same file. Keep in mind how to maintain/version/update the ACLIs and associated shared schemas when plotting how to map your export files, and frequently back up your export files to external devices/machines. You can use source control systems version/maintain ACLI export files, since they are in XML format.

- If I change an ACLI’s schema shared by other ACLIs, do I need to do anything to the other ACLIs?

  If you have multiple ACLIs sharing the same Schema, you should be in the habit of retesting the other ACLIs using that schema for to ensure no unintended side effects occur.

NOTE:

Regularly export all ACLIs with the same schema before modifying the schema by editing any of the ACLIs that use it. Also: Test your ACLI scripts in a telnet or direct access session with the target device(s).

When previewing ACLI and preview form is empty, most likely Perl is not installed. Best practice is to use Perl version 5.10 or later (however not 5.16).

Server Information

You can see mediation and application server in formation in JMX Console. The URLs for this console:

- **Mediation Server JMX**: http://[mediation server IP address]:8089/jmx-console/ (for stand-alone mediation servers), or port 8489 for HTTPS.
- **Application Server JMX**: http://[application server IP ad dress]:8089/jmx-console/, or port 8489 for HTTPS.

Some information visible in these consoles:

Is a mediation server active or standby? —Open the JMX console for the mediation server, then click PollingEngine and view the *Active* attribute. If *true* the mediation server is primary, if *false* it is standby.

To which application server is mediation server posting data? —In the mediation server's console, click ClusterPrimaryDesignator and then view the AppServerPartitionName attribute.

List active subscriptions and targets—Click PollingEngine, then invoke the getSubscriptionAndTargetInfo operation.

Is mediation server writing poll ing results to the spool file? — Click MonitorPollingHandlerMBean and then view the DataBeingWrittenToSpoolFile attribute. While there you can also see the most recent time the mediation server posted data to the application server (item 7) by viewing the MostRecentPostTime attribute.

When does a server skip execution and what is the total number of skips? —Click PollingEngine and then viewing those attributes. While there, you can also see the last time the server rejected execution and the last time that happened

NOTE:

The jmx-console is a Development tool used for troubleshooting and not accessible to the application user. Please request assistance through support channels to investigate any potential application issues.

Environment/Operating SystemIssues

The following are items that have historically ca used some problems. They may not apply to your environment.

CRON Events

CRON events can update Linux Releases, check for linkage errors and run through other tasks. This should occur when server traffic is generally higher. If necessary, change it to run late during off peak hours.

Potential Problem Processes

auditd—This processlogserrors periodically, and can send RESTART or KILLsignals to processes outside of its policies. This could be dangerous if configured wrong.

cpuspeed—Throttles down CPU speeds within the Kernel. Since Cruz's Java runs in a VM it is unaware of sudden increase/decrease in CPU speeds, and this could pose issues in threading and calculating thread counts.

SELINUX

This should be disabled. To check, run the following:

```
[root@AppRedcell01 bin]# selinuxenabled && echo enabled || echo disabled
```

To fix this, if it indicates it is enabled, modify the `/etc/selinux/config` and change `targeted` to `none` so this is preserved on reboots.

Hardware Errors

A fragment found in `dmesg`, triggered further investigation. This detected memory errors.

```
[Hardware Error]: Machine check events logged
```

```
Errors found in /var/log/mcelog:
Hardware event. This is not a software error.
MCE 0
CPU 30 BANK 9
TIME 1370474184 Wed Jun 5 17:16:24 2013
MCA: MEMORY CONTROLLER GEN_CHANNELunspecified_ERR
Transaction: Generic undefined request
STATUS 900000400009008f MCGSTATUS 0
MCGCAP 1000c18 APICID c0 SOCKETID 3
CPUID Vendor Intel Family 6 Model 47
```

DNS Does Not Resolve Public Addresses

DNS must permit application servers to resolve public DNS names like google.com. Web server needs this to determine its public facing interface by determining which route the packet went out on quick test.

Raise User Limits

If User Limits are lowon the Application Servers and possibly Mediation Servers, these can impact threading and normal server behaviors.

Web Server

The memory configuration (heap min/max) should also be the same for all server environments.

Portal Memory Settings

To manually change the web portal heap settings, change the `setenv.sh` (Linux) or `setenv.bat` (Windows) file:

```
set "PORTAL_PERMGEN=512m"
set "PORTAL_MAX_MEM=3072m"
set "PORTAL_INIT_MEM=768m"
set "PORTAL_32BIT_MAX_MEM=768m"
```

These files are in the `Tomcat***/bin` directory. For Linux, restart the portal service to apply new memory settings. In Windows, besides updating `setenv.bat` you must run `service.bat update` in that same directory.

You can increase these to even higher figures if your system has the memory available.

NOTE:

Make sure only one Tomcat process is running, otherwise your web server may exhibit poor performance.

Post Upgrade Web Portal Problems

Afterapplyingpatches, restarting processesorother recentactivity, web portal exhibitsundesired behavior like the following:

- Displaying pink barinapplication portlets indicatingapplications are `temporarily unavailable`.
- Message in application portlets indicating `resource unavailable`.

- Actions screen does not appear after right-click device and choosing *Actions*.
- Inability to expand hierarchiesorapreviouslyselected hierarchyexpands whenclicking another hierarchy.

Solution: These, and other symptoms, can stem from browser caching or attempting to login too soon after starting the application. To resolve, try waiting and/or clearing the browser cache.

Clustering

Youmustenable Clusteringon multipleweb servers, otherwiseindexand usersbecome outofsync. To solve this: enable clustering on the web servers by turning on cluster properties found within `synergy/conf/server-overrides.properties`

Upgrade Installations

The following outlines tasks to execute when you are updating your drivers, extensions or license. Refer to Upgrade/Data Migration Fails on page 657, Post Upgrade Web Portal Problems on page 693 for instructions about how to prevent and/or handle upgrade problems that can occur.

Patch Installation

Updating Driver Patches

1. Shut down your system.
2. On the application server designated as `oware.config.server` in `[installation root]/owareapps\installprops\lib\installed.properties`, copy the update file with the `ocp` or `ddp` extension to the `owareapps` directory.
3. Open a shell or command prompt, and Source the oware environment (Windows: `oware`. Linux: `. /etc/.dsienv`), and execute the following command lines:
4. `cd $OWAREAPPS`
5. `ocpinstall -x <ocp/ddp filename>`
6. `ocpinstall -u <ocp/ddp filename>`
7. ocpinstall -s <ocp/ddp filename >

Repeat steps 1 - 5 on any secondary application or mediation servers.

Adding or Updating Extensions

1. Copy extensions to the extensions folder: `[Installation root]\oware\synergy\extensions`
2. Restart web portal (Synergy) process.

Synergy Portal Updates (netview.war)

△ CAUTION:
Do this when no users are on the system. Apply this to all web servers.

With the portal running and the tomcat catalina log being tailed:

1. Navigate to `[Installation root]/oware/synergy/tomcat-x.x.xx/webapps` and delete the `netview` directory. After a brief pause you should see it being undeployed in the log.

2 Drop the new `netview.war` into the directory `[Installation root]/oware/synergy/deploy`. Wait a few minutes and you should see it hot deploy the new WAR file and load registry items.

3 After this is deployed, shut down the web servers.

△ CAUTION:

Ensure no old copies of netview.war remain in the [Installation root]/oware/synergy/deploy folder. This software automatically deploys any files in this folder. This will cause a conflict.

Synergy Portal Updates (NetviewFactory.war)

1 Shutdown your system (both webserver/appserver processes).

2 Navigate to `[Installation root]/oware/jboss-x.x/server/oware/deploy`. Apply `NetViewFactory.war` by overwriting any existing version with the new one.

License Installation

1 Stop the application or mediation server.

2 Rename the old license file (`[Installation root]\license.xml`)

3 Copy the new license file to your installation's root.

4 Rename it to `license.xml` if it is named anything else.

5 Open a shell and cd to your installation root.

6 Source the application's environment (Windows: `oware`. Linux `. /etc/.dsienv`)

7 Type `licenseimporter license.xml`

SMTP Mail Sender

If you require a sender/reply to e-mail address on mail sent, you can configure that with the following property (as always, it's best to override in `owareapps/installprops/lib/installed.properties`)

```
redcell.smtp.returnaddress.name
```

SMTP Test Failed "Could not convert socket to TLS"

If you encountered "Could not convert socket to TLS" during email test, you may need to import mail server's certificate into Cruz.

To import mail server's certificate into Cruz

1 assume user placed certificate self sign.cer in your desktop, e.g. `C:\Users\qa\Desktop`

open command prompt

cd `C:\Users\qa\Desktop`

type > `oware`

```
keytool.exe -import -trustcacerts -keystore $JAVA_HOME/jre/lib/
security/cacerts -noprompt -alias MyRootCA -file selfsign.cer
```

2 To validate import status:

```
keytool.exe -list -alias MyRootCA -storepass changeit -keystore
$JAVA_HOME/jre/lib/security/cacerts
```

3 restart appserver/webserver services

15

Localization

Cruz text appears in these distinct entities: menus, forms and fields, and messages. The Cruz Internationalization feature lets you alter the text for each of these entities to the language appropriate for a locale. The following sections provide what you need to do this.

Localization Overview

The installation wizard detects the operating system’s default language and installs the software so its default language. If you want the C ruz software installed with English regardless of the installation platform’s default, remove the SynergyI8N.jar file from Synergy.zip file before you install.

NOTE:

This should have been done as part of your pre-installation tasks.

The properties settings that force a locale for messages are the easiest to change. Only English message files ship with the software. You can override them as needed.

Settingthelanguage propertytoanemptystringininstallprops/lib/installed.properties fileapplies the operating system’s defaults.

Override the following Locale properties in that file:

```
oware.resourcebundle.language=en
oware.resourcebundle.country=US
oware.resourcebundle.language.variant=
```

The **language value** must be avalid ISO Language Code. These codes are the lower-case, two-letter codes as defined by ISO-639. Use your preferred search engine to find a full list of these codes at a number of sites.

The **country value** must be a valid ISO Country Code. These codes are the upper-case, two-letter codes as defined by ISO-3166. Use your preferred search engine to find a full list of these codes at a number of sites.

The **variant value** is a vendor or browser-specific code. Forexample, use WIN for Windows, MAC for Macintosh, and POSIX for POSIX. Where two variants exist, separate them with an underscore, and put the most important one first. Forexample, a Traditional Spanish collation might construct a locale with parameters for language, country, and variant as: `es, ES, Traditional_WIN`.

CAUTION:

Multilingual support does not necessarily extend to all application add-ons.

Language and Dictionary Portlets

The Language portlet displays f lags indicating languages available for many non-Cruz portlets and menus. Click a flag to change me nus and titles to the related language.

The Dictionary portlet is a separate portlet that allows you to select which dictionary to use. Follow these similar steps to add the Dictionary portlet.

By default, neither portlet reside on n Cruz page. Both portlets are listed under Portal Applications
> Tools on the Applications List on page 132.

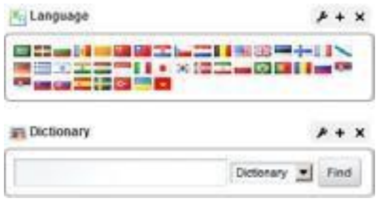

Adding the Language Portlet

The Language portletdisplaysflagsindicatinglanguagesavailablefor many non-Cruz portletsand menus. By default, this portlet does not reside on any Cruz pages.

Add the Language portlet as follows.

1 Select the page on which you want to add this portlet.

2 Select the Add > Applications menu option.

The application list is displayed.

3 Select Portal Applications > Tools > Language.

4 Click Add.

Note that the Dictionary is a separate portlet. Follow these similar steps to add the Dictionary portlet.

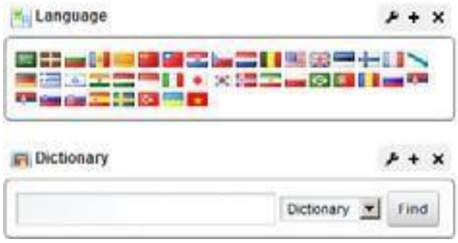

5 Click a flag to change menus and titles to the related language for non-Cruz portlets and menus.

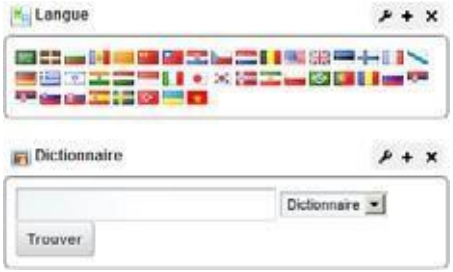

Localizing Events

If your events' MIB is in a language you want translated, the event MIB description is in English. Localize events from the Alarms or Event History portlets as follows.

1 Right-click the event and then select Details.

The Details portlet displays the MIB text.

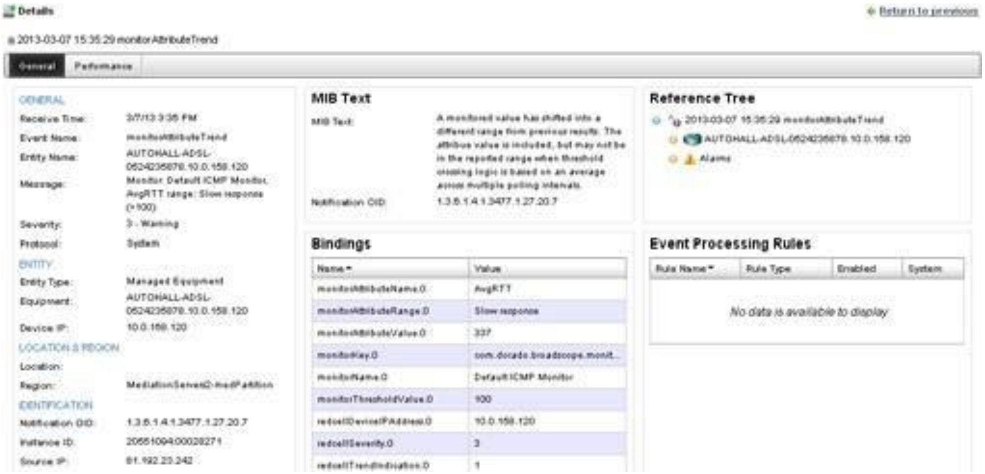

GENERAL

Receive Time:	3/7/13 3:35 PM
Event Name:	monitorAttributeTrend
Entity Name:	AUTOHALL-ADSL-0524236878.10.0.158.120
Message:	Monitor: Default ICMP Monitor, AvgRTT range: Slow response (>100)
Severity:	3 - Warning
Protocol:	System

ENTITY

Entity Type:	Managed Equipment
Equipment:	AUTOHALL-ADSL-0524236878.10.0.158.120
Device IP:	10.0.158.120

LOCATION & REGION

Location:	
Region:	MediationServer0-medPartition

IDENTIFICATION

Notification OID:	1.3.6.1.4.1.3477.1.27.20.7
Instance ID:	20651094.00028271
Source IP:	81.192.23.242

MIB Text

MIB Text:	A monitored value has shifted into a different range from previous results. The attribue value is included, but may not be in the reported range when threshold crossing logic is based on an average across multiple polling intervals.
Notification OID:	1.3.6.1.4.1.3477.1.27.20.7

Reference Tree

2013-03-07 15:35:29 monitorAttributeTrend
AUTOHALL-ADSL-0524236878.10.0.158.120
Alarms

Bindings

Name	Value
monitorAttributeName.0	AvgRTT
monitorAttributeRange.0	Slow response
monitorAttributeValue.0	337
monitorKey.0	com.dorado.broadscope.monit...
monitorName.0	Default ICMP Monitor
monitorThresholdValue.0	100
redcellDeviceIPAddress.0	10.0.158.120
redcellSeverity.0	3
redcellTrendIndication.0	1

Event Processing Rules

Rule Name	Rule Type	Enabled	System

No data is available to display

2 Select the text and copy it to the clipboard.

3 Paste it into your favorite translation site (such as translate.google.com .

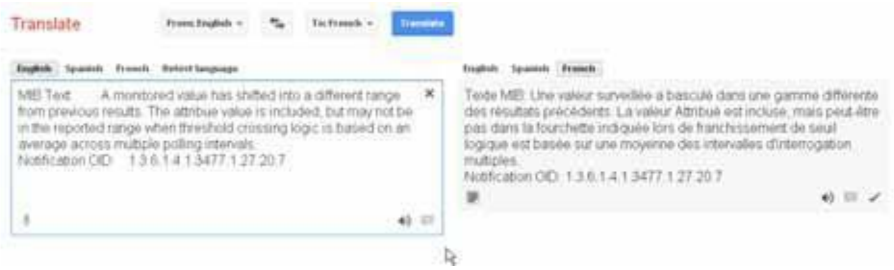

MIB Text A monitored value has shifted into a different range from previous results. The attribue value is included, but may not be in the reported range when threshold crossing logic is based on an average across multiple polling intervals.
Notification OID 1.3.6.1.4.1.3477.1.27.20.7

Texte MIB: Une valeur surveillée a basculé dans une gamme différente des résultats précédents. La valeur Attribué est incluse, mais peut-être pas dans la fourchette indiquée lors de franchissement de seuil logique est basée sur une moyenne des intervalles d'interrogation multiples.
Notification OID: 1.3.6.1.4.1.3477.1.27.20.7

4 Copy the translated text to the clipboard.

5 Return to the portlet used in step 1.

6 Right-click the event and then select Edit > Event Definition.

The Editing Event Definition window is displayed.

7 Paste the translated text into the Advisory Text field.

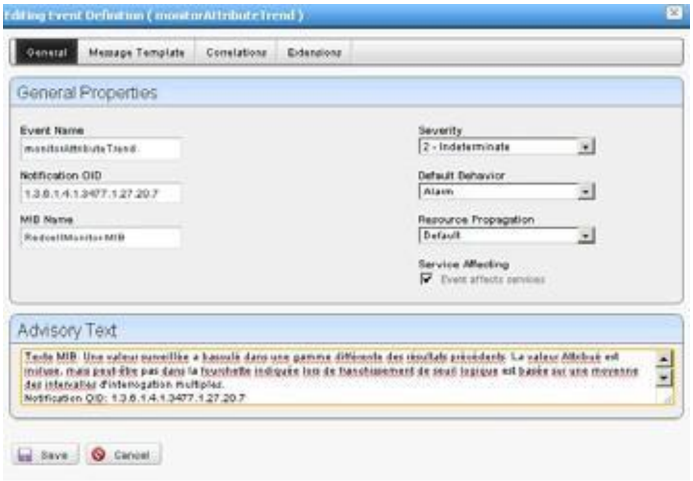

Event Name: monitorAttributeTrend

Notification OID: 1.3.6.1.4.1.3477.1.27.20.7

MIB Name: RedcellMonitor-MIB

Severity: 2 - Indeterminate

Default Behavior: Alarm

Resource Propagation: Default

Service Affecting: ☑ Event affects services

Texte MIB: Une valeur surveillée a basculé dans une gamme différente des résultats précédents. La valeur Attribué est incluse, mais peut-être pas dans la fourchette indiquée lors de franchissement de seuil logique est basée sur une moyenne des intervalles d'interrogation multiples.
Notification OID: 1.3.6.1.4.1.3477.1.27.20.7

8 Click Save.

9 Right-click the event and then select Details.

The Details portlet displays the advisory test.

NOTE:

Advisory Text only appears in the Details of Alarms, not Events.

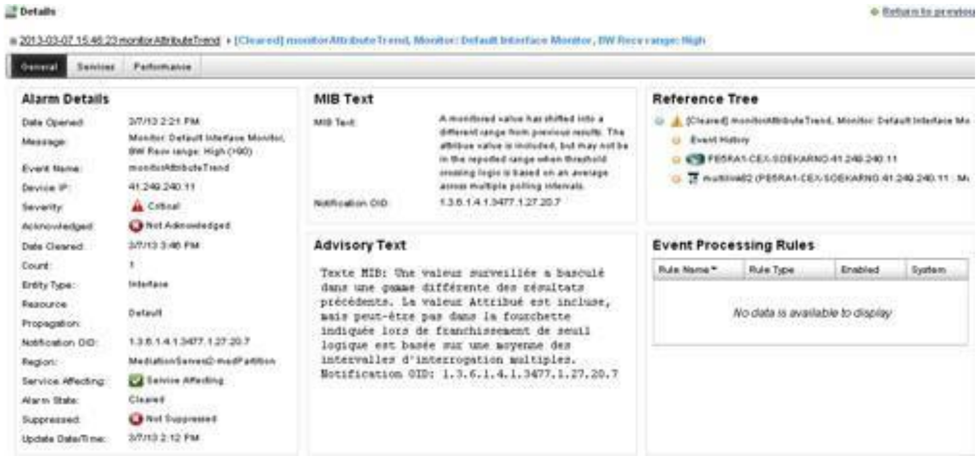

Date Opened:	3/7/13 2:21 PM
Message:	Monitor: Default Interface Monitor, BW Recv range: High (>90)
Event Name:	monitorAttributeTrend
Device IP:	41.249.240.11
Severity:	Critical
Acknowledged:	Not Acknowledged
Date Cleared:	3/7/13 3:46 PM
Count:	1
Entity Type:	Interface
Resource Propagation:	Default
Notification OID:	1.3.6.1.4.1.3477.1.27.20.7
Region:	MediationServer2-medPartition
Service Affecting:	Service Affecting
Alarm State:	Cleared
Suppressed:	Not Suppressed
Update Date/Time:	3/7/13 2:12 PM

MIB Text: A monitored value has shifted into a different range from previous results. The attribue value is included, but may not be in the reported range when threshold crossing logic is based on an average across multiple polling intervals.

Notification OID: 1.3.6.1.4.1.3477.1.27.20.7

[Cleared] monitorAttributeTrend, Monitor: Default Interface Mo

PE6RA1-CEX-SOEKARNO 41.249.240.11

multilink02 (PE6RA1-CEX-SOEKARNO 41.249.240.11 : Mu

Texte MIB: Une valeur surveillée a basculé dans une gamme différente des résultats précédents. La valeur Attribué est incluse, mais peut-être pas dans la fourchette indiquée lors de franchissement de seuil logique est basée sur une moyenne des intervalles d'interrogation multiples.
Notification OID: 1.3.6.1.4.1.3477.1.27.20.7

Rule Name	Rule Type	Enabled	System
No data is available to display			

Localizing Events not in Event History

Localize events that are not in Event History as follows.

1 Note the MIB name for the event you want to localize.

For example, the lldpRemTablesChange event is in LLDP-MIB.

2 Navigate to the Managed Resources portlet.

3 Right-click any resource and then select Direct Access > MIB Browser.

The MIB Browser window is displayed.

4 Locate the MIB note connected to the event.

For example, navigate to the lldpRemTablesChange event in LLDP-MIB by expanding:

```
LLDP-MID > lldpMIB > lldpNotifications > lldpNotificationPrefix
```

5 Click the MIB Information tab.

The description text is displayed.

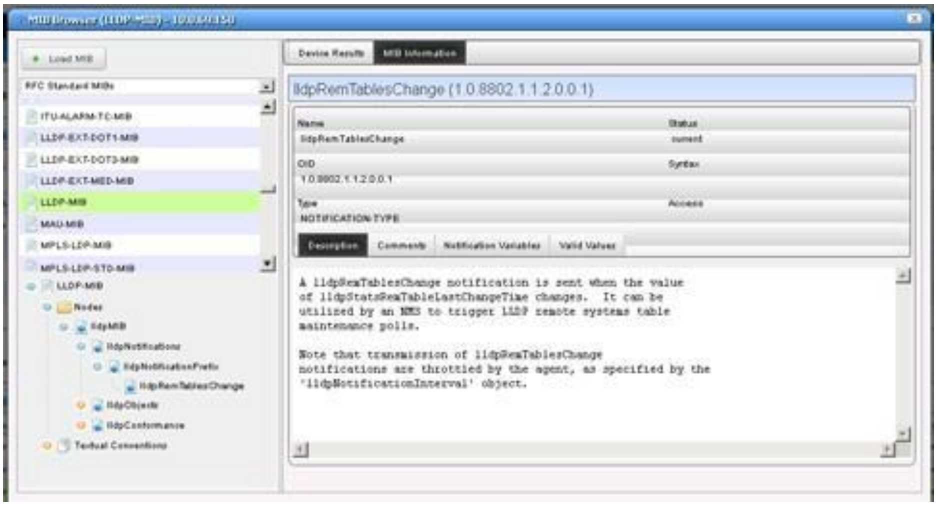

lldpRemTablesChange (1.0.8802.1.1.2.0.0.1)

Name: lldpRemTablesChange
Status: current
OID: 1.0.8802.1.1.2.0.0.1
Syntax:
Type: NOTIFICATION-TYPE
Access:

```
A lldpRemTablesChange notification is sent when the value
of lldpStatsRemTableLastChangeTime changes.  It can be
utilized by an NMS to trigger LLDP remote systems table
maintenance polls.

Note that transmission of lldpRemTablesChange
notifications are throttled by the agent, as specified by the
'lldpNotificationInterval' object.
```

6 Copy the text to the clipboard.

7 Right-click the event and then select Edit > Event Definition.

The Editing Event Definition window is displayed.

8 Paste the translated text into the Advisory Text field.

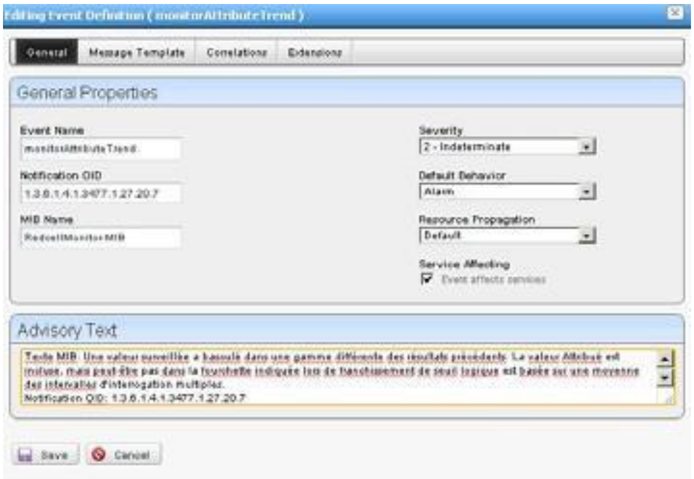

General Properties

Event Name: monitorAttributeTrend
Notification OID: 1.3.6.1.4.1.3477.1.27.20.7
MIB Name: RedcellMonitorMIB
Severity: 2 - Indeterminate
Default Behavior: Alarm
Resource Propagation: Default
Service Affecting: ☑ Event affects services

Advisory Text

Texte MIB: Une valeur surveillée a basculé dans une gamme différente des résultats précédents. La valeur Attribué est incluse, mais peut-être pas dans la fourchette indiquée lors de franchissement de seuil logique est basée sur une moyenne des intervalles d'interrogation multiples.
Notification OID: 1.3.6.1.4.1.3477.1.27.20.7

9 Click Save.

10 Right-click the event and then select Details.

The Details portlet displays the advisory test.

NOTE:

Advisory Text only appears in the Details of Alarms, not Events.

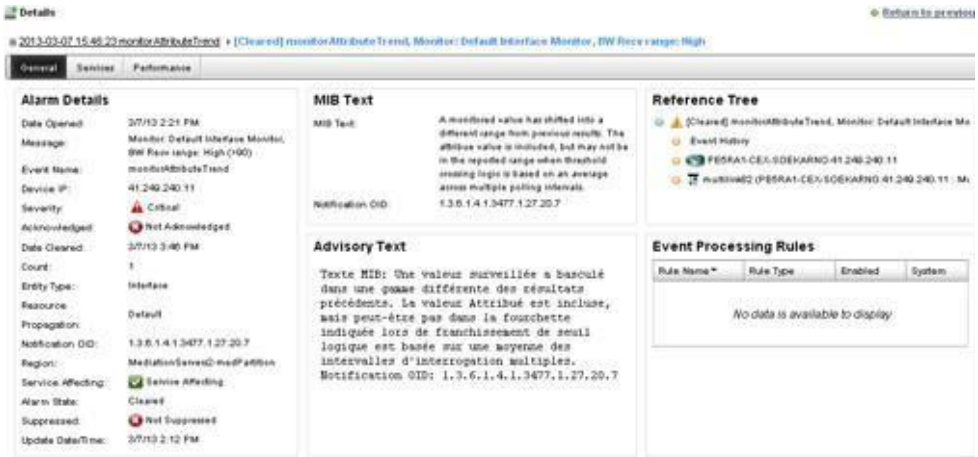

Details
Return to previou
2013-03-07 15:46:23 monitorAttributeTrend ▸ [Cleared] monitorAttributeTrend, Monitor: Default Interface Monitor, BW Recv range: High
General
Services
Performance

Alarm Details

Date Opened:	3/7/13 2:21 PM
Message:	Monitor: Default Interface Monitor, BW Recv range: High (>90)
Event Name:	monitorAttributeTrend
Device IP:	41.249.240.11
Severity:	Critical
Acknowledged:	Not Acknowledged
Date Cleared:	3/7/13 3:46 PM
Count:	1
Entity Type:	Interface
Resource:	Default
Propagation:	
Notification OID:	1.3.6.1.4.1.3477.1.27.20.7
Region:	MediationServer2-medPartition
Service Affecting:	✓ Service Affecting
Alarm State:	Cleared
Suppressed:	Not Suppressed
Update Date/Time:	3/7/13 2:12 PM

MIB Text

MIB Text:	A monitored value has shifted into a different range from previous results. The attribute value is included, but may not be in the reported range when threshold crossing logic is based on an average across multiple polling intervals.
Notification OID:	1.3.6.1.4.1.3477.1.27.20.7

Advisory Text

```
Texte MIB: Une valeur surveillée a basculé
dans une gamme différente des résultats
précédents. La valeur Attribué est incluse,
mais peut-être pas dans la fourchette
indiquée lors de franchissement de seuil
logique est basée sur une moyenne des
intervalles d'interrogation multiples.
Notification OID: 1.3.6.1.4.1.3477.1.27.20.7
```

Reference Tree

- [Cleared] monitorAttributeTrend, Monitor: Default Interface Mo
  - Event History
  - PE5RA1-CEX-SOEKARNO 41.249.240.11
  - multilink2 (PE5RA1-CEX-SOEKARNO 41.249.240.11 : Mu

Event Processing Rules

Rule Name	Rule Type	Enabled	System
No data is available to display			

Localizing Resource Bundles

You can localize some text that is not in Localizing Message Files (described below), but in resource bundles. Resource bundles let you localize static labels and button text on forms. These resource bundles are in .jar files, and load automatically based on the locale settings on the operating system where you installed the Cruz software. See Producing the Properties List on page 708 for more about file naming.

Resource bundles exist inmsgs.jarfileslikeowareapps\redcell\lib\rcform_msgs.ja r.The msgs.jar portion of this file name is common to all existing resource bundles.

The resource bundle for each form contains an ASCII properties file in the same directory as the form class. For example, “form_ A” in the “com.driver” package would be named:

```
com/driver/form_A.properties
```

This properties file contains text like:

```
#Resource bundle for Oware form com.driver.form_A
lblName.text=Name
lblName.tooltipt=Enter your name here
```

To localize this resource bundle into “Spanish/M exican,” the resource bundle for form_A in the com.driver package, is an ASCII properties file in the same directory as the form class, and is named:

```
com.driver.form_A_es_MX.properties
```

With text like:

```
#Resource bundle for Oware form com.driver.form_A
lblName.text=Nombre
lblName.tooltipt= Incorpore su nombre aquí
```

By default, a blank follows the equal sign, and the form displays the text specified when it was created. To override the default, add your text after the equal sign, and create your own .jar with that modified file, named as specified in Producing the Properties List on page 708. Note that not only the .properties files must follow this convention, you must also name the .jar file containing them to reflect the locale.

Localizing Message Files

A message file is a property file. The file name dictates which locales it applies to. The suggested naming convention is as follows:

`<prefix>msgs[_languageCode[_countryCode[_variantCode]]].properties`

Do not provide more precision than necessary. By default, all message files are named as one of the following:

`<prefix>msgs_en.properties (English)`

`<prefix>msgs.properties (language independent)`

No support exists for prepend or append operations.

All entries must follow this syntax: category.number = message text

To find all available message files, search the installation root and all its directories for `*msg*.properties` or `*.msgs` files

Finally, extract the synergy-i18n.jar file into the oware/synergy/extensions directory, edit the appropriate files in the localization subdirectory, then re-compress the .jar file.

△ CAUTION:

If you take the time to translate these files, make sure you keep a copy of any modified files because any upgrade may return them to their original state. You must manually copy the localized files to their original positions to see those translations after any update.

Thelocalization/language functionalitycomesfrom Liferay. Youmayfindadditionalinformation regarding localization on www.liferay.com/documentation.

Overriding MIB Text for Event Names

You can override default MIB text for event name s and descriptions using a text file with the .properties extension in the owareapps/[device driver name]/lib or owareapps/installprops/lib directory to include the override messages. For example:

```
# Below is an example overriding the default redcellDiscoveryJobBegin
  event Name and MIB Text with NEC CX specific information
#
# 1.3.6.1.4.1.3477.1.6.7.3 is the redcellDiscoveryJobBegin Notification
  OID
#
# 1.3.6.1.4.1.3477.1.6.7.3.1=cx2900NMDiscoveryJobBegin
# 1.3.6.1.4.1.3477.1.6.7.3.2=CX2900-NM discovery job begin notification
  indicates a discovery job has been executed
#
1.3.6.1.4.1.3477.1.6.7.3.1=cx2900NMDiscoveryJobBegin
1.3.6.1.4.1.3477.1.6.7.3.2=CX2900-NM discovery job begin notification
  indicates a discovery job has been executed
```

This overrides the redcellDis coveryJobBegin notification.

Withsuch overrides you can alter event names and notification descriptions. These event names and descriptions are visible in the Event Definitions manager.

To do this, create your own properties file in a te xt editor (for example, myfilemsgs.properties) and put it either in the driver's lib directory, or in owareapps/installprops/lib directory if you do not want this file overwritten by any application upgrades.

Caching Message/Property Files

The Cruz loads all message files when it loads properties. The application creates a temporary cache under the oware/temp/msgs directory for the client and the oware/temp/appserver/msgs directory for the server. Subsequent loads use this cache unless a property file or message file was modified (messages are in separate cache).

Creating Property Files for Double-Byte Characters

Properties files must be in escaped unicode format for the properties to appear in a double-byte character set. Create properties files for double-byte character format as follows.

1 Open an editor that is UTF-8 capable.

NOTE:

Notepad UTF-8 files do not work because notepad inserts a Byte Order Mark at the beginning of UTF8 streams.

2 Enter the properties in the text editor using any available double-byte character entry method.

3 Save the file as UTF8 - No Signature.

4 Convert this file to unicode-escaped ANSI form at by running the native2ascii Java utility.

You must specify the source and target file in the command. For example:

```
native2ascii -encoding UTF8 chinese-utf8.properties chinese.properties
```

These steps are just an example. Make sure the final file is named appropriately to match the Oware naming standard previously described. This file displays double-byte characters on Windows double-byte versions.

Understanding the Message Properties Files

Following the convention cited in Localizing Message Files on page 704, the name of a Spanish-language (Español) message properties file intended for use in Mexico would be:

```
rcmsgs_es_mx.properties
```

The entries in this file are grouped according to f unction. If you want to create a local language version of the Cruz messages, edit the message properties file and translate the message portion of each entry into the language appropriate for yoursite.

Edit the Message Properties file with any text editor that renders the characters you need. Double-byte charactersareallowed, but the Cruzapplication supports onlyleft-to-right textrendering. The following example shows an edited message properties file with Italian translations.

```
#
# rcmsgititaliano.properties
# Ciò è l'archvio italiano del messaggio per il nucleo 4.1 di Redcell
#

RC_GENERAL.1=Aggiunta Nuovo
RC_GENERAL.2=Modificando "
RC_GENERAL.3="
RC_GENERAL.4=Prevista Scoperta
RC_GENERAL.5=Modificando
RC_GENERAL.6=Programmazione"
RC_GENERAL.7=Programma
RC_GENERAL.8=Previsto Evento
```

Producing the Properties List

Product a properties list as follows.

1 Shut down the application server.

2 Add the following section to the Cruz log4j.xml file (such as \owareapps\redcell\server\conf\redcell-log4j.xml):

```
<category name="com.dorado.redcell.appframework.nav.RCNavI18NHelper">
  <priority value="TRACE"/>
</category>
```

3 Restart the application server.

4 Open the application server log file once the application is running. T

The server log is in the \oware\jboss-5.1\server\oware\log\server.log directory.

5 Search for the following properties information in the serverlog:

```
--- I8N: Redcell Navigation Tree --------------
You should find something that looks like this:
--- I8N: Redcell Navigation Tree --------------
- Place the following list in a message file and replace each value.
- If the property is already in place, the retrieved value is displayed.
RCMenuObject.I18N.menu_reportManager.1=Reports
RCMenuObject.I18N.leaf_templateMgr.1=Report Templates
RCMenuObject.I18N.leaf_reportsMgr.1=Reports
RCMenuObject.I18N.menu_inventory.1=Inventory
RCMenuObject.I18N.leaf_deviceDiscWiz.1=Resource Discovery
RCMenuObject.I18N.leaf_equipMgr.1=Resources
RCMenuObject.I18N.leaf_equipRoleMgr.1=Resource Roles
…
…
…
RCMenuObject.I18N.leaf_RetentionPolicy.1=Retention Policies
--- I8N: Redcell Navigation Tree --------------
```

NOTE:

The above is an example. Menu item order can differ from this example.

Altering Double-Byte Characters in Audit Trails

Tosee double-byte characters inaudit trail messages, you need to alter the database, asdescribed in the following sections.

- Setting Up MySQL
- Setting Up Oracle

Follow the steps outlined for your database type.

Setting Up MySQL

Setting up your MySQL database for double-byte characters in audit trail messages as follows:

1 Stop MySQL using one of the following commands:

- On Windows, execute this as an administrator;

```
net stop mysql
```

- On Linux, execute this as root:

```
/etc/init.d/owaredb stop
```

2 Open the MySQL configuratio n file in a text editor:

- Windows file is %SYSTEMROOT%/my.ini
- Linux file is /etc/my.cnf

3 Add the following lines to the [mysqld] section:

```
default-character-set=utf8
default-collation=utf8_general_ci
character-set-server=utf8
collation-server=utf8_general_ci
init-connect='SET NAMES utf8'
skip-character-set-client-handshake
```

4 Restart MySql.

- On Windows, execute this as an administrator:

```
net start mysql
```

- On Linux, execute this as root:

```
/etc/init.d/owaredb start
```

5 Run dbevolve.

- On windowb2s execute as the install user:

```
oware dbevolve -x
```

- On Linux,

```
. /etc/.dsienv
dbevolve -x
```

Setting Up Oracle

Set up the Oracle database for double-bytecharacters inaudit trail messages byselecting the UTF8 character set when creating an Oracle database.

```
CREATE DATABASE owbusdb CHARACTER SET UTF8;
```

Forexisting databases, Oracle hasextensive documentation regardingcharacter set selection and conversion. Refer to download.oracle.com/docs/cd/B10501_01/server.920/a96529/ch10.htm.

The simplest form of this is as follows:

ALTER DATABASE owbusdb CHARACTER SET UTF8;

16

Configuring Virtualization

As virtualized network functionality becomes central to a service provider's business model, key Network Functions Virtualization (NFV) capabilities need to be brought under management. With the Cruz NFV applications and the third-party OpenStack operating system, you can quickly bring virtual network functions (VNFs) under full resource management.

This section is intended for system administrators, network administrators, or anyone else using Cruz and OpenStack ® infrastructure to design, configure, and administer your environment. This document assumes that you are familiar with or have knowledge of:

- Linux distribution that supports OpenStack cloud, SQL databases, and virtualization
- Networking configurations
- Concepts, such as DHCP, Linux bridges, VLANs, iptables

Thissection covers the following topicsrelated to the Cruz NFV administration and configuration, such as defining your OpenStack environment, preparing vendor virtual network function (VNF) images for use, and configuring those images:

For a description of the network viritualization-specific portlets see Network Virtualization Portlets on page 866. Cruz NFV Permissions on page 993 and Notifications on page 1001 provide reference information that you might need during configuration.

Overview

Service providers need to accelerate the deployment of new network services to support their revenue and growth objectives. The Cruz Network Functions Virtualization (NFV) feature and its ability to utilize the OpenStack operating system (Figure 16-1) to bring a Virtualized Network Function (VNF) under management gives service providers what they need to:

- Reduce hardware and support costs
- Minimize installation costs
- Optimize scaling and usage
- Enable innovation
- Simplify network services rollout and management
- Reduce risks associated with rolling out new services
- Improve return on investment of new services
- Support NFV packaging of virtual functions and infrastructure resource management

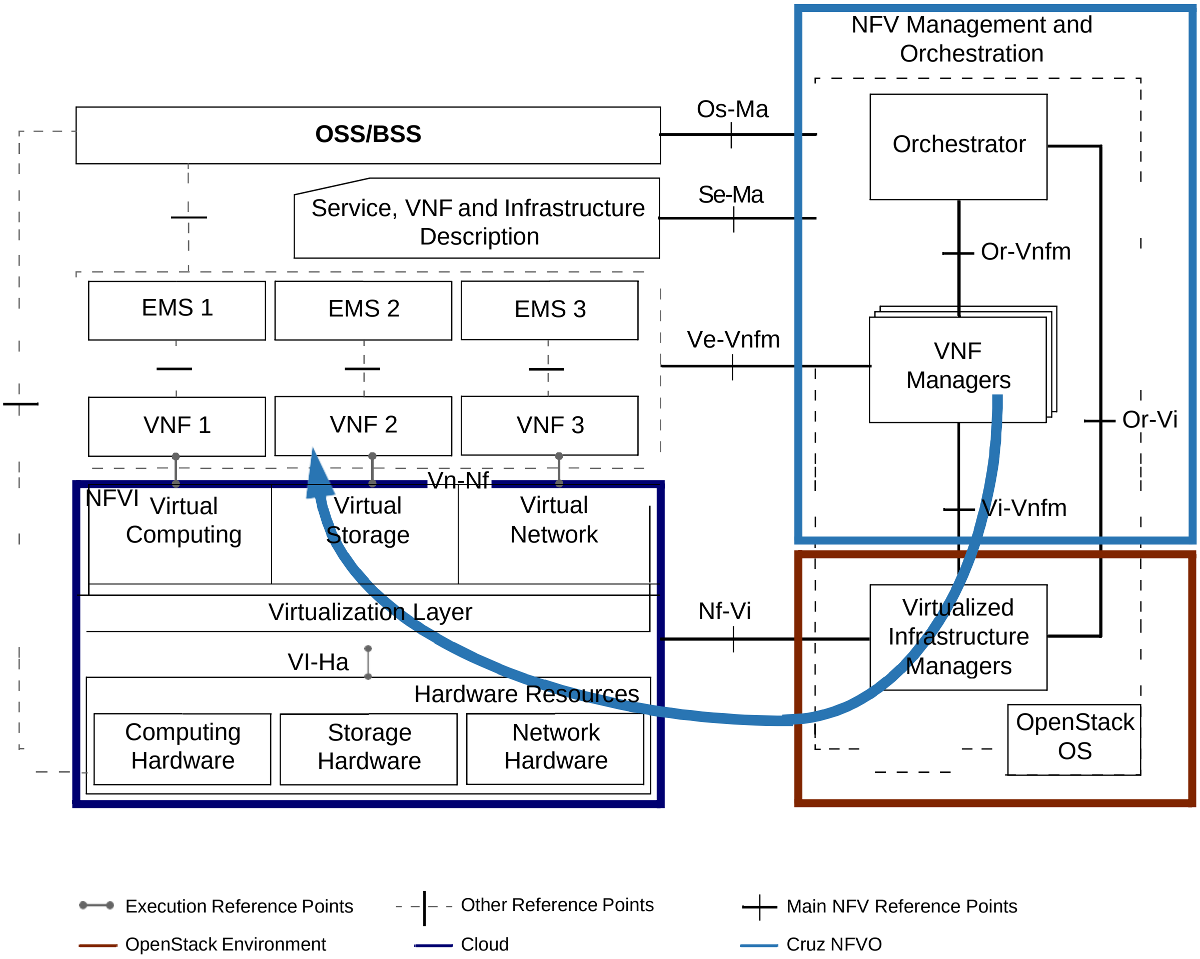

Figure 16-1. Cruz and OpenStack Interaction

The OpenStack operating system provides services used to manage (Figure 16-2):

- Instances' Lifecycle (run, reboot, suspend, resize and terminate instances)
- Compute Resources (CPU, memory, disk, and network interfaces)
- Local Area Networks (Flat, Flat DHCP, VLAN DHCP and IPv6) through programmatically allocated IPs and VLANs
- Who has access to compute resources and prevent users from impacting each other with excessive API utilization
- Virtual Machine (VM) Images (store, import, share, and queryimages)
- Floating IP Addresses (assign and re-assign IP addresses to VMs)
- Projects and Quotas (allocate, track and limit resource utilization)

The following example shows the services interaction in an OpenStack environment.

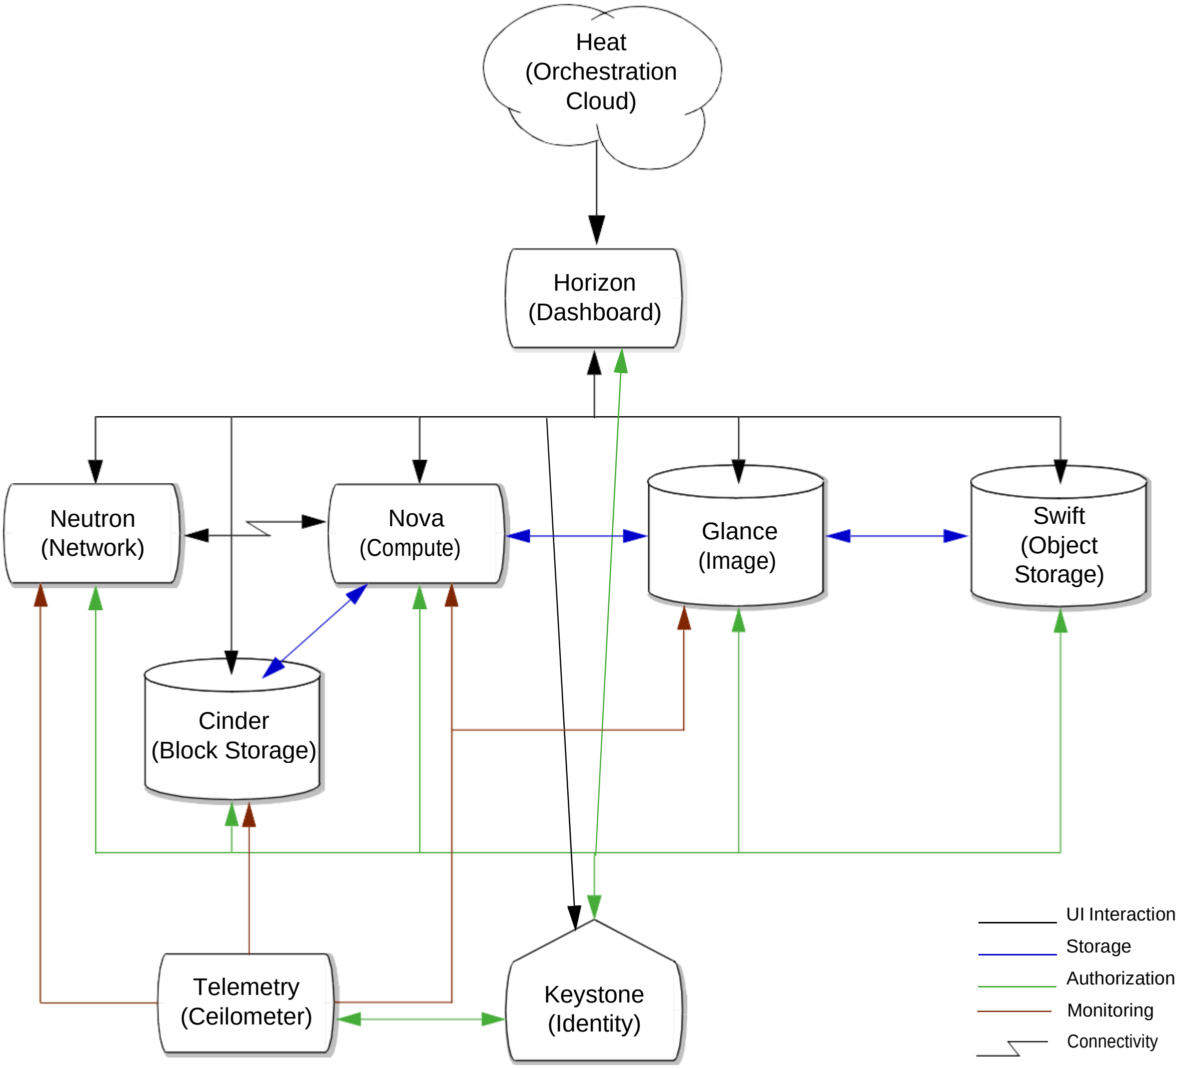

Figure 16-2. OpenStack Services (example)

The following list (Table 16-1) provides a brief description ofeach OpenStack serviceand the Heat Orchestration program provided in the previous example. The identity, compute, and image services are standard dashboard components. The other components areoptional.

Details of the OpenStack system is out of scope for this document. Refer to the OpenStack documentation for more details on these services and other OpenStack services, programs, and so on.

Table 16-1. OpenStack Services

Service Name	Service Type	Description
Heat	Orchestration	Implements an orchestration engine to launch multiple composite cloud applications based on templates in the form of text files that are treated like code.
Horizon	Dashboard	Provides the Web-based user interface used to manage OpenStack services (network, compute, image, block and object storage, and identity).
Neutron	Network	Provides network connectivity between interface devices managed by other OpenStack services, such as Compute.
Nova	Compute	Provides the ability to manage your computing resources, such as CPU, memory, disk, and network interfaces. This allows you to easily migrate workloads and run them at scale with predictability, consistent performance, control, and visibility.
Glance	Image	Stores images from the compute service.
Swift	Object Storage	Provides cloud-storage software so that you can store and retrieve data with an API (application programming interface). Ideal for storing unstructured data that grows without bounds, such as image files.
Cinder	Block Storage	Provides the ability to virtualize block storage devices management and provide a self-service API for end-users to request and consume resources without knowing where their storage is actually deployed or on what type of device. In our example, it stores volumes for the Compute service.
Telemetry	Ceilometer	Monitors, collects, normalizes, and transforms data produced by OpenStack services (network, compute, image, and block storage). Use this data to create different views to help solve telemetry issues.
Keystone	Identity	Provides authentication (authN) and high-level authorization (authZ) for the dashboard, network, compute, image, block storage, and object storage services.

Preparing to Configure NFV Infrastructure

This section provides what you need to prepare for Network Functions Virtualization (NFV) infrastructure configuration and how to start the applications used to do the preparation tasks.

- Verifying Prerequisites
- Signing In/Out Cruz Application
- Starting and Exiting OpenStack Dashboard
- Setting Up an OpenStack Project
- Bringing a VIM Under Management
- Setting Up Cruz Application
- Determining What to Do Next

Verifying Prerequisites

Before you can perform the virtualization, configuration tasks the Network Functions Virtualization (NFV) product, make sure that the following software is installed and operational:

- Cruz 7.4.x application with the NFV feature

  Start the Cruz application and the n select Manage > Show Version to verify your version. See Signing In/Out Cruz Application on page 716 if you need instructions.

  Select Add > Applications and then verify that Network Virtualization applications are listed.

  For a list of NFV applications, see Starting and Exiting OpenStack Dashboard on page 717.

  Refer to the *Cruz Installation Guide* for system requirements and the *Cruz Release Notes* for new features supported, adaptive CLIs, standard change management policies, known issues, and supported devices, switches, and equipment.

- OpenStack operating system (Liberty , Mirantis version 8)

  The OpenStack environment is set up and defined byyoursystem administrator. Obtain login information your system administrator. See Setting Up an OpenStack Project to set up interaction between the Cruz application and the OpenStack operatingsystem.

- One of these Web browsers at the listed version or above: Chrome V45, Firefox V40, Internet Explorer V11, or Safari V8

Make sure that you have access to Cruz and third-party documents that contain additional information related to the configuration tasks described in this guide.

- *Cruz Release Notes*
- *Cruz Installation Guide*
- OpenStack documentation
- Cisco® CSR 1000V-Seriesdatasheet
- F5® BIG-IP® Virtual Edition(VE) datasheet and OpenStack KVMenvironment requirements
- Juniper Networks ® Firefly SRX documentation and release notes
- Sonus TM Network Session Border Controller (SBC) Software Edition (SWe) documentation and release notes

Signing In/Out Cruz Application

To get started using the Cruz NFV features, you need to know how to sign into the Cruz application and sign out when you are finished with the application.

Sign into your Cruz NFV installation from your W eb browser as follows and then sign out when finished.

1 Enter the Cruz URL.

`http://`*`appServerHost`*`:`*`portNumber`*

Replace *appServerHost* with your Cruz NFV server host name or IP address. The default *portNumber* is 8080.

The sign in request is displayed.

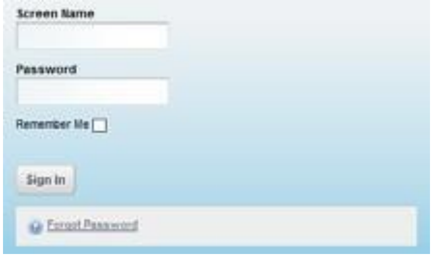

Screen Name

Password

Remember Me □

2 Enter your screen name (username) and password.

3 Click Sign In.

The Cruz NFV Home page is displayed.

If you entered an incorrect usernameor password, a message is displayed. Enter the information again.

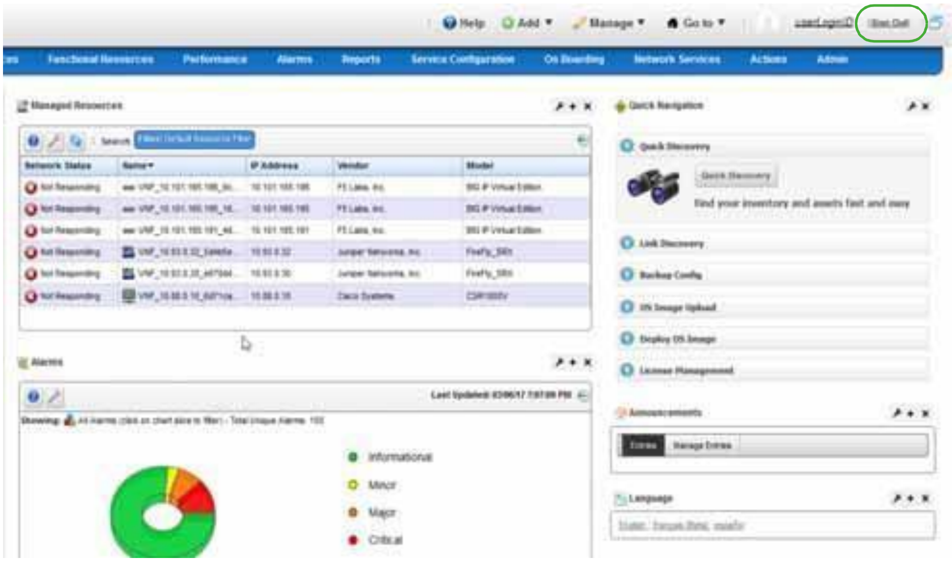

Network Status	Name	IP Address	Vendor	Model
Not Responding	VNF_10.101.185.186_…	10.101.185.186	F5 Labs, Inc.	BIG-IP Virtual Edition
Not Responding	VNF_10.101.185.185_…	10.101.185.185	F5 Labs, Inc.	BIG-IP Virtual Edition
Not Responding	VNF_10.101.185.181_…	10.101.185.181	F5 Labs, Inc.	BIG-IP Virtual Edition
Not Responding	VNF_10.93.0.32_…	10.93.0.32	Juniper Networks, Inc	Firefly_SRX
Not Responding	VNF_10.93.0.30_…	10.93.0.30	Juniper Networks, Inc	Firefly_SRX
Not Responding	VNF_10.88.0.16_…	10.88.0.16	Cisco Systems	CSR1000V

Find your inventory and assets fast and easy

Last Updated: 03/06/17 7:07:09 PM

Showing: All Alarms (click on chart slice to filter) - Total Unique Alarms: 100

4 Click Sign Out when you are finished with the application.

You are returned to the sign in request page.

Starting and Exiting OpenStack Dashboard

Connect to the OpenStack dashboard from your We b browser as follows and then sign out when you are finished.

1 Enter the OpenStack URL.

`http://`*`serverIP`*`/horizon/auth/login/`

Replace *serverIP* with your OpenStack server host name or IP address.

The OpenStack Dashboard sign in page is displayed.

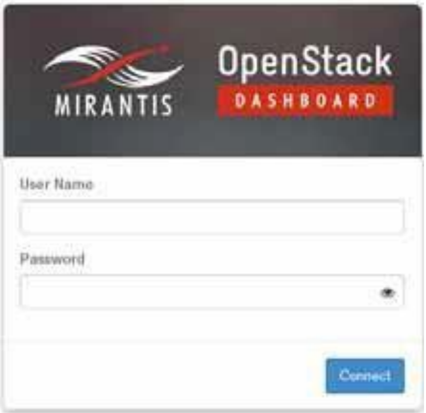

2 Enter your username and password.

3 Click Connect.

The OpenStack dashboard is displayed.

If you entered an incorrect usernameor password, a message is displayed. Enter the information again.

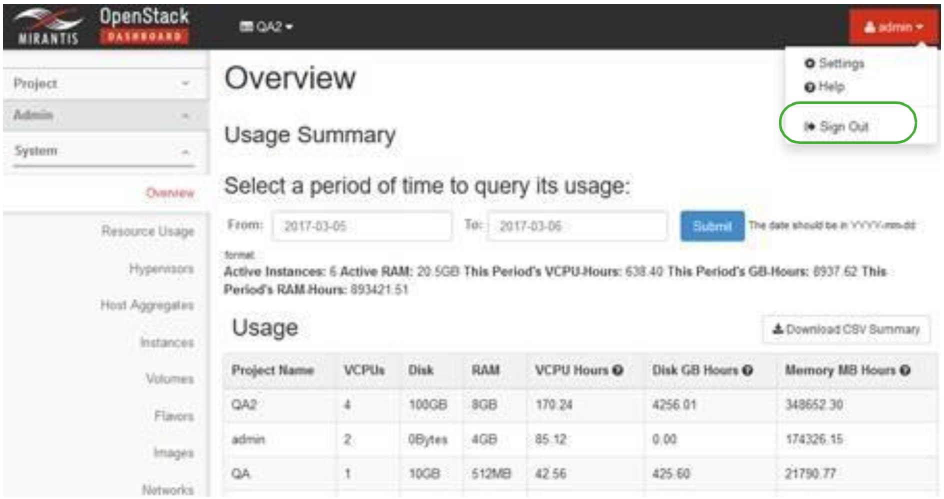

Overview

Usage Summary

Select a period of time to query its usage:

From: 2017-03-05 To: 2017-03-06 The date should be in YYYY-mm-dd format.

Active Instances: 6 **Active RAM:** 20.5GB **This Period's VCPU-Hours:** 638.40 **This Period's GB-Hours:** 8937.62 **This Period's RAM-Hours:** 893421.51

Usage

Project Name	VCPUs	Disk	RAM	VCPU Hours	Disk GB Hours	Memory MB Hours
QA2	4	100GB	8GB	170.24	4256.01	348652.30
admin	2	0Bytes	4GB	85.12	0.00	174326.15
QA	1	10GB	512MB	42.56	425.60	21790.77

4 Select *userName* > Sign Out when finished with the application.

You are returned to the OpenStack Dashboard sign in page.

Setting Up an OpenStack Project

OpenStack software is the third-party cloud computing software that the Cruz application interacts with to manage NFV packagingofvirtual functions and infrastructure resources. Before defining your Virtualized Infrastructure Managers (VIMs), you need to make sure that you have an OpenStack project in which the VIMs will reside and make sure that the project is set up to interact with the Cruz application.

Set up an OpenStack project u sing one of the following set of instructions:

- Creating a Project if the project does not exist.
- Verify Project Setup for Cruz if the project already exists.

NOTE:

Details of the OpenStack system is out of scope for this document. Refer to the OpenStack documentation if you need more OpenStack details than what is provided in this document.

Creating a Project

Create a project from the OpenStack dashboard as follows.

1 Select Identity > Project.

2 Click Create Project.

The Create Project window is displayed.

3 Enter a name and detailed description.

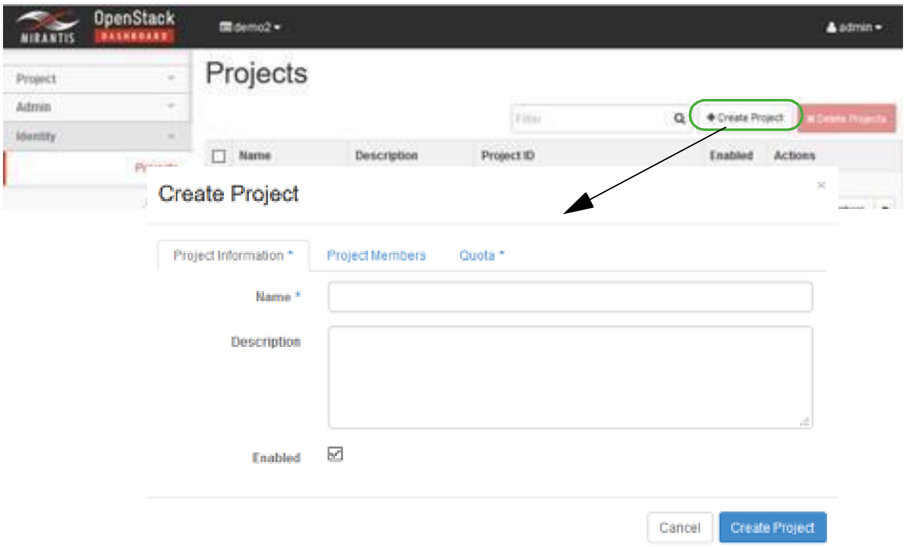

4 Assign users/permissions.

a. Click Project Members.

b. Add the following users to the Project Members list.

- heat-cfn
- heat
- a dmin

c. Set the member permissions to admin.

5 Click Create Project.

The new project is saved to the Projects list.

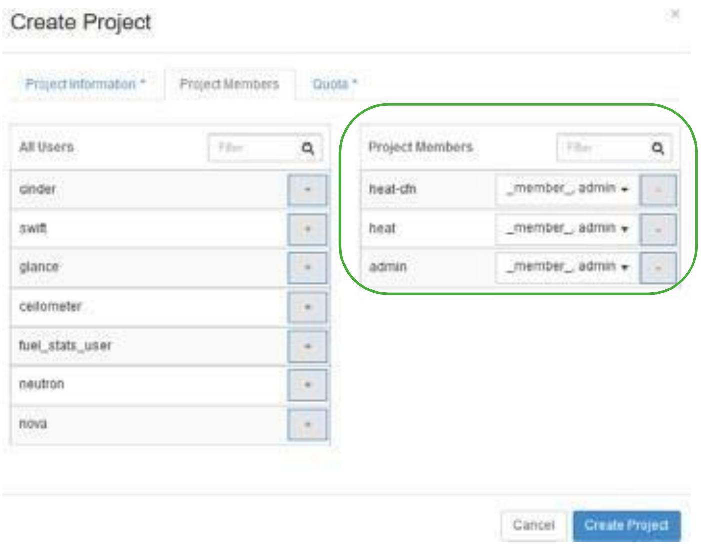

6 Create project networking in an Op enStack VIM (KILO or later release).

a. Select Project > Network > Routers.

b. Click Create Router.

The Create Router window is displayed, where you create a project router with a gateway interface.

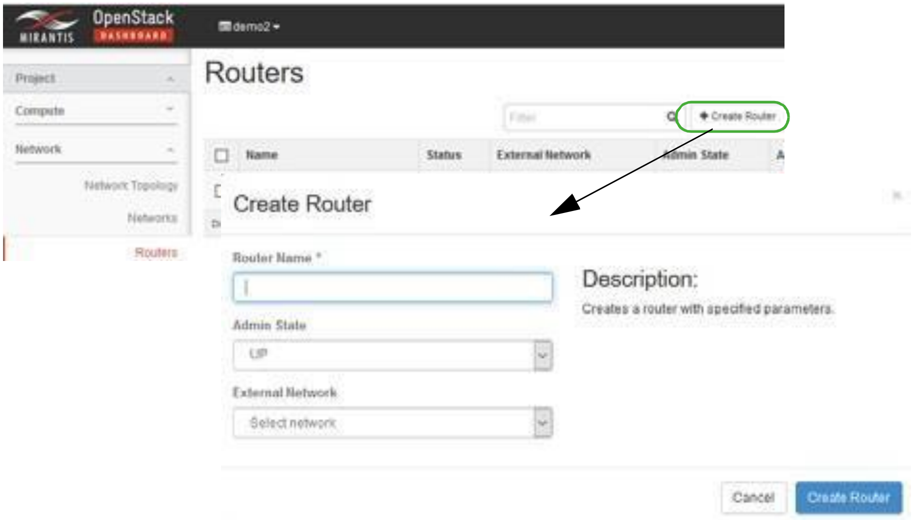

c. Make sure that the router’s gateway IP is reachable from the LAN (your Windows system) following project router creation.

d. Select Project > Network > Networks.

e. Click Create Network.

The Create Network window is displayed.

f. Create and attach the following networks to the project router:
   - net-mgmt
   - net-data1
   - net-data2

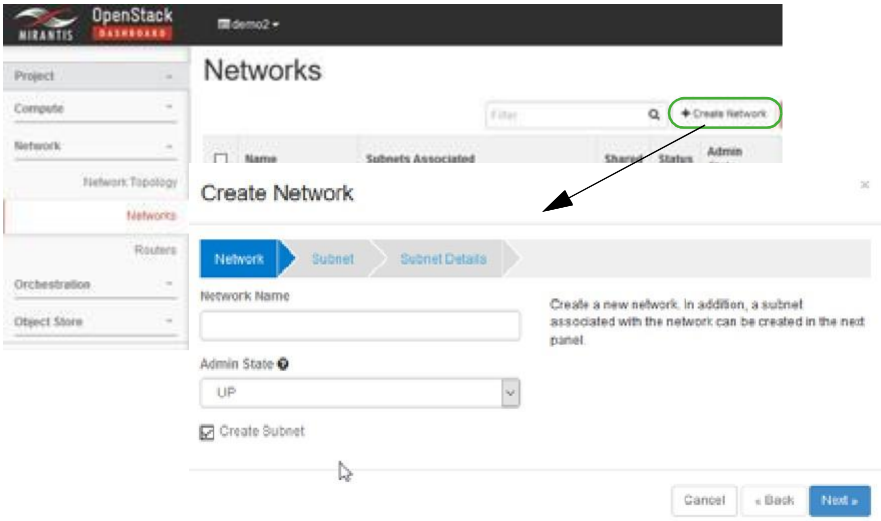

Network Name

Admin State: UP

☑ Create Subnet

Create a new network. In addition, a subnet associated with the network can be created in the next panel.

Now that your OpenStack project is ready, you can get started by creating a VIM from t he Cruz application.

Verify Project Setup for Cruz

Verify that an existing OpenStack project has the proper settings to communicate with the application and make the ap propriate changes as needed.

Verify a project’s setup from the OpenStack dashboard as follows.

1 Select Identity > Projects.

A list of projects is displayed.

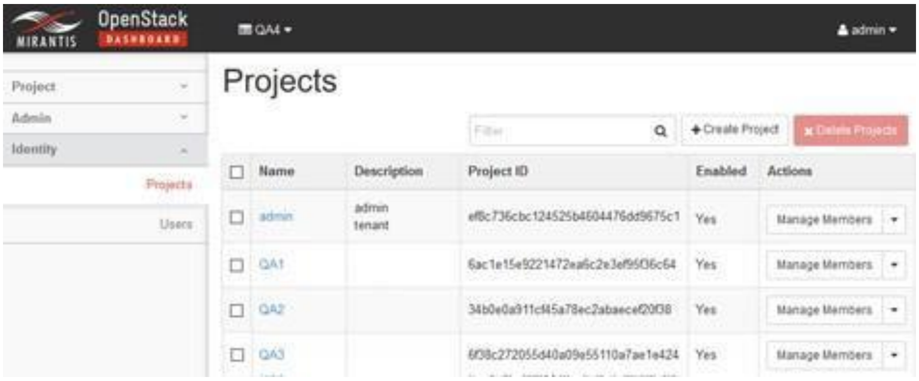

□	Name	Description	Project ID	Enabled	Actions
□	admin	admin tenant	ef8c736cbc124525b4604476dd9675c1	Yes	Manage Members ▾
□	QA1		6ac1e15e9221472ea6c2e3ef95f36c64	Yes	Manage Members ▾
□	QA2		34b0e0a911cf45a78ec2abaecef20f38	Yes	Manage Members ▾
□	QA3		6f38c272055d40a09e55110a7ae1e424	Yes	Manage Members ▾

2 Select the Edit Project action for the appropriate project.

The Edit Project window is displayed.

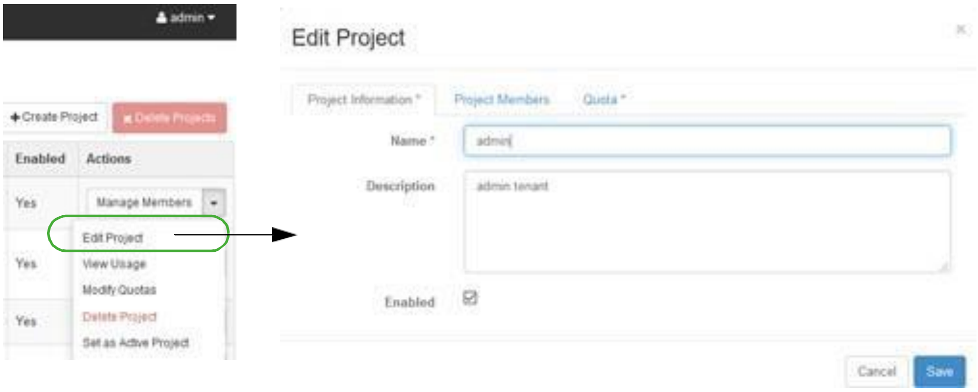

Name * admin

Description admin tenant

Enabled ☑

3 Click Project Members.

4 Add the following users to the Project Members list if they do not exist.

- heat-cfn
- heat
- admin

5 Set the member permissions to admin.

6 Save your changes.

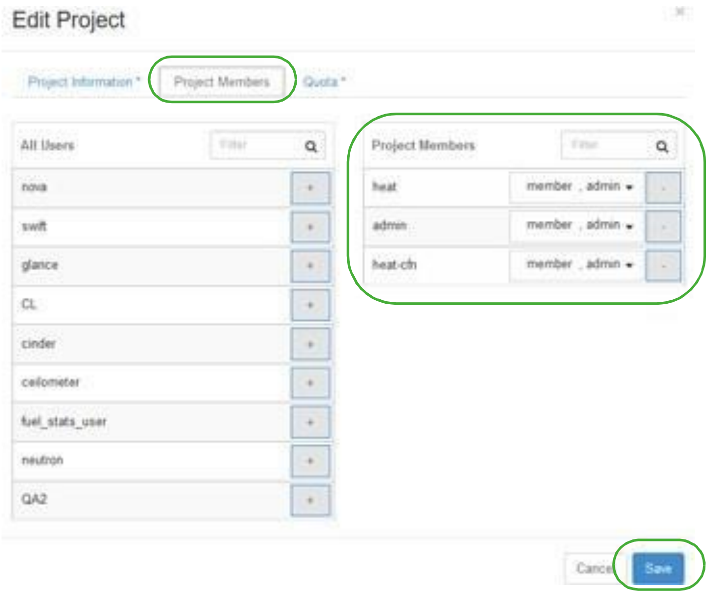

All Users
nova
swift
glance
CL
cinder
ceilometer
fuel_stats_user
neutron
QA2

Project Members	
heat	member , admin
admin	member , admin
heat-cfn	member , admin

7 Create project networking in an OpenStack VIM (KILO or later release) if it does not already exist.

a. Select Project > Network > Routers.

b. Click Create Router.

The Create Router window is displayed, where you create a project router with a gateway interface.

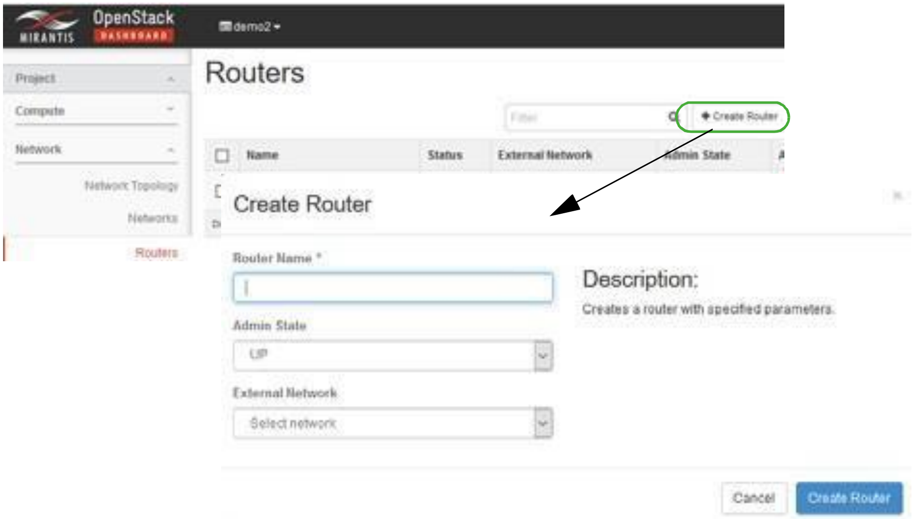

Description:

Creates a router with specified parameters.

Admin State: UP

External Network: Select network

c. Make sure that the router’s gateway IP is reachable from the LAN (your Windows system) following project router creation.

d. Select Project > Network > Networks.

e. Click Create Network.

The Create Network window is displayed.

f. Create and attach the following networks to the project router:

- net-mgmt
- net-data1
- net-data2

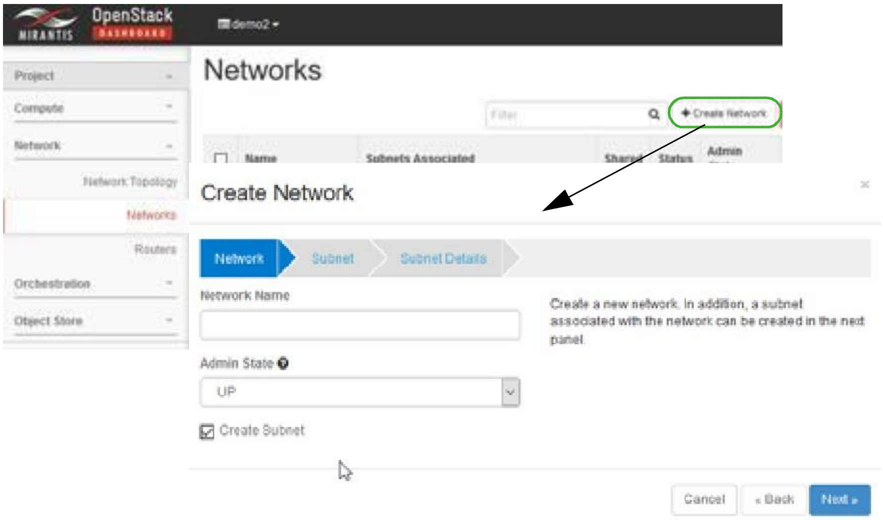

Create a new network. In addition, a subnet associated with the network can be created in the next panel.

Admin State: UP

☑ Create Subnet

Now that your OpenStack project is ready, you can get started by creating a VIM from the Cruz application.

Bringing a VIM Under Management

Before you can deploynetwork servicesorvirtualized network functions (VNFs), you need to bring under management the Virtualized Infrastructure Manager (VIM) to which you want to deploythe network services or VNFs. Bring VIMs under management by:

- Setting Up Authentications
- Defining DiscoveryProfiles
- Adding VIM Resources

Setting Up Authentications

Set up authentications from the Authentication portlet as follows.

1 Right-click > New.

The Creating New Authentication window is displayed.

2 Enter an authentication ID.

3 Enter the remaining information that applies.

4 Save the authentication profile.

5 Repeat steps 1 through 4 for each new authentication required.

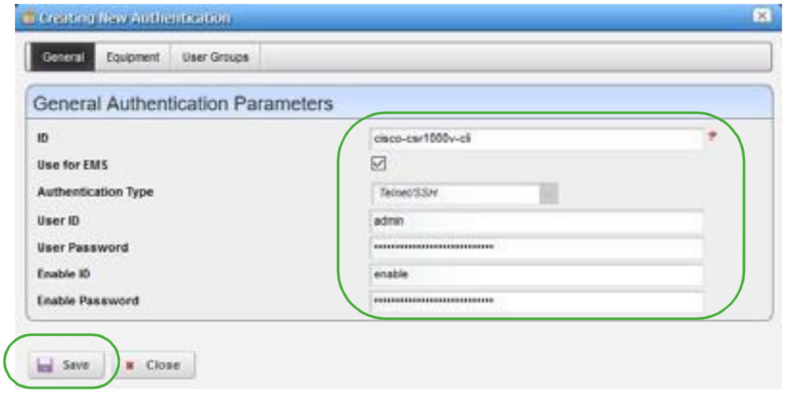

General Authentication Parameters

ID: cisco-csr1000v-cli *
Use for EMS: ☑
Authentication Type: Telnet/SSH
User ID: admin
User Password: ••••••••••••••••••••••••••••
Enable ID: enable
Enable Password: ••••••••••••••••••••••••••••

If you need more details than what is provided here, refer to the Authentication portlet description (Authentications on page 160).

Defining Discovery Profiles

Define discovery profiles from the Discovery Profiles portlet as follows.

1 Right-click > New.

The Creating New Discovery Profile window is displayed.

2 Enter a name and optional description.

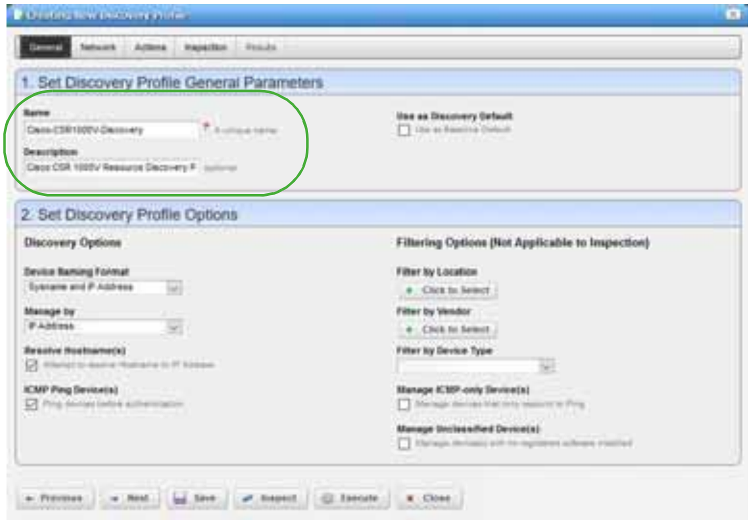

1. Set Discovery Profile General Parameters

Name: Cisco-CSR1000V-Discovery * A unique name

Description: Cisco CSR 1000V Resource Discovery P optional

Use as Discovery Default: □ Use as Baseline Default

2. Set Discovery Profile Options

Discovery Options

Device Naming Format: Sysname and IP Address

Manage by: IP Address

Resolve Hostname(s): ☑ Attempt to resolve Hostname to IP Address

ICMP Ping Device(s): ☑ Ping devices before authentication

Filtering Options (Not Applicable to Inspection)

Filter by Location

Filter by Vendor

Filter by Device Type

Manage ICMP-only Device(s): □ Manage devices that only respond to Ping

Manage Unclassified Device(s): □ Manage device(s) with no registered software installed

3 Select Network.

4 Create a list of authentications.

5 Select authentications for this profile.

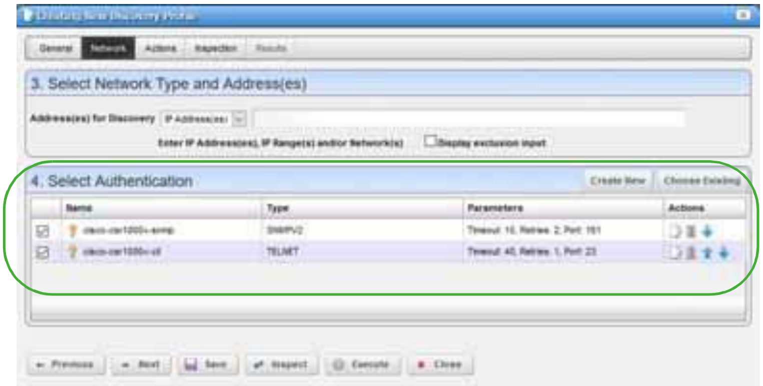

3. Select Network Type and Address(es)

Address(es) for Discovery: IP Address(es)

Enter IP Address(es), IP Range(s) and/or Network(s) □ Display exclusion input

4. Select Authentication

	Name	Type	Parameters	Actions
☑	cisco-csr1000v-snmp	SNMPV2	Timeout: 10, Retries: 2, Port: 161	
☑	cisco-csr1000v-cli	TELNET	Timeout: 40, Retries: 1, Port: 23	

6 Modify the list of actions as needed.

7 Click Execute.

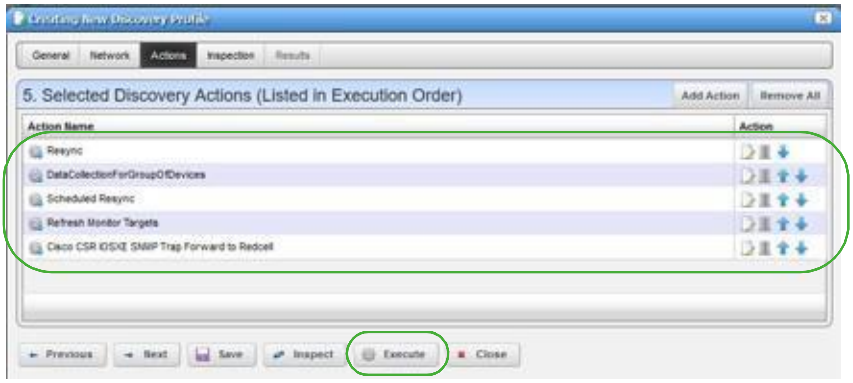

5. Selected Discovery Actions (Listed in Execution Order)

Action Name	Action
Resync	
DataCollectionForGroupOfDevices	
Scheduled Resync	
Refresh Monitor Targets	
Cisco CSR IOSXE SNMP Trap Forward to Redcell	

If you need more details than what is provided here, refer to the Discovery Profile portlet description (Discovery Profiles on page 163).

Adding VIM Resources

Bring a VIM under management from the Virtualized Infrastructure Manager portlet as follows.

1 Right-click > New.

The Editing VIM window is displayed.

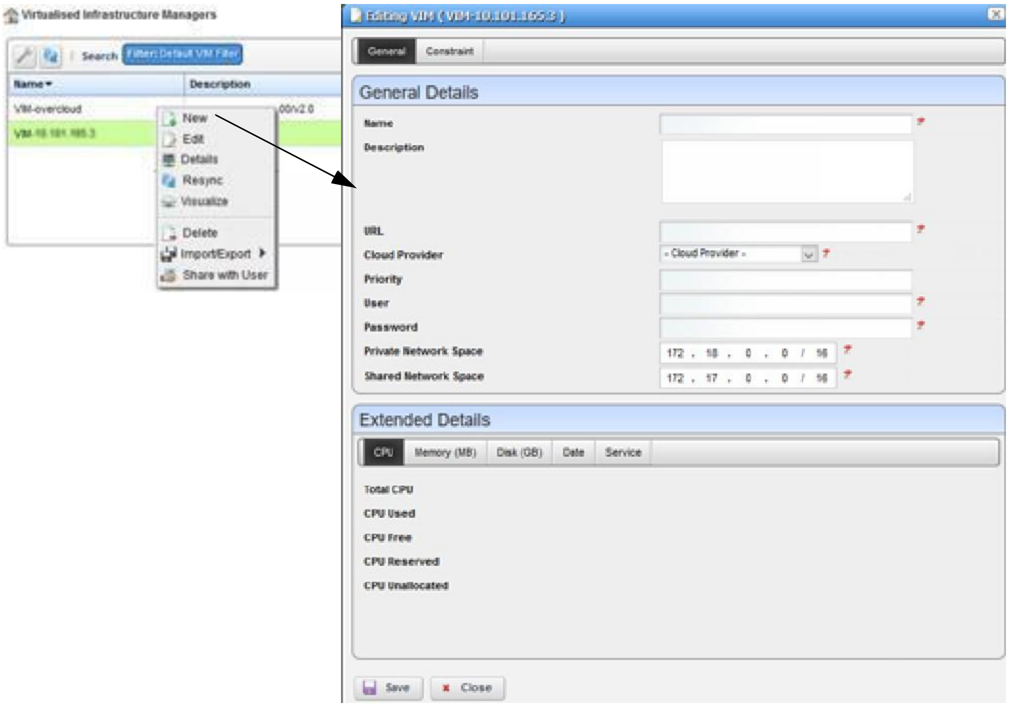

Private Network Space 172 . 18 . 0 . 0 / 16

Shared Network Space 172 . 17 . 0 . 0 / 16

2 Enter the required information.

Optionally, provide a more detailed description and apriority.

3 Specify any constraints.

Constraints specify where to place a network service or VNF if you do not specify a VIM during the staging process.

4 Save the VIM.

This brings the VIM under management.

5 Right-click > Resync the VIM youcreated.

This updates the VIM with the appropriate service models representing the artifacts.

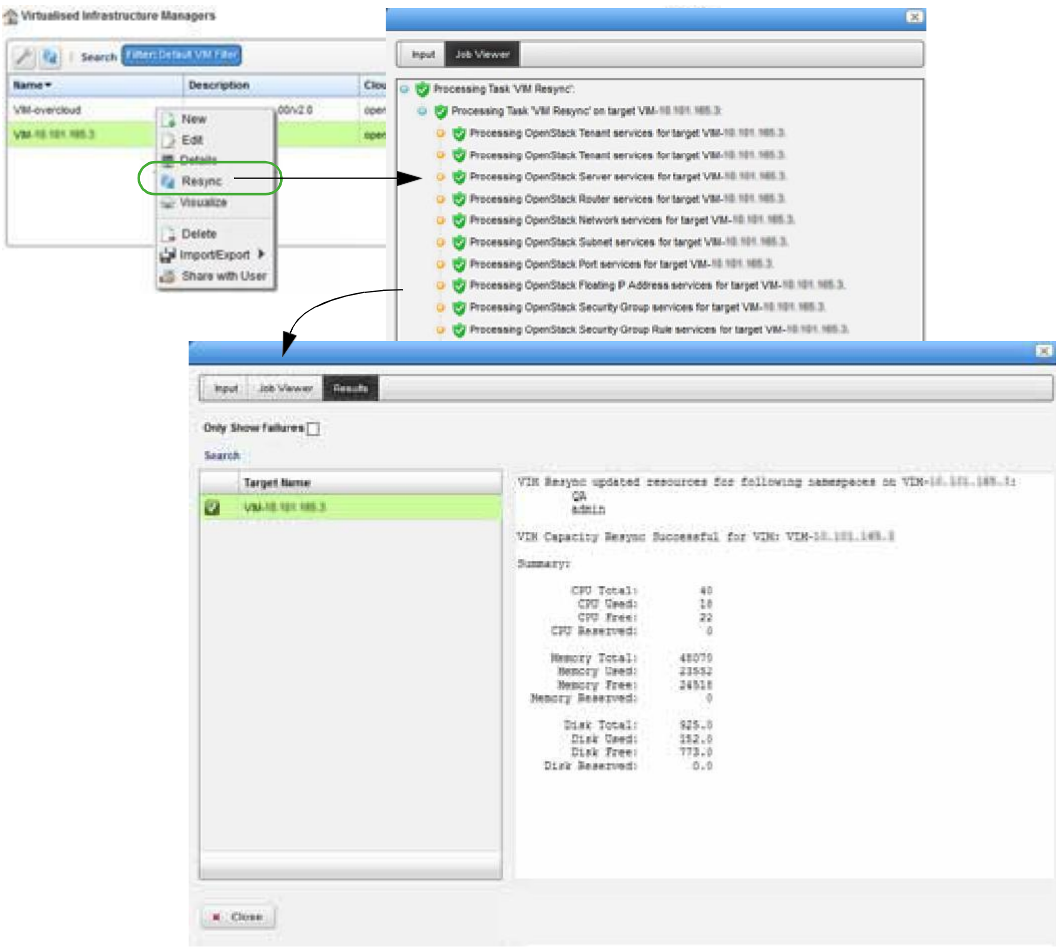

```
Processing Task 'VIM Resync':
  Processing Task 'VIM Resync' on target VIM-[redacted]:
    Processing OpenStack Tenant services for target VIM-[redacted].
    Processing OpenStack Tenant services for target VIM-[redacted].
    Processing OpenStack Server services for target VIM-[redacted].
    Processing OpenStack Router services for target VIM-[redacted].
    Processing OpenStack Network services for target VIM-[redacted].
    Processing OpenStack Subnet services for target VIM-[redacted].
    Processing OpenStack Port services for target VIM-[redacted].
    Processing OpenStack Floating IP Address services for target VIM-[redacted].
    Processing OpenStack Security Group services for target VIM-[redacted].
    Processing OpenStack Security Group Rule services for target VIM-[redacted].
```

☑ VIM-[redacted]

```
VIM Resync updated resources for following namespaces on VIM-[redacted]:
        QA
        admin

VIM Capacity Resync Successful for VIM: VIM-[redacted]

Summary:

        CPU Total:          40
         CPU Used:          18
         CPU Free:          22
     CPU Reserved:           0

     Memory Total:       48070
      Memory Used:       23552
      Memory Free:       24518
  Memory Reserved:           0

       Disk Total:       925.0
        Disk Used:       152.0
        Disk Free:       773.0
    Disk Reserved:         0.0
```

6 View VIM details to see everything related to the VIM, such as:

- Capacity resources (CPU, memory, disk)
- Defined constraints
- Hypervisors
- VIM Images
- Reference details (IP addresses, network, port, router, capacity, and so on)

7 Repeat steps 1 through 6 for each new VIM.

Setting Up Cruz Application

Before you perform the tasks described in this guide, make sure that the NFV-specific portlets you need are added to the existing application.

Which portlets you need, depends on your role (Table 16-2) in the management ofyour virtualized network environment and the deployment of network services and virtualized network functions (VNFs).

Table 16-2. User Roles and Tasks

Role	Tasks
Administrator	Bringing VIM under management Setting up the OpenStack environment (uploading vendor base images to the OpenStack controller, modifying resource flavors, creating VIM snapshots) Preparing vendor-supplied images from Cruz (creating and editing VIM images) Maintaining the VIMs under NFV resource management, descriptors, virtual reservations, and virtual requirements
Network Designer/ Engineer	Defining descriptors Implementing virtual reservations Monitoring virtual requirements
Operator	Instantiating VNFs (stage, deploy) Monitoring virtual requirements Performing other daily operations, such as undeploying, scaling, monitoring system health, resolving issues found, *etc.*

Depending on how your company wants to define the work environment for all users, you have the option to add the Network Virtualization applications (portlets) to existing Cruz pages or create pages and add portlets more appropriately for each users work environment. Once the pages/ portlets setup is complete, verify this setup.

Here is a list of NFV portlets (Table 16-3) and how each portlet is used by the different user roles. For a detailed description of these portlets, see Network Virtualization Portlets on page 866. If any of these portlets are not available, a message is displayed when you try to add it to a page.

Table 16-3. NFV Portlets (Sheet 1 of 2)

Portlet	Operator	Administrator	Network Designer
License Accounts	Maintain a list of licensed accounts	Maintain a list of licensed accounts	Maintain a list of licensed accounts
License Descriptors	View and understand license descriptors	Maintain available license descriptors	Define licensing integration and characteristics
License Records	View and understand license records	View and understand license records	View and understand license records
NFV Monitoring Attributes	View and understand NFV monitoring attributes	Maintain monitoring attributes	Maintain monitoring attributes
Network Service Descriptors	View and understand NSDs	Maintain network service descriptors	Define monitoring attribute characteristics
Network Service Records	Stage, deploy, and undeploy (manage) network services	No activity performed by the administrator	Test an implemented network service
OSS Instance	View and understand OSS instances	Implement OSS instances	No activity performed by the network designer

Table 16-3. NFV Portlets (Sheet 2 of 2)

Portlet	Operator	Administrator	Network Designer
Physical Network Function Descriptors	View and understand PNFDs	Maintain PNF descriptors	Implement physical network functionality
Physical Network Function Records	Stage, deploy, and undeploy (manage) PNFs	Test new or modified descriptors	Test implemented PNFs
SDN Controllers	View and understand SDN controllers	Maintain software-defined network (SDN) controllers.	Test SDN controllers
Software Images	View and understand software images	Create a software image descriptor for a snapshot image, deploy software images, and edit the VIM software descriptor parameters	Create a software image descriptor for a snapshot image, deploy software images, and edit the VIM software descriptor parameters
VIM Images	View and understand VIM images	Maintain consistency across the VIM instances	Test that software images deployed/undeployed were added/removed from the VIM images list
Virtual Network Function Descriptors	View and understand VNFDs	Maintain VNF descriptors	Implement virtualized network functions
Virtual Network Function Records	Stage, deploy, and undeploy (manage) VNFs	No activity performed by the administrator	Test an implemented network service
Virtual Requirements	View virtual domain resource requirements and usage (such as memory, CPU, and disk). You also have the option to modify a domain’s description.		
Virtual Reservations	Verify that a VDU’s resources were reserved after they stage a VNF record	Maintain virtual reservations	Implement virtual reservations
Virtualized Infrastructure Managers	View available VIMs and their resources before deploying services or VNFs	Maintain the VIMs under resource management	No activity performed by the network designer

Adding a Page

Pages allow you to organize commonly used portlets into a single location for easy access or to group portlets by tasks you perform. For example, group tasks used to monitor and resolve network health on a single page. You have the option to add NFV pages to the me nu bar, add child pages to existing pages, or a combination of both.

Add any needed pages to the application as follows.

1 Select Manage > Page.

The Manage Page window is displayed.

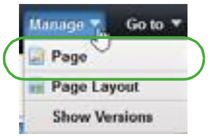

2 Add pages you want to access from the navigation panel.

a. Select Private Pages.

b. Specify the look and feel.

c. Click Add Page.

The Add Page window is displayed.

d. Enter page name.

e. Select a type.

f. Click Add Page.

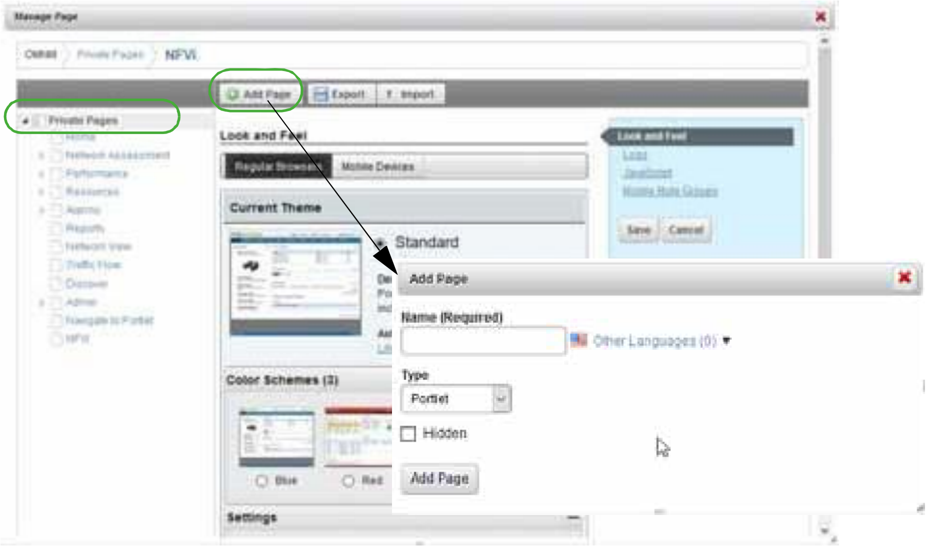

Name (Required)

Type: Portlet

□ Hidden

3 Repeat steps 1 and 2 for each page you add to the navigation panel.

4 Add a child page to an existing page.

a. Select Manage > Page.

The Manage Page window is displayed.

b. Select the page under which to add the child page.

c. Click Add Child Page.

The Add Child Page window is displayed.

d. Enter page name.

e. Select the type.

f. Select whether to copy the parent or make the page hidden.

g. Click Add Page.

h. Repeat these steps for each child page you want to add.

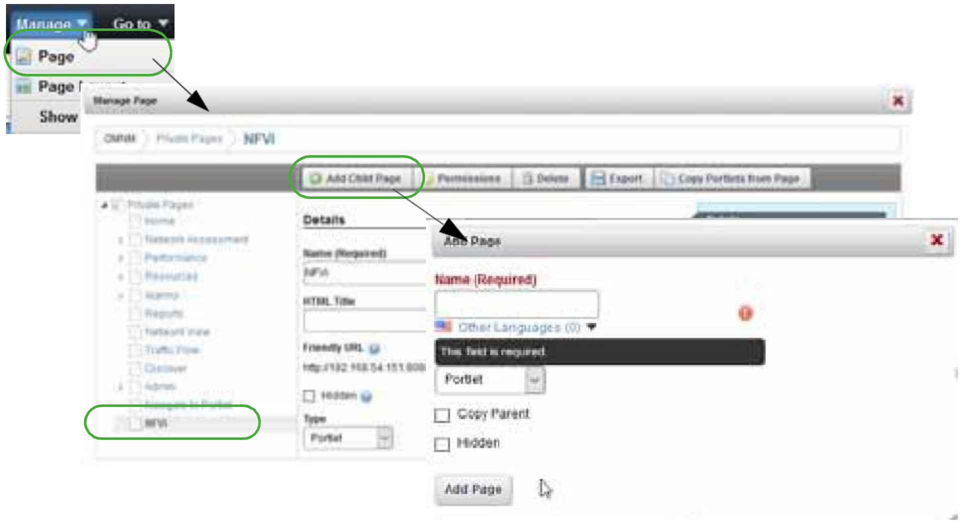

Now you are ready to add portlets (applications) to the appropriate pages.

Adding Portlets to a Page

These instructions assume that the pages to which you plan to add portlets already exist. If the pages do not exist, see Adding a Page before continuing.

Add portlets to a page from the application as follows.

1 Click the page label to which you want to add a portlet.

The selected page is displayed.

2 Select Add > Applications.

The applications list is displayed.

3 Expand the Network Virtualization list.

See Setting Up Cruz Application on page 726 for a brief description of each portlet, or Network Virtualization Portlets on page 866 more details.

4 Click Add for each portlet you want to add to the selectedpage.

Note that the purple indicator means that you can add the portlet only once to a page. All other portlets you can add multiple instances to a page.

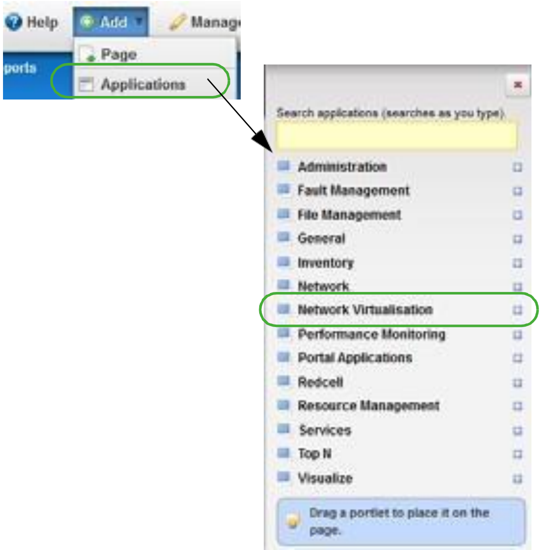

5 Refresh the page to show the portlet.

If you do not have permissions to view the portlet, a message is displayed. Delete the portlet or consult your system administrator if you feel you should have access to the portlet.

6 Repeat steps 4 and 5 for each portlet you want to add to the selected page.

7 Rearrange the portlets as needed using the drag-and-drop method.

8 Repeat steps 1 through 7 for each page to which you want to add portlets.

Testing Your Cruz Setup

You should run **some** tests tomake sure you have what you need, and you can perform some basic tasks, such create, edit, view details, delete, and so on.

Test your setup from the application as follows.

1 Sign out and then sign back in if you are already logged in.

2 Verify that your pages are accessible from the menu bar.

3 Make any necessary changes.

4 Make sure that you can do some basic tasks (add, edit, delete, and view details) from each added portlet, such as these:

- Network Service Descriptors
- Physical Network Function Descriptors
- Software/VIM Images
- Virtual Network Function Descriptors
- Virtual Requirements/Reservations
- Virtualized Infrastructure Managers

5 Navigate to the Virtual Network Function Records portlet.

a. Verify that you can **discover** VNF records.

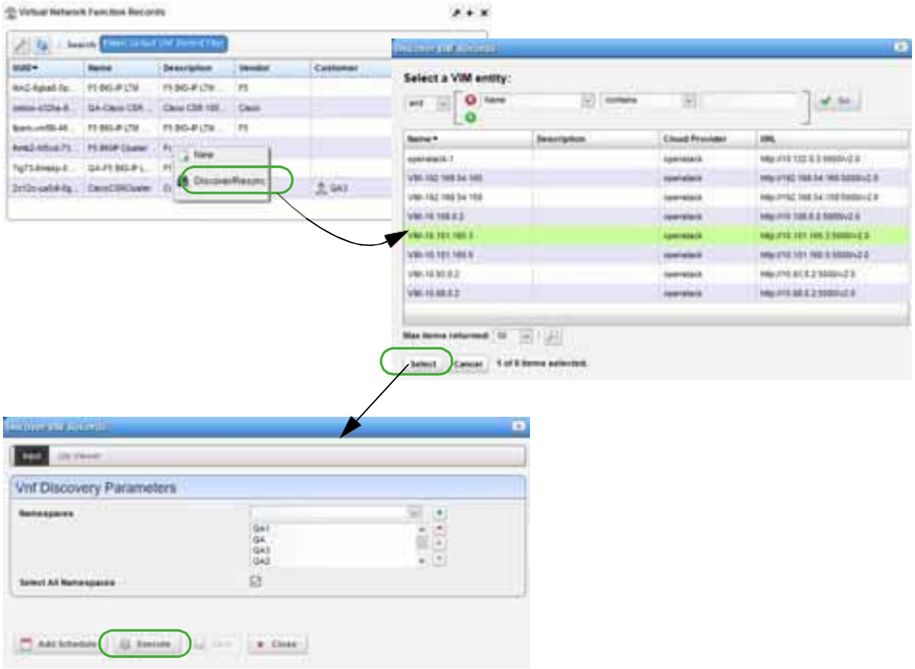

b. Verify that you can **stage** a record.

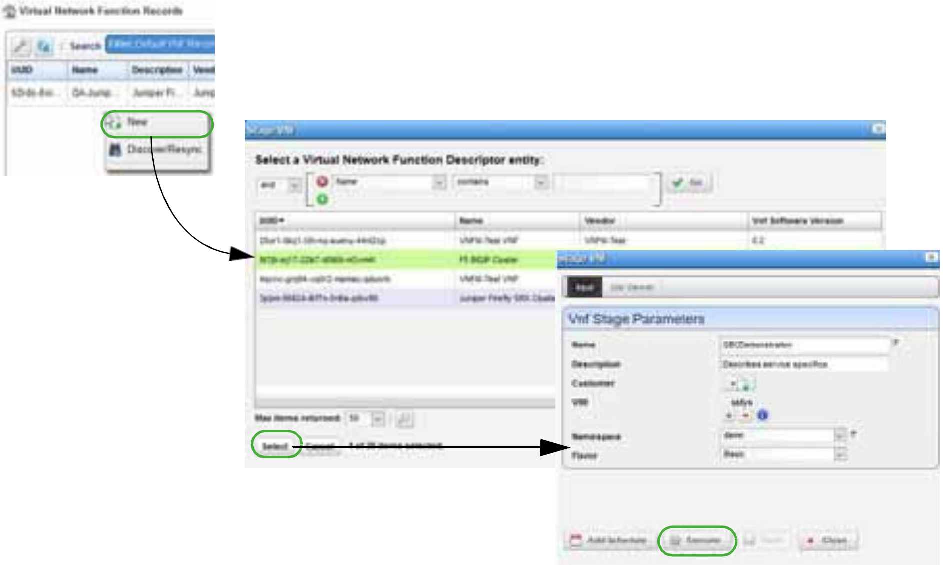

c. Verify that you can **deploy** a staged record staged.

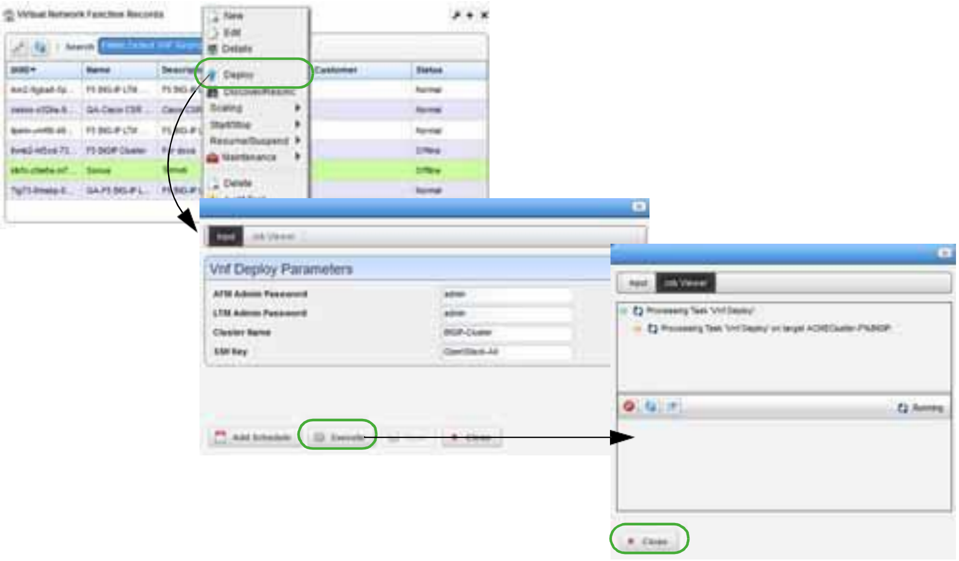

d. Verify that you can **modify** a record.

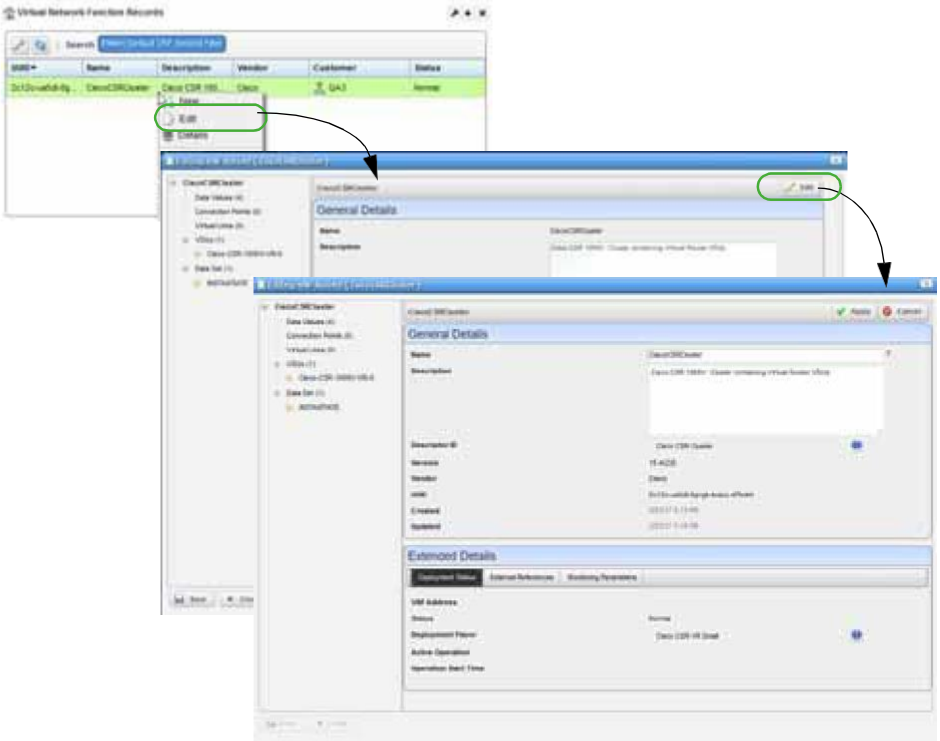

e. Verify that you can **undeploy** the record deployed in step c.

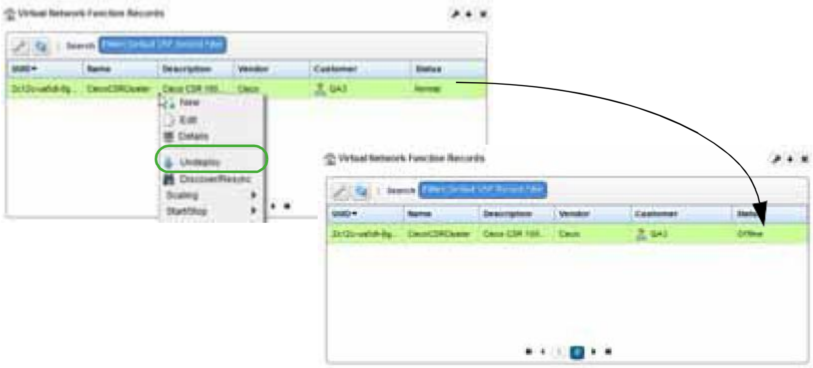

f. Verify that you can **delete** a record.

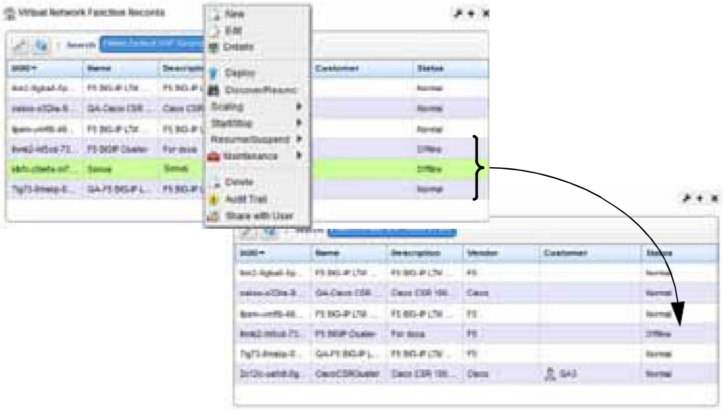

Determining What to Do Next

Nowthat you havepreparedyour system to configure your NFVinfrastructure, you youarereadyto get started with the following tasks that apply to your environment (Table 16-4).

Table 16-4. Configuration Tasks (Sheet 1 of 2)

Task	High-Level Steps	Detail Steps
Create Cisco CSR100V VNF descriptor	1. ReviewVNF Packagecontent. 2. Deployvendor-providedimage. 3. Create resourceflavors. 4. CreateVMsnapshots. 5. Create softwareimagedescriptor. 6. ModifyVIMsoftware imagedescriptor.	Setting Up Cisco CSR100V VNF Package

Table 16-4. Configuration Tasks (Sheet 2 of 2)

Task	High-Level Steps	Detail Steps
Create F5 BIGIP VNF descriptors	1. ReviewVNFpackagecontent. 2. Createan OpenStackproject. 3. Deploy vendor-providedimage. 4. Modify running VM booted fromoriginal image 5. Create resourceflavors. 6. Create VMsnapshots. 7. Createsoftware image descriptor. 8. ModifyVIM softwareimagedescriptor. 9. Applylicense.	Setting Up F5 BIGIP VNF Package .
Create Juniper Firefly VNF descriptor	1. ReviewVNFpackagecontent. 2. Convertvendor. jvato qcow2 format. 3. Createa project. 4. Deploy vendor-providedimage. 5. Create resourceflavors. 6. Modify snapshotimagechanges. 7. Create VMsnapshots. 8. Createsoftware imagedescriptor. 9. ModifyVIM softwareimagedescriptor.	Setting Up the Juniper Firefly VNF Package
Bring Sonus SBC under management	1. ReviewVNFpackagecontent. 2. Make sure thatyouunderstandthe VNF configurationmethods. **Note:** Sonus VSBCVDUsare automatically brought under Cruz management.	Understanding the Sonus VSBC VNF Package
Manage resources	1. Create, modify, and/ordelete VIMs. 2. Resyncing VIMs. 3. DiscoveringVNFrecords. 4. Managing softwareimages. 5. Maintaining resourcemonitors.	Managing NFVI/VIM Resources
Manage descriptors	1. Createor Importdescriptors. 2. Modify descriptorsas needed. 3. Remove descriptors no longer needed. **Note:** Thiscovers networkservice,VNF, and PNF descriptors.	Managing Descriptors
Automate deployment	**Note:** Only includes referenceinformationat this tim	eAutomating SMB vCPE Deployment

Setting Up Cisco CSR100V VNF Package

Thissectionexplainsthe scope, functions, and processesassociated with Virtualized Network Functions (VNF) onboarding and instantiation.

- Understanding the Cisco CSR100V VNF Package
- Preparing an Image

Understanding the Cisco CSR100V VNF Package

The Cisco CSR 1000v model is a virtual router that can be deployed in the OpenStack environment. Cloud Services Router (CSR) software images are downloadable and portable between on-premises virtualized data centers, public, and hybrid cloud environments. Th e CSR allows you to rapidly provision routers as needed across the data center and into the cloud.

The Cisco CSR100V VNF package provides a sample VNF Descriptor and associated artifacts to deploy the Cisco CSR 1000v virtual router.

Cisco Systems® provides qcow2 images (.qcow2) and licenses that are readyfor the OpenStack environment. You need these files to deploy and enable the Cisco CS R 1000v features.

The following list provides the VNF package vendor and validation information:

Vendor and Model Information	
Vendor	Cisco Systems, Inc.
Product Family	CSR
Models	Type 1: 1000v
Operating Systems	IOSXE
Version	15.4(2)S
Validated Environment	
VIM	OpenStack
Version	Liberty
Installer	Mirantis
Version	8
Network Configurations	VXLAN, VLAN Segmentation
Cruz **MANO (Management and Operation) Validated Functions**	
MANO	NSD, VNFD, NSR, VNFR Link LCEs
VNFM	Generic (MANO Lite) Independent/Specific
NFVO	MANO Lite
Cruz **RMO (Resource Management and Operation) Validated Functions**	
Inventory	Resource Discovery and Inventory Resource Creation Method (NFV-Model)
Monitoring	SNMP Monitoring CPU & Memory KPIs TFA (flows)
Config	ACLI

NetRestore File Management	NR Backup NR Restore NR Firmware (not validated)

This section covers the:

- Cisco CSR VNF Package Contents
- Cruz VNF Package - Install Component Files
- License Requirements
- Known Issues
- References

Once you have a good understanding of the VNF package, you can get started preparing an image .

Cisco CSR VNF Package Contents

The Cisco CSR (Cloud Services Router) VNF package contains artifacts necessary to install and deploy Virtualized Network Functions (VNFs) for the Cisco CSR 1000v type Virtualization Deployment Unit (VDU).

The sample VNF descriptor created has basic flavors to deploy the VNF VDU as part of a standalone VNF record. The Network Service descriptor is not currently available for the CSR software.

The sample Cisco CSR Cluster VNF descriptor is located in the following file:

```
owareapps/nfv-vnfm-cisco/db/vnf.cisco-csr-cluster-v1.xml
```

This sample VNF descriptor defines the packaging and behavior of one or more Cisco CSR 1000v VDUs and can be included as part of a network descriptor.

Here are the **Cisco CSR Cluster** descriptor deploy parameters used:

Field Name	Data Type	Description	Default	Valid Values	Required
Cluster Name	String	VNF’s Cluster Name	Cisco-Cluster	String	N
DNS IP	IP Address	IP of DNS Server	8.8.8.8	IP Address	N

Here are the **Cisco CSR Cluster** descriptor flavors used:

Flavor Name	VDU Name	VDU Count	CPU	Disk (GB)	Memory (MB)
Cisco CSR VR Small	Cisco-CSR-1000V-VR	1	2	0	4096

View the VNF descriptor’s details from th e Cruz VNF Descriptor portlet.

Cruz VNF Package - Install Component Files

The following Virtualized Network Function (VNF) package installation components are required for the Cisco CSR100V product:

VNF (OCP)	nfv-vnfm-cisco.ocp
Device Driver (DDP)	cisco.ddp

The Cruz VNF package descriptors for the **Cisco CSR100V** VNF product includes:

NVF VNF Package			Resource Management		
VIM Software Image Records	VNF Descriptors	Sample NSD	Discovery Authentication Records	Discovery Profiles	Device Driver
Y	Y	N	Y	Y	Y

The VNF package descriptors include:

- Discovery Authentication Profiles
- Discovery Profiles
- NFV VIM Software Image Descriptors

Discovery Authentication Profiles

The Cisco CSR100V Discovery Authentication profiles are located in the following file:

`owareapps/nfv-vnfm-cisco/db/rc.disc-auths-cisco.xml`

This file contains the following discovery authentication profiles:

Authentication Record Name	Description
cisco-csr1000v-cli	Telnet Credentials
cisco-csr1000v-snmp	SNMP v1/v2 Credentials

Discovery Profiles

The Cisco CSR100V Discovery profiles are located in the following file:

`owareapps/nfv-vnfm-cisco/db/rcdisc-profile-cisco.xml`

Thisfilecontains the Cisco-CSR1000v-Discoveryprofile that automatically discoversand populates resource records during VNF recordinstantiation.

The following tasks are run once the target VM in the OpenStack environment is fully booted and can respond to SNMP and CLI requests from the management system:

1. Resync
2. DataCollectionForGroupOfDevices
3. Scheduled Resync
4. Refresh Monitor Targets
5. Cisco CSR IOSXE SNMP Trap Forwarding to Cruz

NFV VIM Software Image Descriptors

The original Cisco vendor image was used to produce the following workable snapshot image used to create Network Functions Virtualization (NFV) Virtualized Infrastructure Manager (VIM) software image descriptors from the application. See Preparing an Image on page 740 for detailed instructions.

Name	Scope/Description	Container Format	Disk Format	Min RAM (MB)	MIN Disk (GB)	Resource Discovery Profile
CiscoCSR1000v-snapshot-2-16Ports	Dorado Snapshot of original csr1000v-universalk9 .03.12.00.S.154-2.S-std plus 16 Interfaces per iosxe-cfg-16Ports.txt file	BARE	QCOW2	4096	0	Y

License Requirements

Contact your vendor for licensing requirements.

Known Issues

The OpenStack flavor for the Cisco CSR 1000v VMs must contain a value of 0 for the Root, Ephemeral, and Swap Disk fields. Otherwise, the VM does not complete the bootup process.

References

This section provides references to:

- Vendor Documentation (URLs)
- VNF Configuration Methods

Vendor Documentation (URLs)

Refer to the Cisco website (http://www.cisco.com) forallvendor-specificdocumentation.

For the Cisco CSR 1000v-Series datasheet, refer to the following Cisco website page:

`http://www.cisco.com/c/en/us/products/collateral/routers/cloud-services-router-1000v-series/datasheet-c78-733443.html`

VNF Configuration Methods

Acombination of image (snapshot) preparation, along with Virtualization Deployment Unit (VDU) Post Instantiate Configuration Lifecycle Events (LCEs), is implemented in this VNF package to ensure the Cisco CSR VDU is automatically brought under Cruz management.

Customized Image Preparation. The Cisco CSR image is provided by Cisco Resource Management applications.

See Preparing an Image on page 740 for a description of how to create an updated image (snapshot) for the Cisco CSR 1000v model series to use to boot up VMs that are managed by the Cruz application. The snapshot produced is used to create Cisco CSRs in the OpenStack environment, so that no manual intervention is required to apply the management system communication settings.

Cruz Resource Configuration Tasks. Once the target VM in the OpenStack environment is fully booted and can respond to SNMP and CLI requests from the management system, post discovery tasks are run.

Formore details about the Cisco-CSR1000v-Discoveryprofile and the tasks that run, see Discovery Profiles on page 738.

Preparing an Image

Before you can discover and manage a Virtualized Network Function (VNF) resource from the application, you need to:

1. Set up your OpenStack environment by preparing a Cisco CSR image with the necessary configuration.
2. Create a software image record from the Cruz application.

Setting Up Your OpenStack Environment

To set up your OpenStack environment, you need to:

1. Deploy the vendor's base image by uploading the original qcow2 image file received from your Cisco supplier to a target Bare Metal OpenStack controller.
2. Create flavors representing the smallest supported flavor to use. The vendor (Cisco) provides a flavor (Figure 16-3).
3. Create VM snapshots that are used to create and export a VIM software image descriptor within the Cruz application.

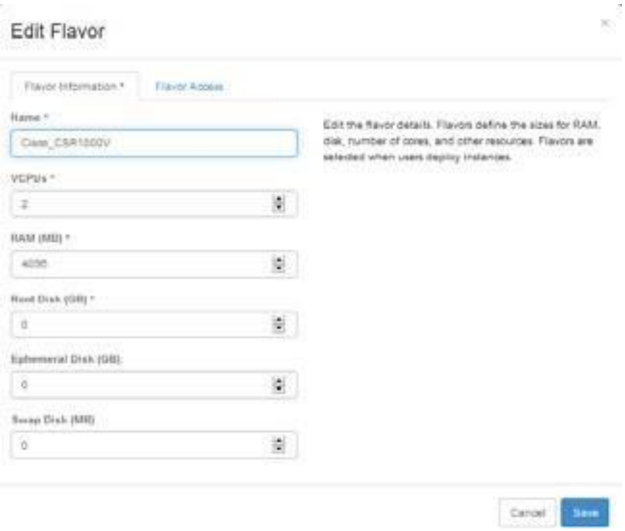

Name *: Cisco_CSR1000V

Edit the flavor details. Flavors define the sizes for RAM, disk, number of cores, and other resources. Flavors are selected when users deploy instances.

VCPUs *: 2

RAM (MB) *: 4096

Root Disk (GB) *: 0

Ephemeral Disk (GB): 0

Swap Disk (MB): 0

Figure 16-3. Cisco_CSR1000v Flavor

Instantiating a VM From a Vendor Image

Instantiate a VM (virtual machine) inyour OpenStackenvironment fromthevendor-provided image as follows.

1 Create a default Cisco IOS-XEconfiguration file (iosxe_config.txt) with your operational settings for your Cisco CSR 1000v virtual routers.

Refer to your Cisco product documentation for instructions on how to create and populate a Cisco configuration file for your Cisco CSR VNFs.

a. Add configuration for the number of interfaces you need to utilize (GigabitEthernet **1** GigabitEthernet *N*).

The following example shows configuration for four interfaces:

```
interface GigabitEthernet1
 description "Data Port 1"
 ip address dhcp
 no shutdown
!negotiation auto
!
interface GigabitEthernet2
 description "Data Port 2"
 ip address dhcp
 no shutdown
!negotiation auto
!
interface GigabitEthernet3
 description "Data Port 3"
 ip address dhcp
 no shutdown
!negotiation auto
!
interface GigabitEthernet4
 description "Data Port 4"
 ip address dhcp
 no shutdown
!negotiation auto
```

b. Configure more GigabitEthernet interfaces than areattached to the VM when booted later in from the OpenStack system.

This allows you to use the same image in multiple network topologies that have a varying number of networks to which you will attach the Cisco CSR 1000v virtual routers.

c. Name the file iosxe_config.txt.

d. Upload the iosxe_config.txt file to your target's local disk under the /rootdirectory.

2 Obtain an image from your vendor in qcow2 format.

3 Upload the qcow2 image to the OpenStack VIM.

For example, Cisco provided the following image file for the Cisco CSR 1000v VNF package:

```
csr1000v-universalk9.03.12.00.S.154-2.S-std.qcow2
```

4 Name the image record in the OpenStack system as `CiscoCSR1000v`.

5 Double-click the image from the Project > Compute > Images page.

The image details overview is displayed.

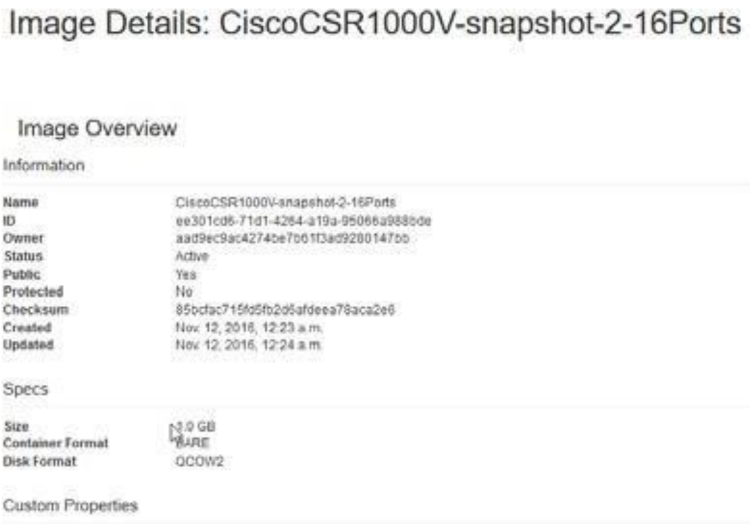

Image Details: CiscoCSR1000V-snapshot-2-16Ports

Image Overview

Information

Name	CiscoCSR1000V-snapshot-2-16Ports
ID	ee301cd6-71d1-4264-a19a-95066a988bde
Owner	aad9ec9ac4274be7b61f3ad9280147bb
Status	Active
Public	Yes
Protected	No
Checksum	85bcfac715fd5fb2d5afdeea78aca2e6
Created	Nov. 12, 2016, 12:23 a.m.
Updated	Nov. 12, 2016, 12:24 a.m.

Specs

Size	1.0 GB
Container Format	BARE
Disk Format	QCOW2

Custom Properties

6 Boot up a VM from the OpenStack system using this image.

Here is an example command syntax:

```
Image file Used:\\192.168.51.11\it\VMs\vnf\vendor\cisco\CSR\csr1000v-
  universalk9.03.12.00.S.154-2.S-std.qcow2
```

```
Config file used - placed on OpenStack controller 10.101.170.3 (bare metal
  ICEHOUSE via Fuel) in the root directory:
```

```
\\192.168.51.11\it\VMs\vnf\vendor\cisco\CSR\iosxe_config_MgtPlus8DataPort
  s.txt -> copied to /root/iosxe_config_MgtPlus8DataPorts.txt on the
  OpenStack Controller.
```

7 Log into the OpenStack controller.

8 Boot a new VM using the following nova command linesyntax.

Make sure that you are using the image along with your iosxe_config.txt file and that the instance is operational.

```
nova --os-username admin --os-password <OpenStack admin password> --os-
  <OpenStack Tenant Name> admin --os-auth-url <OpenStack Auth URL> boot
  <VM Name> --image <OpenStack Image Name> --flavor <OpenStack Flavor
  Name> --nic net-id=<OpenStack Network 1 UUID> --nic net-
  id=<OpenStack Network 2 UUID> --config-drive=true --file
  iosxe_config.txt=<default config file disk location/filename>
```

Thisexample command hastwo network connections. Youcan include additional network connections by adding repetitions of the following:

```
--nic net-id=<value>
```

9 Verify that the VM:

- Booted successfully.
- Has the proper Cisco configuration file contents applied to its running configuration
- Is fully operational.

10 Suspend the VM instance from the OpenStack dashboard.

11 Take a snapshot of the instance.

12 Boot a new VM using the snapshot.

13 Make sure that the VM is operational and has the proper contents from the original configuration file created in step 1 applied to its running configuration.

Create a flavor before creating the VM snapshots.

Creating a Flavor

Cisco provides a flavor. However, VIMs require some additional settings to boot properly.

Create a flavor to use from the OpenStack dashboard as follows.

△ CAUTION:

It is critical that the Root, Ephemeral, and Swap Disk settings are all set to 0. Otherwise, the VMs created later using this flavor and the Cisco CSR 1000v image do not boot properly.

1 Select Admin > System > Flavors.

2 Click Create Flavor.

The Create Flavor window is displayed.

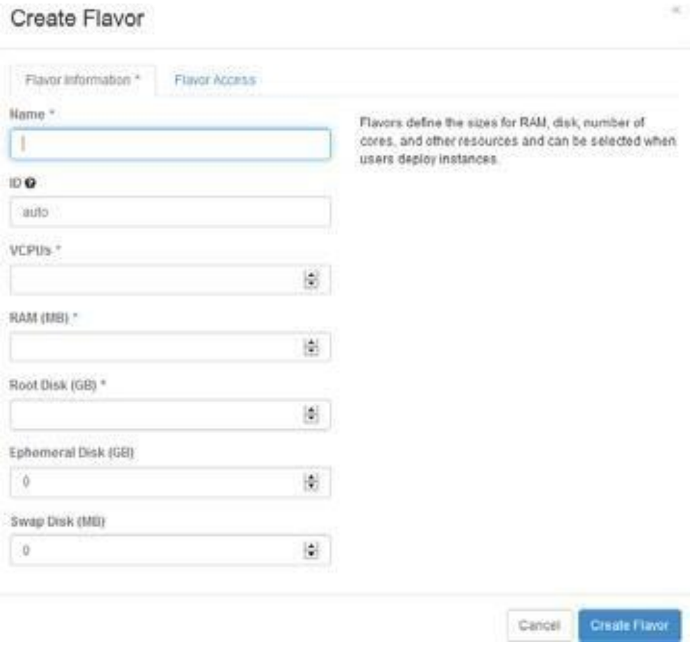

Create Flavor

Name *

ID ❓ auto

VCPUs *

RAM (MB) *

Root Disk (GB) *

Ephemeral Disk (GB) 0

Swap Disk (MB) 0

Flavors define the sizes for RAM, disk, number of cores, and other resources and can be selected when users deploy instances.

3 Enter a name, such as Cisco_CSR 1000v, and the number of VCPUs.

4 Set the memory required to 4096.

5 Set Root, Ephemeral, and Swap Disk values to zero.

6. Specify which projects to apply the flavor.

   If you do not specify a project, the flavor is available to all projects.

7. Click Create Flavor.

   Your flavor is created.

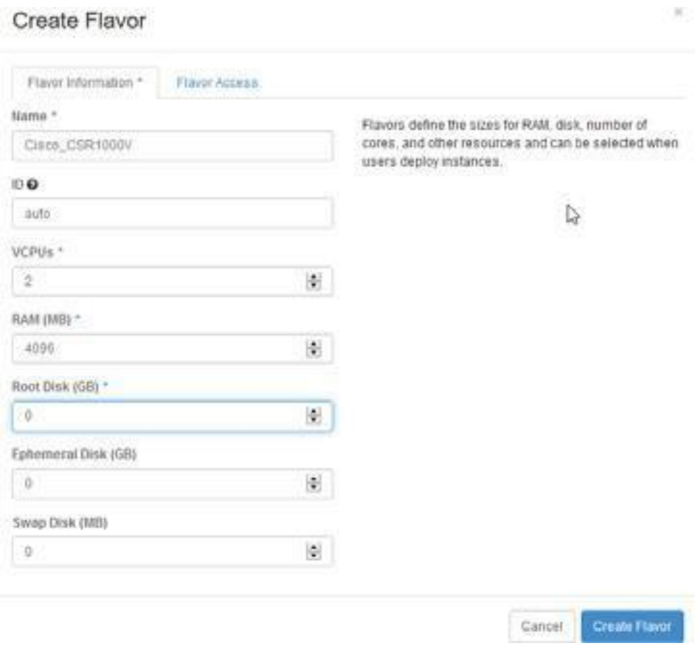

Create Flavor

Name *: Cisco_CSR1000V

ID: auto

VCPUs *: 2

RAM (MB) *: 4096

Root Disk (GB) *: 0

Ephemeral Disk (GB): 0

Swap Disk (MB): 0

Flavors define the sizes for RAM, disk, number of cores, and other resources and can be selected when users deploy instances.

Next you need to create a VM snapshot.

Creating a VM Snapshot

The VM snapshot is used to create and export a VIM software image descriptor from the Cruz application and then the VIM software image descriptor is used to manage, deploy, or migrate images.

Create a VM snapshot from the OpenStack dashboard as follows.

1. Shut off the instance.

   a. Select Project > Compute > Instances.

      A list of instances is displayed.

   b. Select the Shut Off Instance action for the appropriate instance.

      A confirmation message is displayed.

   c. Click Shut Off Instance.

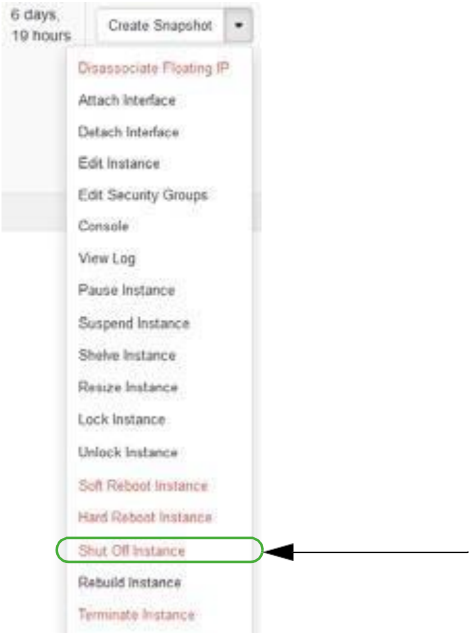

If you prefer to shut down the Cisco CSR 1000v VM from the Cisco CLI, refer to your Cisco documentation for instructions.

2 Create a snapshot image.

a. Select Admin > System > Instances.

A list of instances is displayed.

b. Select the Create Snapshot action for the instance.

The Create a Snapshot window is displayed.

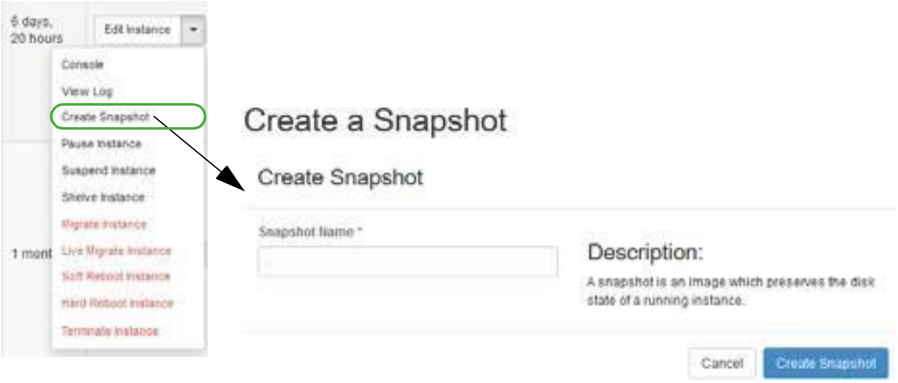

c. Enter a name, such as `CiscoCSR1000v-snapshot` followed by any additional text based on the configuration file you created and applied to the VM (such as Cisco IOS-XE).

d. Click Create Snapshot.

3 Assign OpenStack image properties to the snapshot image created in step 2.

a. Select Admin > System > Images.

b. Select the Edit Image action for the instance.

The Update Image window is displayed.

△ CAUTION:
It is critical that the Root, Ephemeral, and Swap Disk settings are all set to 0. Otherwise, the VMs created later using this flavor and the Cisco CSR 1000v image do not boot properly.

c. Set the disk and memory requirements to zero (0).

d. Select Public.

This is required so that the Cruz can access and manage the image later on.

e. Click Update Image.

Your changesare saved.

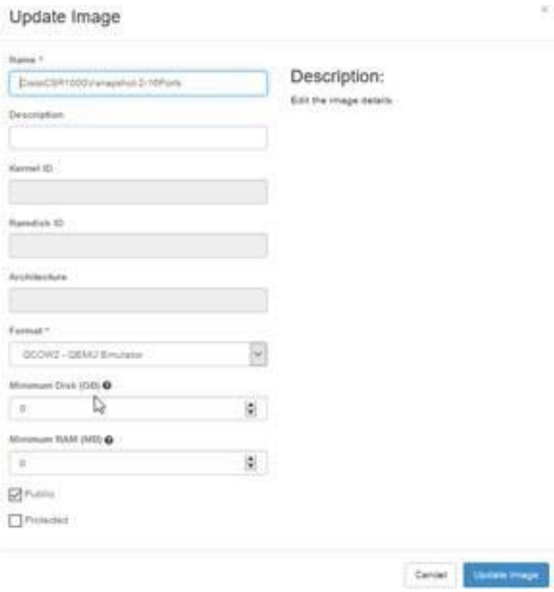

Update Image

Name *: CiscoCSR1000Vsnapshot-2-10Ports

Description:

Kernel ID:

Ramdisk ID:

Architecture:

Format *: QCOW2 - QEMU Emulator

Minimum Disk (GB): 0

Minimum RAM (MB): 0

☑ Public

☐ Protected

Description:

Edit the image details

Once you create the VM snapshot image in the OpenStac k environment, you are ready to create a software image descriptor from the Cruz application.

Creating a Software Image Descriptor

The Network Functions Virtualization (NFV) feature uses the software image descriptor to manage, deploy, or migrate images. Create a software image descriptor record from the Cruz application by:

1 Creating a Software Image Descriptor for a Snapshot

2 Modifying the VIM Software Image Descriptor

The steps provided in this section assume that you have already created a VIM that houses the snapshot image in the OpenStack environment. If a VIM that houses the snapshot image does not exist, see Setting Up Your OpenStack Environment on page 740 before continuing.

Creating a Software Image Descriptor for a Snapshot

Create a software image descriptor for a snapshot image from the Cruz application as follows.

1 Navigate to the Software Images portlet.

The portlet’s location depends on your company’s configuration.

2 Right-click an image.

3 Select Download.

The Save VIM Image to Disk Parameters input window is displayed.

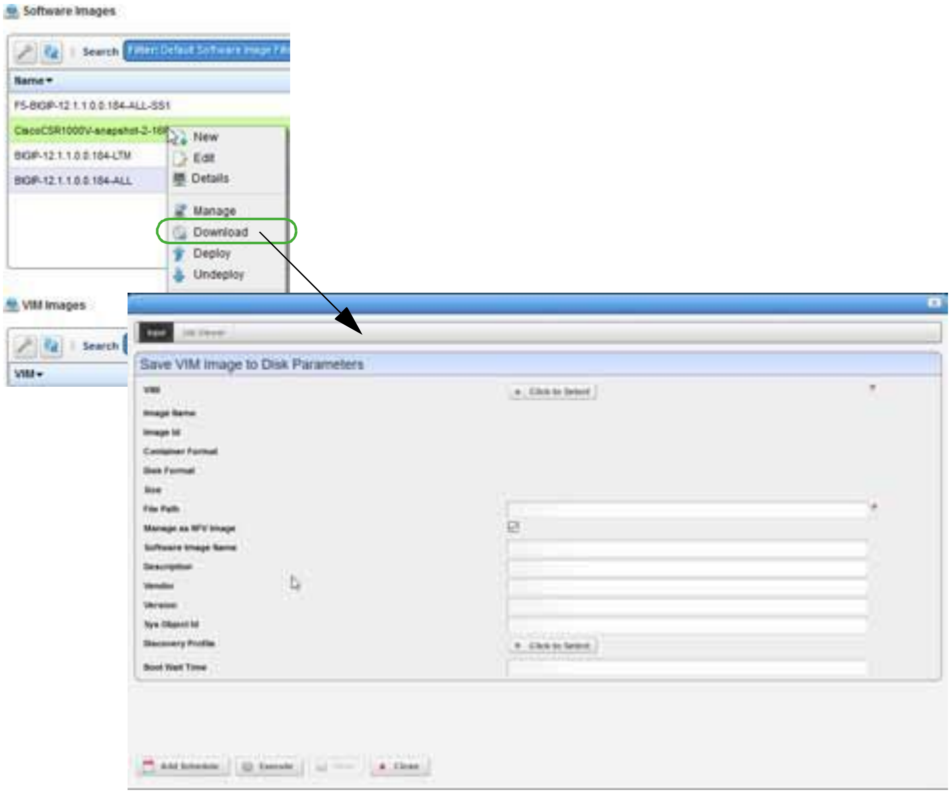

4 Select the Virtualized Infrastructure Manager (VIM) housing the CiscoCSR1000v-snapshot-2-16Ports image created earlier.

5 Select the snapshot image to download.

For example: `CiscoCSR1000v-snapshot-2-16Ports.qcow2`

The fields are populated based on your selection.

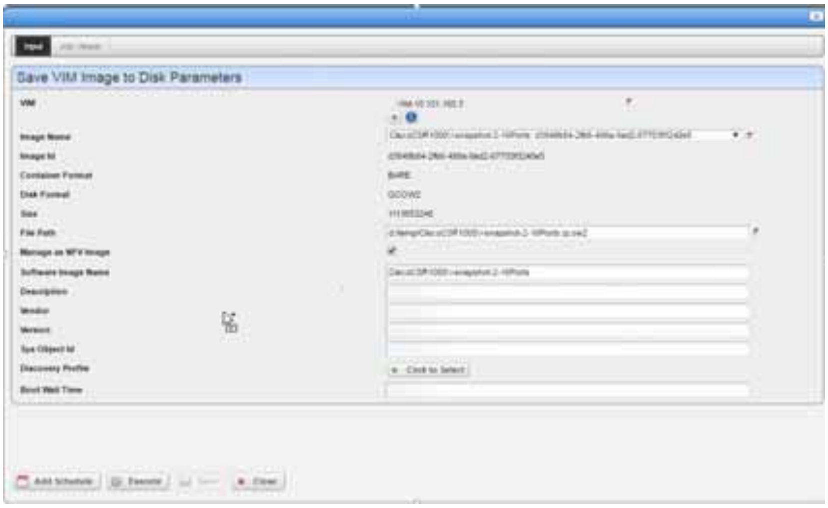

Save VIM Image to Disk Parameters

VIM

Image Name

Image Id

Container Format: BARE

Disk Format: QCOW2

Size

File Path

Manage as NFV Image: ✓

Software Image Name

Description

Vendor

Version

Sys Object Id

Discovery Profile: Click to Select

Boot Wait Time

6 Click Execute.

7 Copy the image downloaded from your local disk to your company's centralized file server (such as an FTP Server) to redistribute the new snapshot image across additional OpenStack VIMs.

Before the image descriptor is complete and ready to export, you need to modify the image descriptor.

Modifying the VIM Software Image Descriptor

These instructions assume that you have already created a VIM image descriptor. If you have not created the image descriptor, see Creating a Software Image Descriptor for a Snapshot on page 747.

Modify the VIM software image descriptor from the Cruz application as follows.

1 Navigate to the Software Images portlet.

The portlet's location depends on your company's configuration.

2 Right-click an image.

3 Select Manage.

The Manage new Software Image from VIM Parameters window is displayed.

4 Select the VIM that houses the snapshot image.

5 Select the snapshot image name.

For example: `CiscoCSR1000v-snapshot-2-16Ports:` *`uniqueImageID`*

The image ID, container format, disk format, size, and software image name fields are populated.

6 Select a discovery profile.

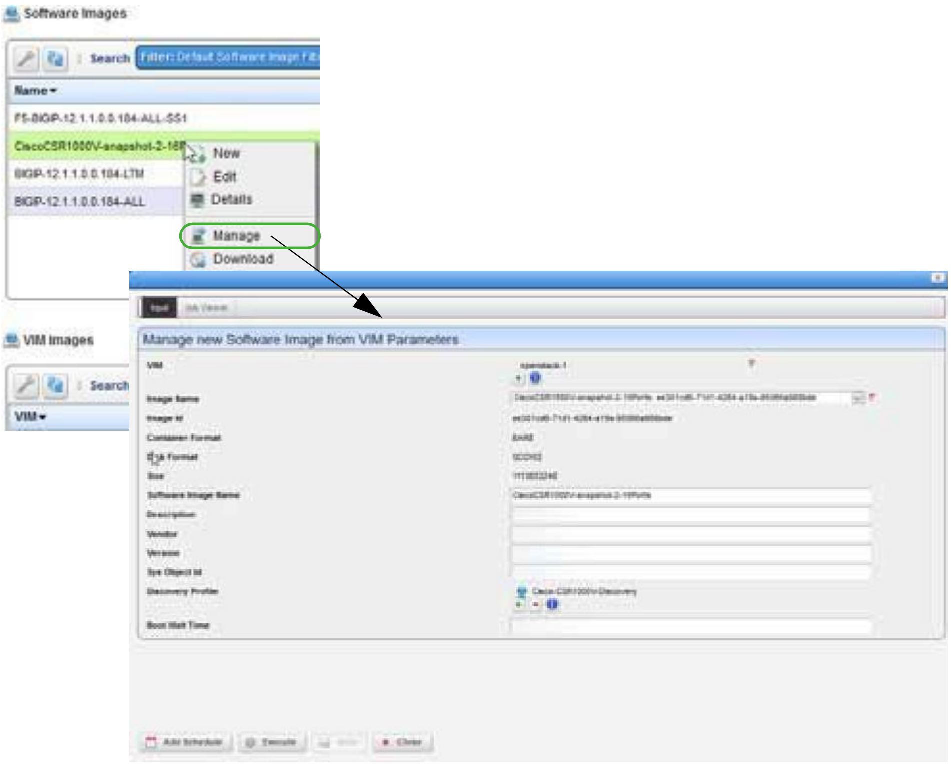

Manage new Software Image from VIM Parameters

VIM: openstack-1
Image Name: CiscoCSR1000V-snapshot-2-16Ports
Image Id
Container Format: BARE
Disk Format: QCOW2
Size
Software Image Name: CiscoCSR1000V-snapshot-2-16Ports
Description
Vendor
Version
Sys Object Id
Discovery Profile: Cisco-CSR1000V-Discovery
Boot Wait Time

7 Click Execute.

A newsoftware image record is created matching the image name in your OpenStack environment.

8 Right-click an image.

9 Select Edit.

The Editing Software Image window is displayed.

10 Select the snapshot image from the Image Name list.

The Editing Software Image window is displayed.

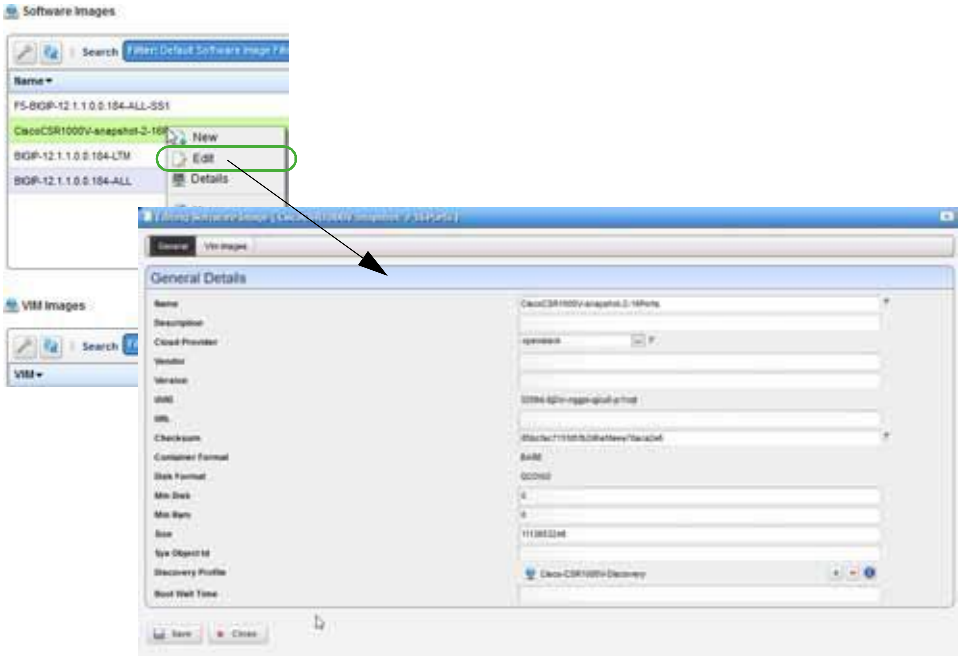

Name	CiscoCSR1000V-snapshot.2.16Ports
Description	
Cloud Provider	openstack
Vendor	
Version	
UUID	[illegible]
URL	
Checksum	[illegible]
Container Format	BARE
Disk Format	QCOW2
Min Disk	0
Min Ram	0
Size	[illegible]
Sys Object Id	
Discovery Profile	Cisco-CSR1000V-Discovery
Boot Wait Time	

11 Enter the missing information as follows.

a. Add a meaningful description.

b. Specify the vendor, such as Cisco.

c. Add a meaningful version number, such as 2.

d. Enter avalid URLthat points to yourlocalenvironment’s file serveror disklocation where the qcow2 file is stored and accessible. Forexample:

ftp://*caserverIP*/cisco/csr/CiscoCSR1000v-snapshot-2-16Ports.qcow2

e. Enter the SNMP SysObject ID for the Cisco CSR 1000v model, suchas:

`1.3.6.1.4.1.9.1.1537`

f. Select the discovery profile containing the required ssh and SNMP authentication objects.

Discovery profiles are created by your system administrator or someone within your organization with administrative permissions.

g. Leave the Boot Wait Time field blank.

The Cruz assigns a default integer value of 0 after you save the profile.

h. Reviewyourinputs to ensureaccuracy.

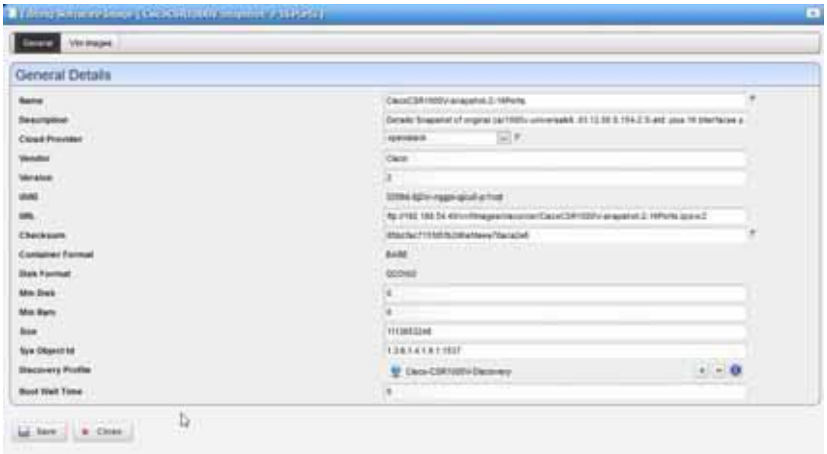

General Details

Name	CiscoCSR1000V-snapshot.2-16Ports
Description	[illegible]
Cloud Provider	openstack
Vendor	Cisco
Version	2
UUID	[illegible]
URL	[illegible]
Checksum	[illegible]
Container Format	BARE
Disk Format	QCOW2
Min Disk	0
Min Ram	0
Size	[illegible]
Sys Object Id	1.3.6.1.4.1.9.1.1537
Discovery Profile	Cisco-CSR1000V-Discovery
Boot Wait Time	[illegible]

12 Save the software image record.

13 Locate and Select the newly created Software Image Record and export it to an archive disk location for backup purposes.

Setting Up F5 BIGIP VNF Package

Thissectionexplains thescope, functions, and processesassociated with Virtualized Network Functions (VNF) onboarding and instantiation:

- Understanding the F5 BIGIP VNF Package
- Preparing an Image
- Manually Applying an F5 License

Understanding the F5 BIGIP VNF Package

The F5 BIG -IP Virtual Edition (VE) is a virtual application delivery controller that is deployable in an OpenStack environment. Th e BIG-IP VE provides rich features, such as advanced traffic management, acceleration, DNS, firewall, and access management. The BIG-IP VE software images are downloadable and portable between on -premises virtualized datacenters, public, and hybrid cloud environments. With BIG-IP virtual editions, you can rapidly provision consistent application services across the data center and into the cloud.

The application F5 BIGIP Virtualized Network Function (VNF) package provides sample VNF descriptors and associated artifacts used to deploy the L TM (Local Traffic Manager, Load Balancer) and AFM (Advanced Firewall Manager) models.

F5 Networks provides qcow2 images (.qcow2) and licenses that are ready for the OpenStack environment. You need these files to deploy and enable the F5 BIGIP VNF features.

The following list provides the VNF package vendor and validation in formation:

Vendor and Model Information	
Vendor	F5 Networks, Inc.
Product Family	BIGIP VE
Models	Type 1: LTM (Load Balancer) Type 2: AFM (Firewall)

Operating Systems	BIGIP Traffic Management Shell (tmsh)
Version	12.1.1
Validated Environment	
VIM	OpenStack
Version	Liberty
Installer	Mirantis
Version	8
Network Configurations	VXLAN, VLAN Segmentation
Cruz **MANO (Management and Operation) Validated Functions**	
MANO	NSD, VNFD, NSR, VNFR Link LCEs
VNFM	Generic (MANO Lite) Independent/Specific (N-N/A)
NFVO	MANO Lite
Cruz **RMO (Resource Management and Operation) Validated Functions**	
Inventory	Resource Discovery and Inventory Resource Creation Method (NFV-Model)
Monitoring	SNMP Monitoring CPU & Memory KPIs TFA (flows)
Config	ACLI
NetRestore File Management	NR Backup NR Restore NR Firmware (not validated)

This section covers the:

- F5 BIGIP VNF Package Contents
- Cruz VNF Package - Install Component Files
- License Requirements
- Known Issues
- References

Once you have a good understanding of the VNF package, you can get started preparing an image.

F5 BIGIP VNF Package Contents

The F5 BIGIP VNFpackagecontainsartifacts necessary to install and deploy Virtualized Network Functions (VNFs) for the following Virtualization Deployment Unit (VDU) types the F5 BIG-IP VE product family offers:

- Type 1: BIGIP VE - LTM (Load Balancer)
- Type 2: BIGIP VE - AFM (Firewall)

Multiple sample VNF descriptors were created with basic flavors to deploy the VNF VDUs, either as part of a standalone VNF record, or as part of a higher level, end-to-end Network Service record:

- NFV VNF Descriptors
- Network Service Descriptor

NFV VNF Descriptors

Thissectiondescribesthefollowing Cruz NetworkFunctions Virtualization(NFV) VNF (virtualized network functions) descriptors:

- F5 BIGIP LTM Virtual Appliance
- F5 BIGIP ALL AFM Virtual Appliance
- F5 BIGIP Cluster.

F5 BIGIP LTM Virtual Appliance

The sample F5 BIGIP LTM Virtual Appliance de scriptor is located in the following file:

`owareapps/nfv-vnfm-f5/db/vnfd.f5-bigip-ltm-v12.xml`

Thissample Virtualized Network Function (VNF) descriptor defines the packagingand behaviorof a single F5 BIGIP LTM (Load Balancer) Virtualization Deployment Unit (VDU) and can be included as part of a network descriptor.

Here are the **F5 BIGIP LTM Virtual** descriptor deploy parameters used:

Field Name	Data Type	Description	Default Value	Valid Values	Required
LTM Admin Password	String	VNF’s Admin User Password	admin	Valid Linux password format	N

Here are the **F5 BIGIP LTM Virtual** descriptor flavors used:

Flavor Name	VDU Name	VDU Count	CPU	Disk (GB)	Memory (MB)
F5-BIGIP-TTM-Small	BIGIP LTM VA	1	2	50	4096

View the VNF descriptor’s composition details from the Cruz VNF Descriptor portlet.

F5 BIGIP ALL AFM Virtual Appliance. The sample F5 BIGIP ALL AFM Virtual Appliance descriptor is located in the following file:

`owareapps/nfv-vnfm-f5/db/vnfd.f5-bigip-all-afm-v12.xml`

Thissample Virtualized Network Function (VNF) descriptor defines packagingand behavior ofa single F5 BIGIP AFM (Firewall) Virtualization Deployment Unit (VDU) and can be included as part of a Network descriptor.

Here are the **F5 BIGIP ALL AFM Virtual** descriptor deploy parameters:

Field Name	Data Type	Description	Default	Valid Values	Required
AFM Admin Password	String	VNF’s Admin User Password	admin	Valid Linux password format	N

Here are the **F5 BIGIP ALL AFM Virtual** descriptor flavors:

Flavor Name	VDU Name	VDU Count	CPU	Disk (GB)	Memory (MB)
F5-BIGIP-AFM-Small	BIGIP AFM VA	1	4	160	8192

View the VNF descriptor’s composition details from the Cruz VNF Descriptor portlet.

F5 BIGIP Cluster. . The sample F5 BIGIP Cluster descriptor is located in the following file:

`owareapps/nfv-vnfm-f5/db/vnfd.f5-bigip-cluster-v12.xml`

Thissample Virtualized Network Function (VNF) descriptor definesastandalone type, which means that no higher-level network service is required to establish the network/virtual linksand deploy stage/deploy the VNF.

This sample VNF descriptor includes lifecycle event (LCE) tasks to create the project's router and networks as well as deploy into the project' s network a varied composition of F5 BIGIP Virtualization Deployment Units (VDUs), based upon theassigned VNF flavor.

Here are the **F5 BIGIP Cluster** descriptor deploy parameters used:

Field Name	Data Type	Description	Default Value	Valid Values
Cluster Name	String	Name of the Cluster identifying the group of networks and interconnected VDUs	F5-BIGIP-Cluster	String (1-255 characters)
Shared Key	String	SSH Shared Key	OpenStack-All	Name of a valid/existing OpenStack Key Pair Object usable by the selected VIM and project.

Here are the **F5 BIGIP Cluster** descriptor flavors used for the VNF scope. There are no VDU scope deploy parameters:

Flavor Name	VDU Name	VDU Count	CPU	Disk (GB)	Memory (MB)
F5-BIGIP-LTM-Only	F5 BIGIP LTMBIGIP-C-LTM VDU	1	2	50	4096
F5-BIGIP-Small	BIGIP-C-LTM F5 BIGIP LTM VDU	1	2	50	4096
	BIGIP-C-AFM F5 BIGIP ALL (AFM) VDU	1	4	160	8192

View the VNF descriptor's composition details from the Cruz VNF Descriptor portlet.

Network Service Descriptor

Thesample Network Service descriptor (NSD) provided shows how to define a networkserviceand utilize any combination of the Virtualized Network Function (VNF) descriptors provided in this package (those that are not standalone), to compose various flavors with different Virtualization Deployment Unit (VDU) composition for a single network service definition.

View the Network Service descriptor contents from the Cruz Network Service Descriptors portlet.

NOTE:

The Network Service Descriptors portlet location is dependent on your company's Cruz NFV installation and configuration. By default, the Network Service Descriptors portlet is located on the On-Boarding page and the Network Services page.

This sample NSD is located in the following file:

`owareapps/nfv-vnfm-f5/db/nsd.f5-bigip- v12-sample.xml`

Here are the **VNF Scope** deploy parameters used:

Field Name	Data Type	Description	Default Value	Valid Values	Required
Cluster Name	String	Name of the Cluster identifying the group of networks and interconnected VDUs	F5-BIGIP-Cluster	String (1-255 characters)	Y

Here are the **BIGIP LTM VA VDU Scope** deploy parameters used:

Field Name	Data Type	Description	Default Value	Valid Values	Required
LTM Admin Password	String	VNF’s Admin User Password	admin	Valid Linux password format	N

Here are the **BIGIP AFM VA VDU Scope** deploy parameters used:

Field Name	Data Type	Description	Default Value	Valid Values	Required
AFM Admin Password	String	VNF’s Admin User Password	admin	Valid Linux password format	N

Here are the **Network Service** descriptor flavors used:

NS Flavor Name	VNF Flavor Names	VDU Name	VDU Count	CPU	Disk (GB)	Memory (MB)
Standard-Small	F5 BIGIP LTM Small	F5 BIGIP LTM Virtual Appliance 1 x BIGIP LTM VA VDU	1	2	50	4096
Standard-Medium	F5 BIGIP LTM VA Small	BIGIP LTM VA VDU 1 x BIGIP LTM VDU	1	2	50	4096
	F5 BIGIP AFM VA Small	F5 BIGIP LTM Virtual Appliance 1 x BIGIP AFM VA VDU	1	2	160	8192

Cruz VNF Package - Install Component Files

The following Cruz Virtualized Network Function (VNF) package installation components are required for the F5 BIGIP product:

VNF (OCP)	nfv-vnfm-f5.ocp
Device Driver (DDP)	f5bigip.ddp

The Cruz VNF package descriptors for the **F5 BIGIP** VNF product include:

NVF VNF Package			Resource Management		
VIM Software Image Records	VNF Descriptors	Sample NSD	Discovery Authentication Records	Discovery Profiles	Device Driver
Y	Y	Y	Y	Y	Y

The VNF package descriptors include:

- Discovery Authentication Profiles
- Discovery Profiles
- NFV VIM Software Image Descriptors

Discovery Authentication Profiles

The F5 BIGIP Discovery Authentication pro files are located in the following file:

`owareapps/nfv-vnfm-f5/db/rc.disc-auths-f5-bigip.xml`

This file contains the following discovery authentication profiles:

Authentication Record Name	Description
f5bigip-cli	SSH v2 Credentials
f5bigip-snmp	SNMP v1/v2 Credentials
f5bitip-https	HTTPS Credentials

Discovery Profiles

The F5 BIGIP Discovery profiles are located in the following file:

`Fowareapps/nfv-vnfm-f5/db/rcdisc-profile-f5-bigip.xml`

Thisfilecontains the F5 BIGIP Discovery profile that automatically discovers and populates resourcerecordsduring Virtualized Network Functionality(VNF) recordinstantiation. The discovery profile name is F5-BIGIP-VE-Discovery_v12x.

The following tasks are run once the target VM in the OpenStack environment is fully booted and can respond to SNMP and CLI requests from the management system:

1 Resync

2 DataCollection ForGroupOfDevices

3 Scheduled R esync

4 Refresh Monitor Targets

5 F5 BIGIP SNMP Trap Forwarding to Cruz

NFV VIM Software Image Descriptors

The F5-BIGIP-12.1.1.0.0.184-LTM-SS1 and F5-BIGIP-12.1.1.0.0.184-ALL-SS1 snapshot records were used to create and test the Network Functions Virtualization (NFV) descriptors in this package. See Preparing an Image on page 759 for image customization required to produce and use these snapshots.

Theoriginal F5 vendorimagescovered bythe BIGIP-12.1.1.0.0.184-LTMand BIGIP-12.1.1.0.0.184-ALL descriptors were used to produce the workable snapshot images.

Name	Scope/Description	Container Format	Disk Format	Min RAM (MB)	MIN Disk (GB)	Resource Discovery Profile
F5-BIGIP-12.1.1.0.0.184-LTM-SS1	F5 BIGIP LTM - snapshot created from F5 original image BIGIP-12.1.1.0.0.184-LTM	BARE	QCOW2	4096	50	Y

Name	Scope/Description	Container Format	Disk Format	Min RAM (MB)	MIN Disk (GB)	Resource Discovery Profile
F5-BIGIP-12.1.1.0.0.184-ALL-SS1	F5 Big IP Firewall - snapshot created from F5 original image BIGIP-12.1.1.0.0.184-ALL	BARE	QCOW2	8192	160	Y
BIGIP-12.1.1.0.0.184-LTM	F5 BIGIP VE LTM (Load Balancer) Image - Original **Note:** No RC Discovery/ Management Enabled	BARE	QCOW2	4096	50	Y
BIGIP-12.1.1.0.0.184-ALL	F5 BIGIP VE ALL Image (Use for AFM-Firewall) F5 BIGIP VE LTM (Load Balancer) Image - Original **Note:** No RC Discovery/ Management Enabled	BARE	QCOW2	8192	160	Y

License Requirements

Toenable either the LTM(Load Balancer) or AFM (Firewall) feature set, you must manuallyapply a license to each F5 BIGIP VM.

Manuallyapply the license after creating the Virtualized Network Function (VNF) record and successfully deploying the Virtualization Deployment Units (VDUs) to the OpenStack environment.

Youcanregister and applyalicense directly to an F5 BIG-IP VE VM whosestatus is Online (Active) by logging into the VNF’s Web interface and following the instructions F5 provides.

See Manually Applying an F5 License on page 774 regarding an example process for obtaining/ applying the license provided by F5 using the manual license registration process. Contact F5 directly for further assistance.

Known Issues

The following issues are known:

- F5 BIGIP VE license registration and activation requires manual steps.

  License activation is not automatically assigned or applied by Cruz NFV -MANO. License activation requires some manual steps. After instantiating a VNF/VDU, you must perform the manual steps every time you instantiate a VM that represents an F5 BIGIP VE instance in the OpenStack environment.

  See Manually Applying an F5 License on page 774 for detailed instructions.

- Management Traffic using SNMP is blocked by default.

  By default, the F5 BIGIP VE is configured to block SNMP traffic.

  Cruz Resource Discovery and Resource Monitoring activities require that SNMP access is turned on and enabled on the running F5 BIGIP VEs, in order to discover and manage the VEs.

  Cruz applications **require** a setting before creating the bootable snapshot that enables SNMP access.

See Preparing an Image on page 759 for details on creating an updated F5 BIGIP VE snapshot image.

- Pagination preference for F5's TMSH CLI is only availableglobally.
  The F5 BIGIP TMSH command line interface (CLI) Pagination preference setting is not available on a per CLI session or per user account basis. Because the setting is global and it is required to be set to false for Cruz to issue scripts to the F5 BIGIP VM via the TMSH CLI without errors in an automated fashion, the global pagination setting must be reset to `false` (off).
  Turn off the pagination setting before creating the bootable snapshot that contains other settings.
  See Preparing an Image on page 759 for details on creating the updated F 5 BIGIP VE snapshot image.
- The nfv-vnfm-f5 package is a codeless VNFM implementation and therefore **does not support** NFV VNF record and NFV Network Service record Discovery/Resync.

References

This section provides references to:

- Vendor Documentation (URLs)
- VNF Configuration Methods

Vendor Documentation (URLs)

Referto thefollowingvendor websitesforassistance, downloads, documentation, product datasheet, and so on:

- F5 Vendor website:
  `http://www.f5.com`
  `http://ask.f5.com`
  `http://devcentral.f5.com`
- Product Description:
  `https://www.f5.com/pdf/products/BIGIP-virtual-editions-datasheet.pdf`
- OpenStack KVM Environment Requirements:
  `https://support.f5.com/kb/en-us/products/BIGIP_ltm/manuals/product/bigip-ve-setup-linux-kvm-12-0-0.html`

VNF Configuration Methods

A combination of image (snapshot) preparation, along with Virtualization Deployment Unit (VDU) Post Instantiate Configuration Lifecycle Events (LCEs), is implemented in this VNF package to ensure that the F5 BIGIP LTMand AFM VNF VDUs are automatically brought under Cruz management.

Customized Image Preparation. The base BIGIP modelfamilyimagesprovided by F5 requiresadditional manual configuration before producing VMs in the OpenStack environment that can be discovered and managed using the Cruz Resource Management applications.

See Preparing an Image on page 759 for a description of how to create an updated image (snapshot) for the F5 BIGIP LTM and AFM model series to use to boot up VMs that the application manages. Use the snapshot produced to create F5 BIGIP VEs in the OpenStack environment, so that no manual intervention is required to apply the management system communication settings.

Cruz Resource Configuration Tasks. Once the target VM in the OpenStack environment is fully booted and can respond to SNMP and CLI requests from the management system, post discovery tasks are run.

For more details about the F5-BIGIP-VE-Discover_v12x profile and the tasks that run, see Discovery Profiles on page 756.

Preparing an Image

Beforeyoucandiscoverand manage a Virtualized Network Function (VNF) resource from the application, you need to:

1 Set up your OpenStack environment bypreparing BIGIP VE VIM images with the necessary configuration.

2 Create a software image record from the Cruz application.

A snapshot is ta ken to capture the following settings, so that the Cruz system automatically discovers the resources using SNMP and SSH (tmsh CLI) access when a new BIGIP VE VM is created/booted in the OpenStack environment.

To enable Cruz Resource Management for the BIGIP VE VNF VDUs (VMs), use these management se ttings:

- SNMP Access Enabled
- Pagination Turned Off
- Mgmt (eth0) Port MTU - Set to recommended value of 1400 instead of 1500 (default)

Setting Up Your OpenStack Environment

To set up your OpenStack environment, you need to:

1 Create aproject.

2 Deploy the vendor’s base image by uploading the original qcow2 image file received from the F5 vendor to a target Bare Metal OpenStack controller.

3 Make configuration changes to the running VM booted from the F5 vendor’s original image.

4 Create flavors representing the smallest supported flavor to use. The F5 vendor provides the following flavors (Figure 16-4).

F5 LTM v12.1.1 For LTM Models	Flavor Name: f5.medium	CPU – 2, Memory – 4 GB (4096 MB), Disk - 50 GB
F5 BIGIP ALL v12.1.1 For AFM (Firewall) Models	Flavor Name: f5.xxlarge	CPU – 4, Memory – 8 GB (8192 MB), Disk – 160 GB

5 Create VM snapshots that are used to create and export a VIM software image descriptor within the application.

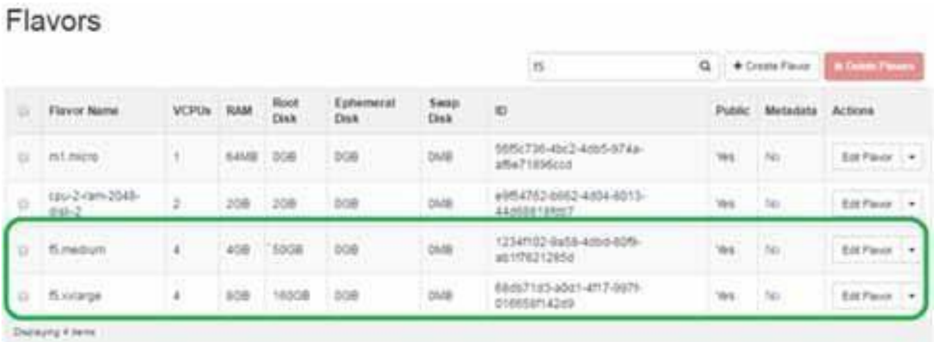

Flavors

f5

□	Flavor Name	VCPUs	RAM	Root Disk	Ephemeral Disk	Swap Disk	ID	Public	Metadata	Actions
□	m1.micro	1	64MB	0GB	0GB	0MB	56f5c736-4bc2-4db5-974a-af6e71896ccd	Yes	No	Edit Flavor
□	cpu-2-ram-2048-disk-2	2	2GB	2GB	0GB	0MB	e9f54762-b662-4d04-8013-44d688189bf7	Yes	No	Edit Flavor
□	f5.medium	4	4GB	50GB	0GB	0MB	1234f102-9a58-4dbd-80f9-ab1f76212b5d	Yes	No	Edit Flavor
□	f5.xxlarge	4	8GB	160GB	0GB	0MB	68d571d3-a0d1-4f17-997f-0166581142d9	Yes	No	Edit Flavor

Displaying 4 items

Figure 16-4. F5 BIGIP Flavors Example

Creating a Project

Create a project in an OpenStack VIM (Virtualized Infrastructure Manager) as follows.

1 Select Identity > Project.

2 Click Create Project.

The Create Project window is displayed.

3 Enter a name and detailed descriptions.

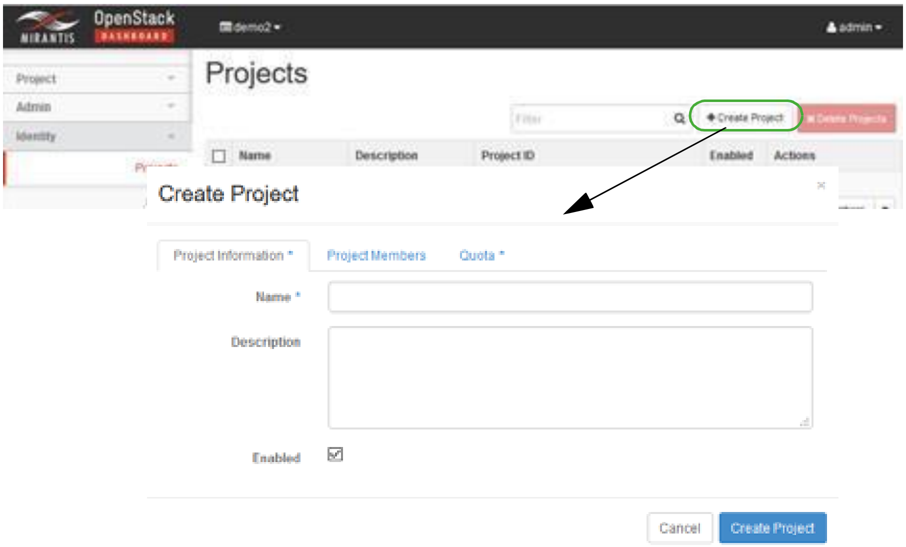

Create Project

Name *

Description

Enabled ☑

4 Assign users/permissions.

a. Click Project Members.

b. Add the following users to the Project Members list.

- heat-cfn
- heat
- admin

c. Set the member permissions to admin.

5 Click Create Project.

The new project is saved to the Projects list.

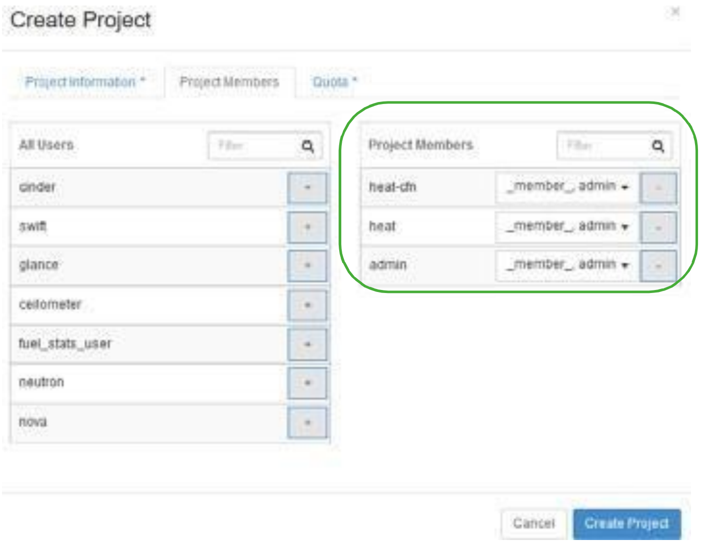

Create Project

Project Information * | Project Members | Quota *

All Users	
cinder	+
swift	+
glance	+
ceilometer	+
fuel_stats_user	+
neutron	+
nova	+

Project Members		
heat-cfn	_member_, admin	-
heat	_member_, admin	-
admin	_member_, admin	-

6 Create project networking in an OpenStack VIM (KILO or later release).

a. Create a project router with a gatewayinterface.

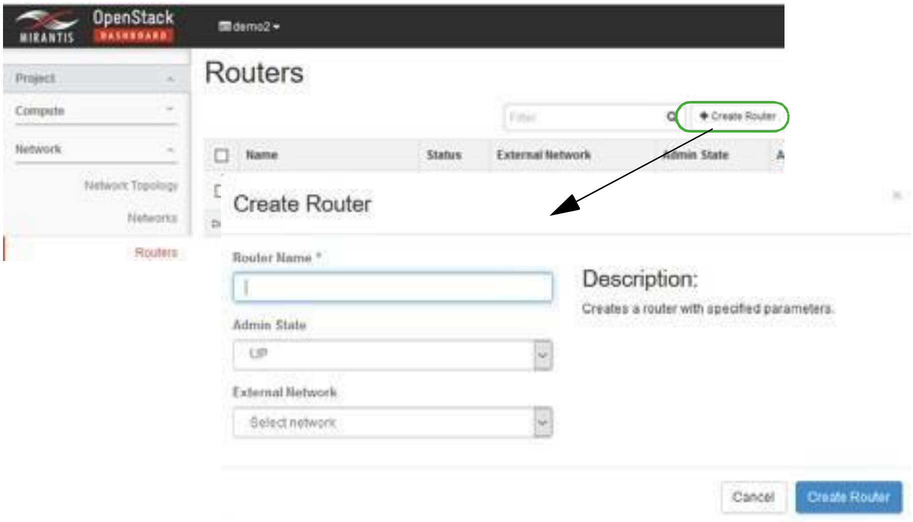

Routers

Create Router

Router Name *

Admin State: UP

External Network: Select network

Description:

Creates a router with specified parameters.

b. Make sure that the router's gateway IP is reachable from the LAN (your PC) following project router creation.

c. Create and attach the following networks to the project router:

- net-mgmt
- net-data1
- net-data2

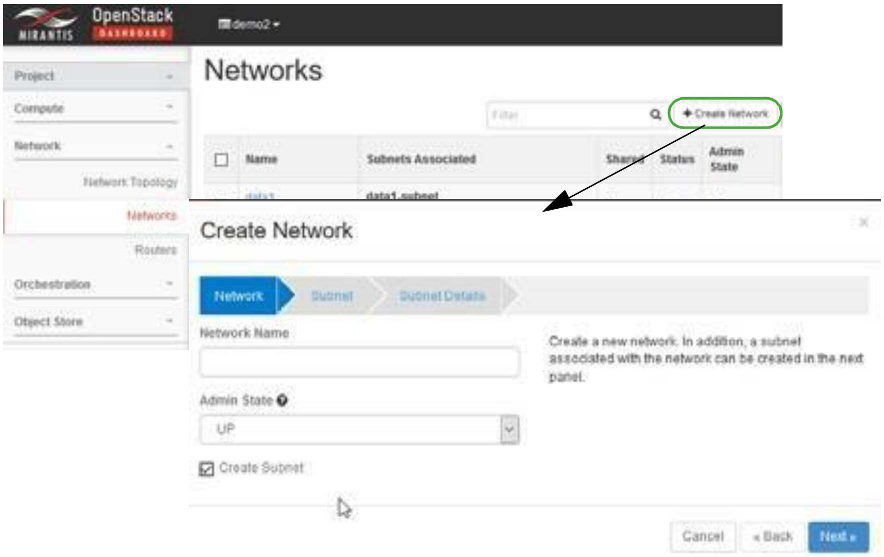

Admin State: UP

☑ Create Subnet

Create a new network. In addition, a subnet associated with the network can be created in the next panel.

7 Create the F5 BIGIP security group in the project.

The new security group allows utilization of the F5 BIGIP VE VMs created during Image Preparation and other manual activities for the F5 BIGIP VE performed in the OpenStack environment.

a. Navigate to the project, for example the F5Dev project.

b. Select Project > Compute > Access & Security.

The Access & Security options are displayed.

c. Click Create Security Group.

The Create Security Group window is displayed.

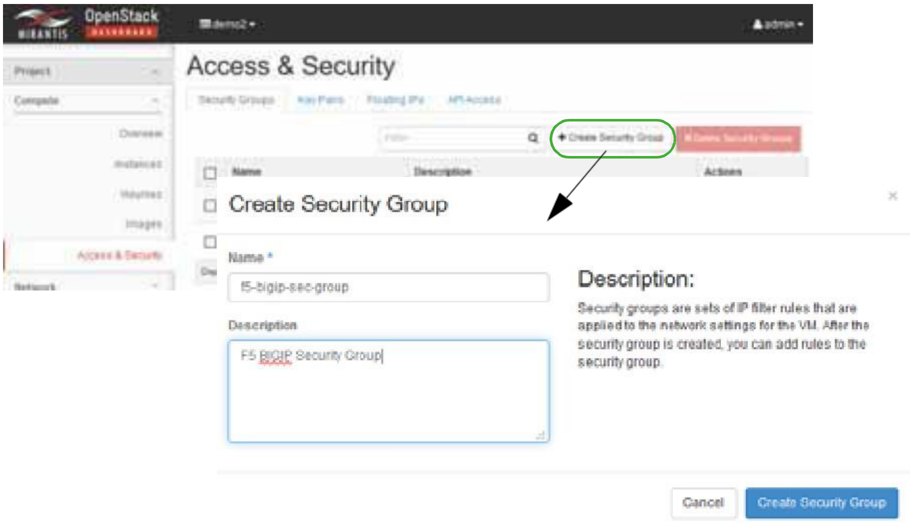

Name *: f5-bigip-sec-group

Description: F5 BIGIP Security Group

Description:

Security groups are sets of IP filter rules that are applied to the network settings for the VM. After the security group is created, you can add rules to the security group.

d. Enter the security group name and description.

e. Click Create Security Group.

The new security group object is saved in the OpenStack environment.

f. Select the Manage Rules action for the new security group.

The two default rules added by the OpenStack environment are listed.

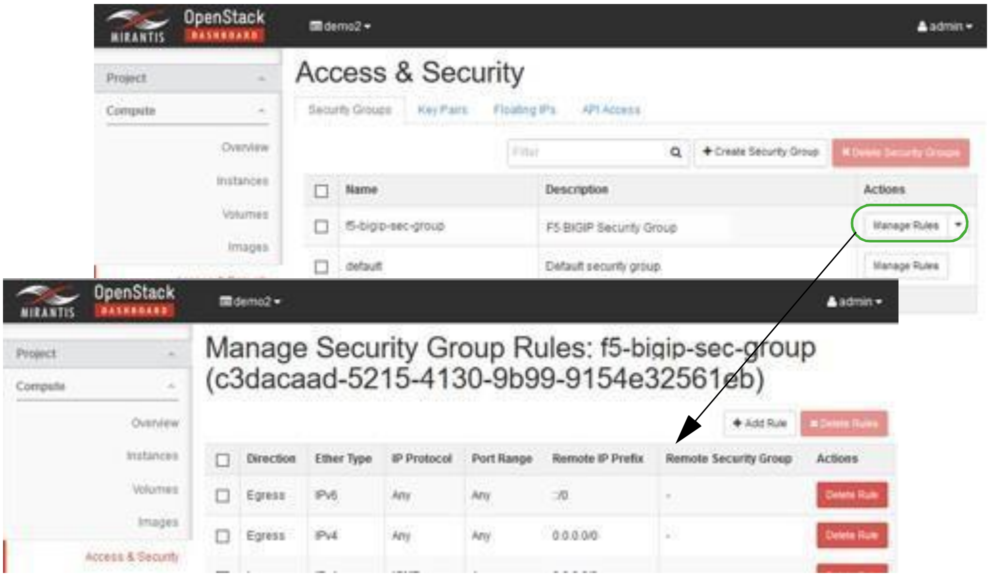

Access & Security

□	Name	Description	Actions
□	f5-bigip-sec-group	F5 BIGIP Security Group	Manage Rules
□	default	Default security group	Manage Rules

Manage Security Group Rules: f5-bigip-sec-group (c3dacaad-5215-4130-9b99-9154e32561eb)

□	Direction	Ether Type	IP Protocol	Port Range	Remote IP Prefix	Remote Security Group	Actions
□	Egress	IPv6	Any	Any	::/0	-	Delete Rule
□	Egress	IPv4	Any	Any	0.0.0.0/0	-	Delete Rule

g. Click Add Rule.

The Add Rule window is displayed.

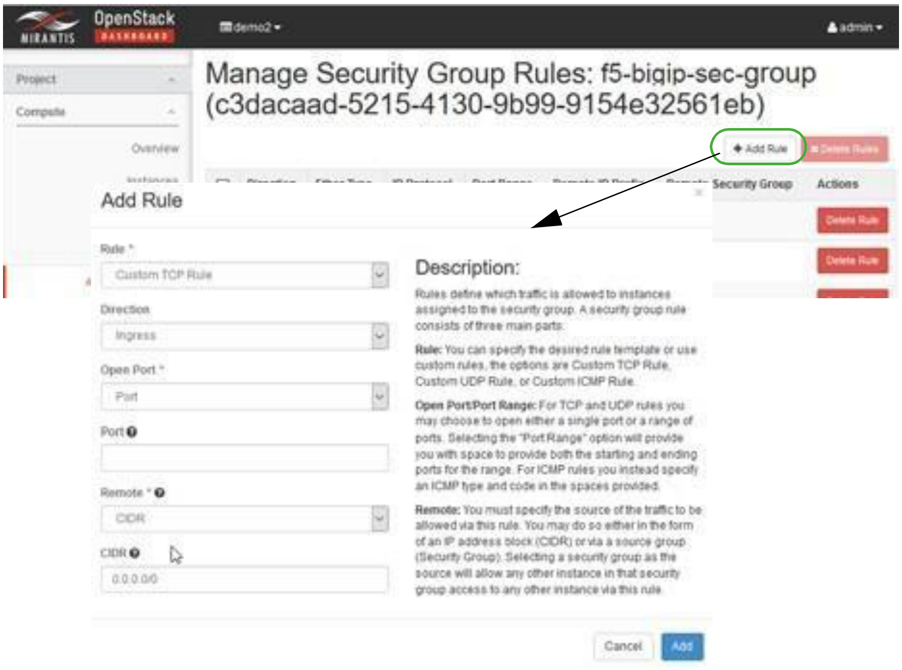

Manage Security Group Rules: f5-bigip-sec-group (c3dacaad-5215-4130-9b99-9154e32561eb)

Add Rule

Rule *: Custom TCP Rule

Direction: Ingress

Open Port *: Port

Port:

Remote *: CIDR

CIDR: 0.0.0.0/0

Description:

Rules define which traffic is allowed to instances assigned to the security group. A security group rule consists of three main parts:

Rule: You can specify the desired rule template or use custom rules, the options are Custom TCP Rule, Custom UDP Rule, or Custom ICMP Rule.

Open Port/Port Range: For TCP and UDP rules you may choose to open either a single port or a range of ports. Selecting the "Port Range" option will provide you with space to provide both the starting and ending ports for the range. For ICMP rules you instead specify an ICMP type and code in the spaces provided.

Remote: You must specify the source of the traffic to be allowed via this rule. You may do so either in the form of an IP address block (CIDR) or via a source group (Security Group). Selecting a security group as the source will allow any other instance in that security group access to any other instance via this rule.

h. Add the following rules to enable all ICMP, TCP, and UDP traffic:

Direction	Either Type	IP Protocol	Port Range	Remote IP Prefix	Remote Security
Ingress	IPv4	ICMP	Any	0.0.0.0/0	-
Egress	IPv4	ICMP	Any	0.0.0.0/0	-
Ingress	IPv4	TCP	1-65535	0.0.0.0/0	-
Egress	IPv4	TCP	1-65535	0.0.0.0/0	-
Ingress	IPv4	UDP	1-65535	0.0.0.0/0	-
Egress	IPv4	UDP	1-65535	0.0.0.0/0	-

When done, the resulting rule set for the f5-bigip-sec-group security group should contain the two default rules added by th e OpenStack environment and the traffic rules you added, as in:

Direction	Either Type	IP Protocol	Port Range	Remote IP Prefix	Remote Security
Egress	IPv4	Any	Any	0.0.0.0/0	-
Egress	IPv6	Any	Any	::/0	-
Ingress	IPv4	ICMP	Any	0.0.0.0/0	-
Egress	IPv4	ICMP	Any	0.0.0.0/0	-
Ingress	IPv4	TCP	1-65535	0.0.0.0/0	-
Egress	IPv4	TCP	1-65535	0.0.0.0/0	-
Ingress	IPv4	UDP	1-65535	0.0.0.0/0	-
Egress	IPv4	UDP	1-65535	0.0.0.0/0	-

Now it is time to instantiate a VM from a vendor image.

Instantiating a VM from a Vendor Image

Instantiate a VM (virtual machine) in the new OpenStack project from the F5 vendor-provided image as follows.

1 Boot up the F5 BIGIP LTM or F5 BIGIP ALL (AFM) image with the following network connections.

Linux interface name	F5 Interface Name	Network
eth0	mgmt	net-mgmt
< None > *	1.1	net-data-1
< None > *	1.2	net-data-2
* The < None> interface names require CLI configuration along with a proper license before the interface can be configured/enabled and pass traffic. This is done as post-bootup and post device discovery configuration applied by the Cruz system (not part of the base image snapshot).		

2 Assign a floating IP Address to the eth0/mgmt connection.

3 Log into the new VMs Web interface using the floating IP address.

For example: `https://10.122.0.244`

4 Verify that the status changes to ONLINE (ACTIVE) and the proceeding status bar is green.

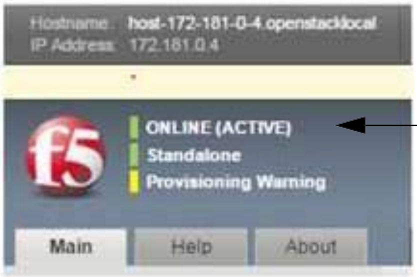

Hostname: host-172-181-0-4.openstacklocal
IP Address: 172.181.0.4

ONLINE (ACTIVE)
Standalone
Provisioning Warning

Status Information

Now you are ready to make VM configurations changes.

Manually Making Configuration Changes

Toenable automated resourcemanagement, the Cruzapplicationrequiresadditionalconfiguration to the VM post bootup and this configuration captured in the new snapshot.

Manually make the required configuration changes to the running VM booted from the F5 vendor’s original image as follows.

NOTE:

DoradoSoftwareprovidesanexecutablescriptintheF5BIGIPDeviceDrivercomponentthatautomates the following configurationadditions.

1 Go to a command line prompt.

On a Windows system, enter Cygwin Linux shell.

2 Run the following bash script to apply the management configuration.

```
$OWARE_USER_ROOT/oware/bin/oware.cmd
```

3 Run the f5vmconfig.sh script, which logsinto the target VM using the F5 vendor’s default ssh v2 user credentials and applies the following commands to the VM as well as saves the running configuration to be permanent.

`f5vmconfig.sh -i <`*`Management IP Address`*`> [-m <`*`MTU`*`>]`

Where:

- **-i** is the target F5 BIGIP device management IP address (required)
- **-m** is the MTU integer value (optional)

  The Dorado Software LAN environment requires that MTU value of 1400 be assigned to the mgmt (eth0) port. Omit or change to a different valid MTU value based upon your specific network environment requirements.
- **-h** is the help message

For example: `f5vmconfig.sh -i 10.122.15.44 -m 1400`

4 Reboot the VM for the MTU change to take affect by logging into the Linux shell using the `ssh` command and issuing the `reboot` command.

The ssh connectivity is lost, until the VM reboots and completely starts up again.

5 Manually log into the VM using the `ssh` command.

6 Verify that the MTU setting was applied or retainedproperly:

```
# ifconfig eth0
config eth0 Link encap:Ethernet HWaddr FA:16:3E:9E:E7:4D
   inet addr:172.85.0.5 Bcast:172.85.0.255 Mask:255.255.255.0
```

```
    inet6 addr: fe80::f816:3eff:fe9e:e74d/64 Scope:Link
    UP BROADCAST RUNNING ALLMULTI MULTICAST MTU:1400 Metric:1
    RX packets:817 errors:0 dropped:0 overruns:0 frame:0
    TX packets:843 errors:0 dropped:0 overruns:0 carrier:0
    collisions:0 txqueuelen:1000
    RX bytes:167119 (163.2 KiB) TX bytes:148788 (145.3 KiB)
```

7. Verify that the **mtu** setting of **1400** for the **eth0** interface command was appended to the /config/startup file by running the following command.

```
config # cat /config/startup
#
# NOTE:
# This file will be installed in /config/startup and it will
# be called by /etc/rc.local.
#
# - /config/startup is for customer config additions and
# will be saved in UCS and synced by tmsh run sys sync-sys-files
#
# - /etc/rc.local should *not* be used by customers and
# can/will be changed by F5
#
ifconfig eth0 mtu 1400
```

8. Make sure that the SNMP allowed addresses tmsh cli setting of **ALL** was persisted after reboot:

```
config # tmsh list sys snmp allowed-addresses
sys snmp {
    allowed-addresses { ALL }
}
```

Create flavors before creating the VM snapshots.

Creating Flavors

Create the following F5 flavors from the OpenStack dashboard.

F5 LTM v12.1.1 For LTM Models	Flavor Name: f5.medium	CPU – 2, Memory – 4 GB (4096 MB), Disk - 50 GB
F5 BIGIP ALL v12.1.1 For AFM (Firewall) Models	Flavor Name: f5.xxlarge	CPU – 4, Memory – 8 GB (8192 MB), Disk – 160 GB

1. Select Admin > System > Flavors.
2. Click Create Flavor.

   The Create Flavor window is displayed.

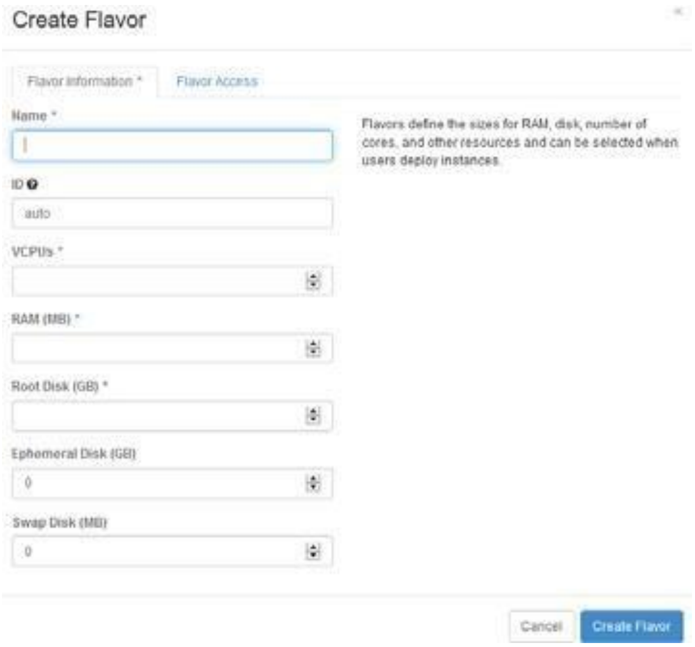

Create Flavor

Flavor Information * | Flavor Access

Name *

ID ?

auto

VCPUs *

RAM (MB) *

Root Disk (GB) *

Ephemeral Disk (GB)

0

Swap Disk (MB)

0

Flavors define the sizes for RAM, disk, number of cores, and other resources and can be selected when users deploy instances.

3 Enter a name, such as f5.medium.

4 Set VCPUs, memory, and disk values.

5 Specify which projects to apply the flavor.

If you do not specify a project, the flavor is available to all projects.

6 Click Create Flavor.

Yourflavor iscreated.

7 Repeat step 2 through step 6 for the f5.xxlarge flavor.

Next you need to create a VM snapshot.

Creating a VM Snapshot

The VM snapshot is used to createand exporta Virtualized Infrastructure Manager (VIM) software image descriptor from the Cruz application and then the VIM software image descriptor is used to manage, deploy, or migrate images.

This section uses the F5 BIGIP VE LTM image as an example. The same process applies for creating a snapshot for the F5 BIGIP VE ALL image (for the AFM model).

Gracefullyshutdown the VM manually in the OpenStack environment and then take a snapshot produced as a qcow2 image file.

Create a VM snapshot from the OpenStack dashboard as follows.

1 Shut off the instance.

a. Select Project > Compute > Instances.

A list of instances is displayed.

b. Select the Shut Off Instance action for the appropriate instance.

A confirmation message is displayed.

c. Click shut Off Instance.

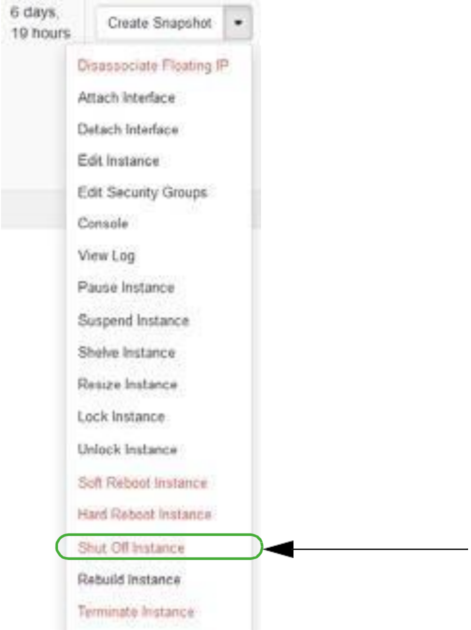

If you prefer to shutdown F5 BIGIP VM processes from the Linux shell, log in as root and enter `shutdown now`. Refer to your F5 BIG-IP documentation for more details.

2 Create a snapshot image.

a. Select Admin > System > Instances.

A list of instances is displayed.

b. Select the Create Snapshot action for the instance.

The Create a Snapshot window is displayed.

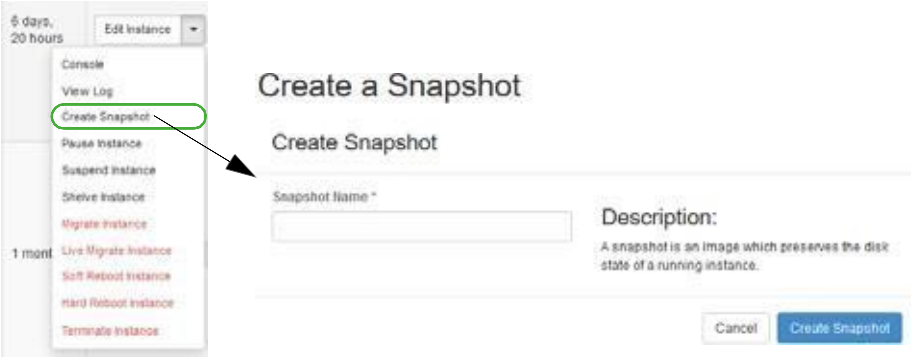

c. Enter a name, such as the original image name with the `F5-` vendor prefix and a suffix of `-SS1`.

d. Click Create Snapshot.

3 Assign OpenStack image properties to the snapshot image created in step 2.

   a. Select Admin > System > Images.

   b. Select the Edit Image action for the instance.

      The Update Image properties are displayed.

   c. Set the disk and memory requirements to match the smallest flavor that the image supports.

      See Creating Flavors on page 766 for BIGIP-VE Version 12.x.x flavors.

   d. Select Public.

      This is required so that the Cruz can access and manage the image later on.

   e. Click Update Image

      Your changes are saved.

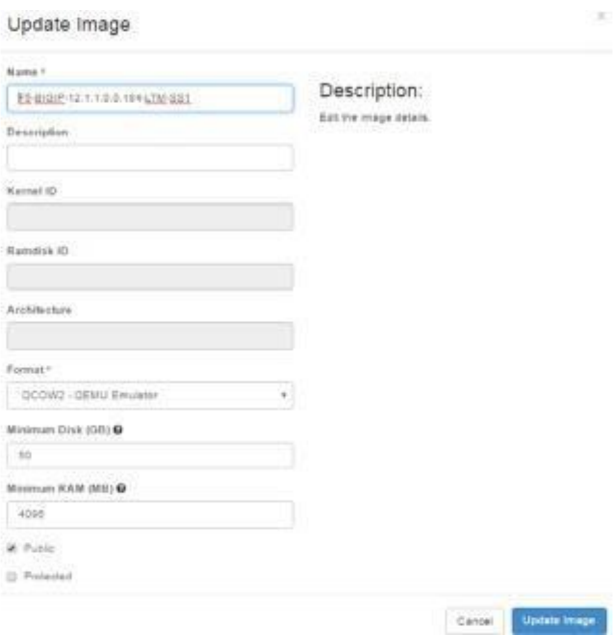

Update Image

Name *

F5-BIGIP-12.1.1.0.0.184-LTM-SS1

Description

Kernel ID

Ramdisk ID

Architecture

Format *

QCOW2 - QEMU Emulator

Minimum Disk (GB)

50

Minimum RAM (MB)

4096

☑ Public

☐ Protected

Description:

Edit the image details.

Once you create the VM snapshot image in the OpenStack environment, you are ready to create a software image record from the application.

Creating a Software Image Descriptor

The Network Functions Virtualization (NFV) feature uses the software image descriptor to manage, deploy, or migrate images. Create a software image descriptor record from the Cruz application by:

1 Creating a Software Image Descriptor for Snapshot Image

2 Modifying the VIM Software Image Descriptor

The steps provided in this section assume that you have already created a VIM that houses the snapshot image in the OpenStack environment. If a VIM that houses the snapshot image does not exist, see Setting Up Your OpenStack Environment on page 759 before continuing.

Creating a Software Image Descriptor for Snapshot Image

Create a software image descriptor for a snapshot image from the Cruz application as follows.

1 Navigate to the Software Images portlet.

The portlet’s location depends on your company’s configuration.

2 Right-click an image.

3 Select Download.

The Save VIM Image to Disk Parameters input window is displayed.

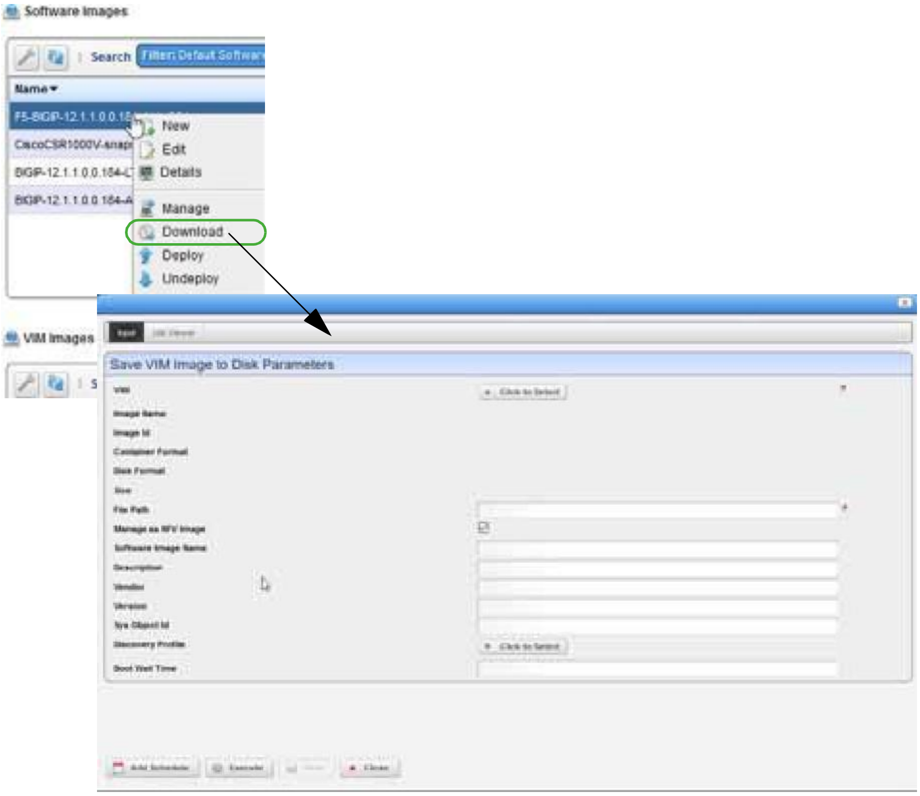

4 Select the Virtualized Infrastructure Manager (VIM) housing thesnapshot qcow2 image created earlier.

5 Select the snapshot image to download.

For example: `F5-BIGIP-12.1.1.0.0.184-LTM-SS1.qcow2`

The fields are populated based on your selection.

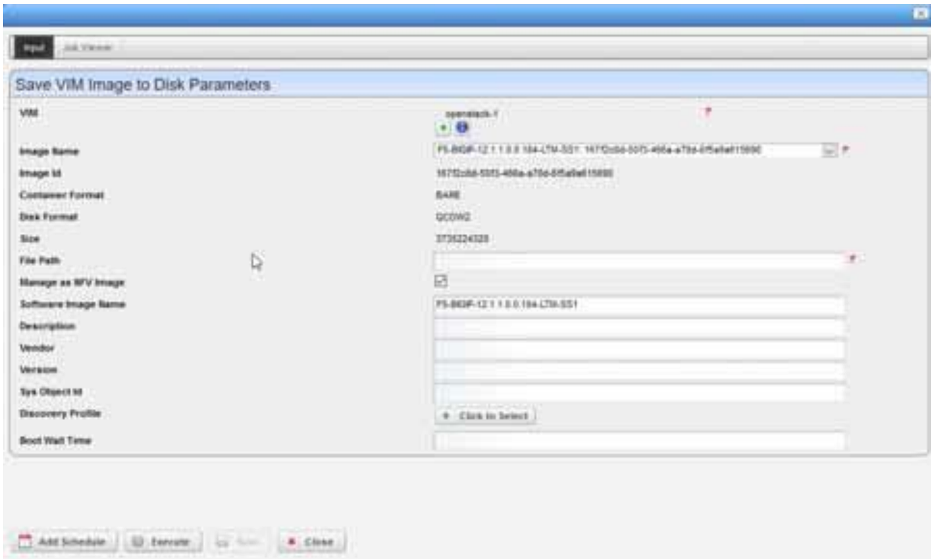

Save VIM Image to Disk Parameters

VIM	openstack-1
Image Name	F5-BIGIP-12.1.1.0.0.184-LTM-SS1: 167f2cdd-50f3-466a-a7dd-6f5a6a615690
Image Id	167f2cdd-50f3-466a-a7dd-6f5a6a615690
Container Format	BARE
Disk Format	QCOW2
Size	3736224328
File Path	
Manage as NFV Image	☑
Software Image Name	F5-BIGIP-12.1.1.0.0.184-LTM-SS1
Description	
Vendor	
Version	
Sys Object Id	
Discovery Profile	+ Click to Select
Boot Wait Time	

6 Click Execute.

7 Copy the image downloaded from your local disk to your company’s centralized file server (such as an FTP Server) to redistribute the new snapshot image across additional OpenStack VIMs.

Before the image descript or record is complete and ready toexport, you need to edit the image descriptor.

Modifying the VIM Software Image Descriptor

These instructions assume that you have already created a Virtualized Infrastructure Manager (VIM) image descriptor. If you have not created the image descriptor, see Creating a Software Image Descriptor for Snapshot Image on page 770 before continuing.

Modify the VIM software image descriptor from the application as follows.

1 Navigate to the Software Images portlet.

The portlet’s location depends on your company’s configuration.

2 Right-click an image.

3 Select Manage.

The Manage new Software Image from VIM Parameters window is displayed.

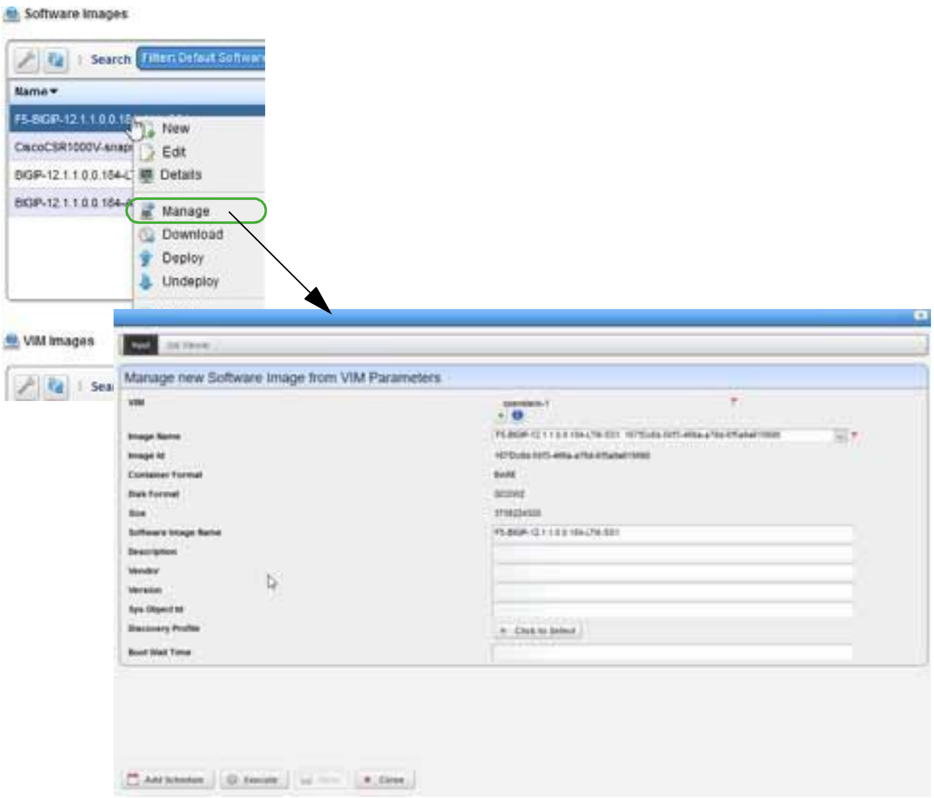

4 Select the VIM that houses the snapshot image.

5 Select the snapshot image name.

For example: `F5-BIGIP-12.1.1.0.0.184-LTM-SS1:` *`uniqueImageID`*

The image ID, container format, disk format, size, and software image name fields are populated.

6 Select a discovery profile.

7 Click Execute.

A new software image record is created matching the image name in your OpenStack environment.

8 Right-click an image.

9 Select Edit.

The Editing Software Image window is displayed.

10 Select the snapshot image from the Image Name list.

The Editing Software Image window is displayed.

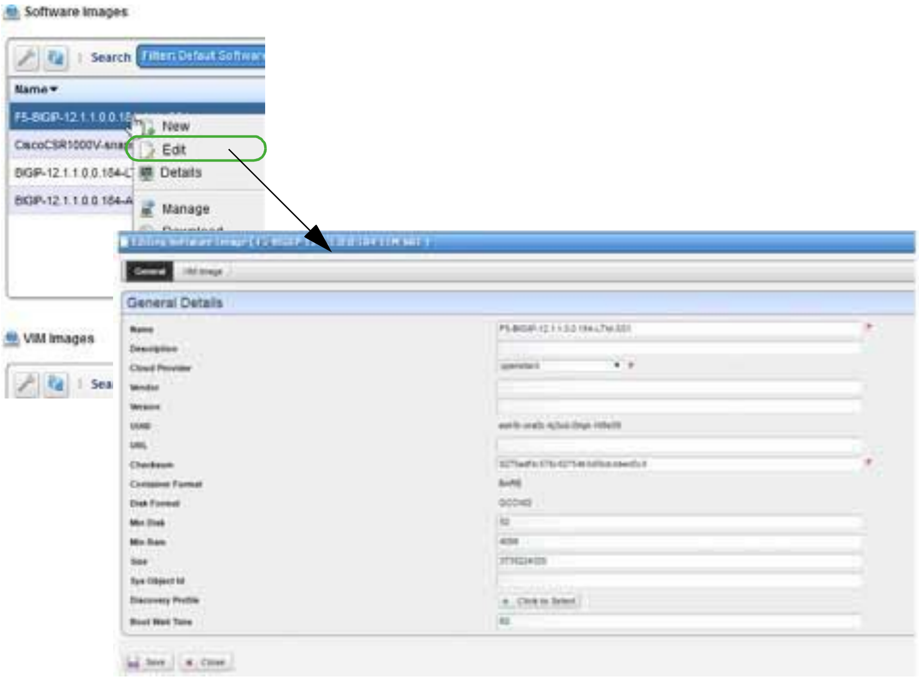

General Details

Name	F5-BIGIP-12.1.1.0.0.184-LTM-SS1
Description	
Cloud Provider	openstack
Vendor	
Version	
UUID	[illegible]
URL	
Checksum	[illegible]
Container Format	BARE
Disk Format	QCOW2
Min Disk	[illegible]
Min Ram	4096
Size	[illegible]
Sys Object Id	
Discovery Profile	Click to Select
Boot Wait Time	[illegible]

11 Enter the missing information as follows.

a. Add a meaningful description.

b. Specify the vendor, such as F5.

c. Add a meaningful version number, such as 12.1.1.

d. Enter a valid URL that points to your local environment’s file server or disk location where the qcow2 file is stored and accessible. For example:

`ftp://serverIP/bigipveltm/F5-BIGIP-12.1.1.0.0.184-LTM-SS1.qcow2`

e. Enter the SNMP SysObject ID for the F5 VE instances (for LTM, AFM and other VE models), which is:

`1.3.6.1.4.1.3375.2.1.3.4.43`

f. Select the discovery profile containing the required ssh and SNMP authentication objects.

Discovery profiles are created by your system administrator or someone within your organization with administrative permissions.

g. Leave the Boot Wait Time field blank.

The Cruz assigns a default integer value of 0 after you save the profile.

h. Review your inputs to ensure accuracy.

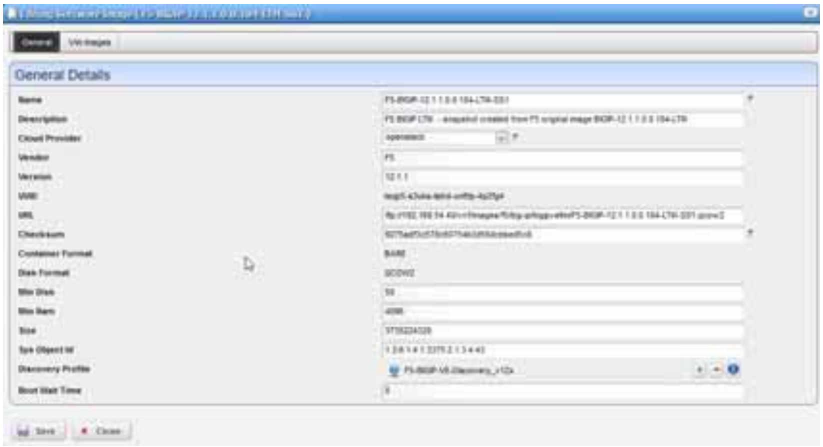

General Details

Name	F5-BIGIP-12.1.1.0.0.184-LTM-SS1
Description	F5 BIGIP LTM - snapshot created from F5 original image BIGIP-12.1.1.0.0.184-LTM
Cloud Provider	openstack
Vendor	F5
Version	12.1.1
UUID	[illegible]
URL	[illegible]
Checksum	[illegible]
Container Format	BARE
Disk Format	QCOW2
Min Disk	50
Min Ram	4096
Size	9735224320
Sys Object Id	1.3.6.1.4.1.3375.2.1.3.4.43
Discovery Profile	F5-BIGIP-VE-Discovery_v12x
Boot Wait Time	8

12 Save the software image record.

13 Locate and Select the newly created Software Image Record and export it to an archive disk location for backup purposes.

Manually Applying an F5 License

Due to license requirements and logistics for the F5 BIGIP VEinstances, you need to work directly with F5 to obtain and apply a license. This section provides an example on how to apply a license obtained from F5. The example covers applying the AFM (Firewall product feature set) license after you successfully booted a VM from the F5-BIGIP-12.1.1.0.0.184-ALL-SS1 snapshot created in Creating a VM Snapshot on page 744.

Consult the F5 Networks product documentation or contact the F5 Support team directly for further instructions on configuring the BIGIP VE VM following license activation.

Manually apply an F5 AFM (Firewall) features license from the OpenStack system as follows.

1 Login using the admin user account to the running VM housed in the OpenStack system to be licensed using https.

2 Navigate to the Main Tab.

A Setup Utility wizard is displayed.

3 Click Next.

The Setup > Utility > License information is displayed.

4 Click Activate.

5 Locate the registration key value from the text provided by F5.

6 Enter the license properties.

a. Enter the base registration key provided by F5.

b. Select Manual activation method.

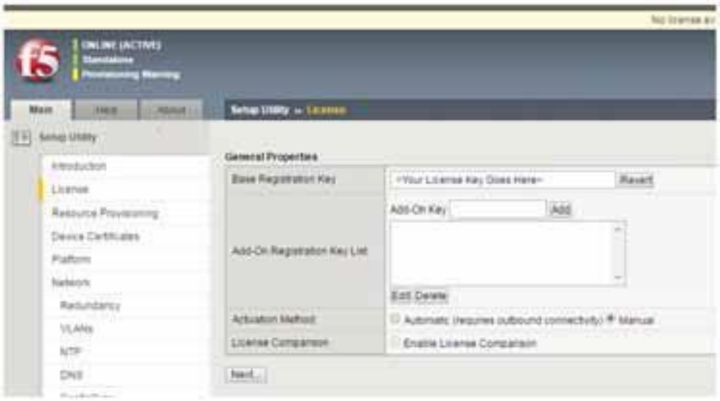

General Properties

Base Registration Key	<Your License Key Goes Here>
Add-On Registration Key List	Add-On Key
Activation Method	○ Automatic (requires outbound connectivity) ● Manual
License Comparison	□ Enable License Comparison

c. Click Next.

The General Properties, Dossier is populated.

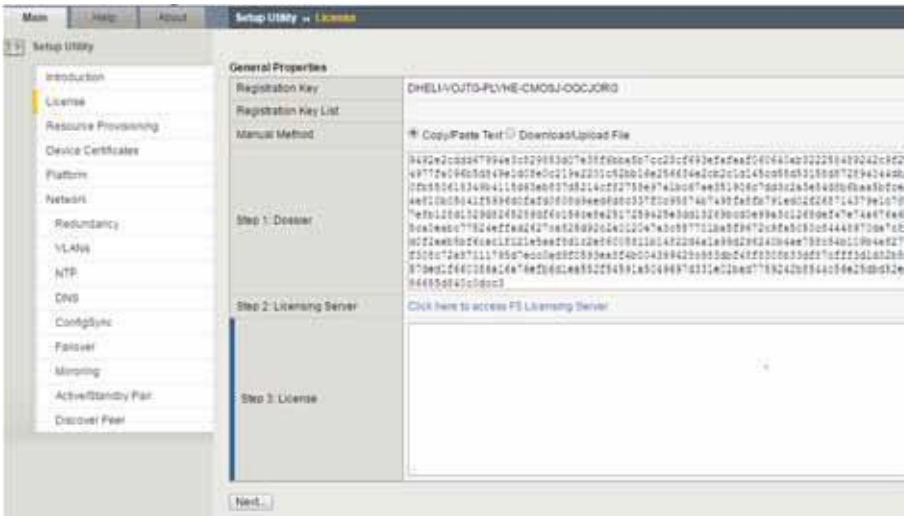

General Properties

Registration Key	DHELI-VOJTG-PLVHE-CMOSJ-OQCJORG
Registration Key List	
Manual Method	● Copy/Paste Text ○ Download/Upload File
Step 1: Dossier	[illegible]
Step 2: Licensing Server	Click here to access F5 Licensing Server
Step 3: License	

d. Select the Download/Upload File option.

e. Select “Click Here to Download Dossier File.”

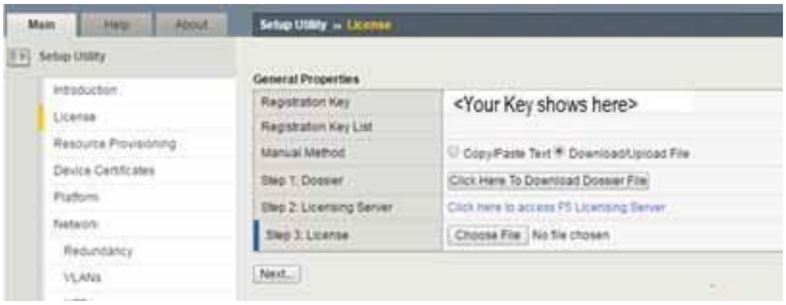

General Properties

Registration Key	<Your Key shows here>
Registration Key List	
Manual Method	○ Copy/Paste Text ● Download/Upload File
Step 1: Dossier	Click Here To Download Dossier File
Step 2: Licensing Server	Click here to access F5 Licensing Server
Step 3: License	Choose File No file chosen

f. Save the file to an accessible/writable disk location.

7 Specify the licensing server.

a. Select “Click here to access F5 Licensing Server.”

The Activate F5 Product webpage is displayed.

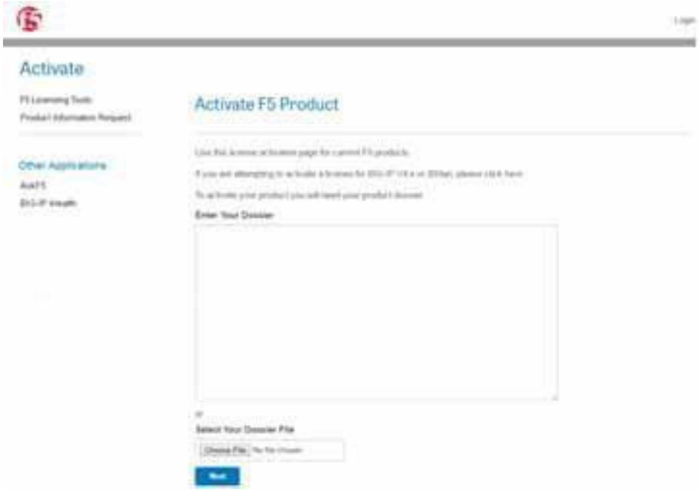

b. Click Choose File > Next.

The license agreement is displayed.

c. Select the option to accept and agree to the license terms.

d. Click Next.

A file activation page is displayed.

e. Click Download license.

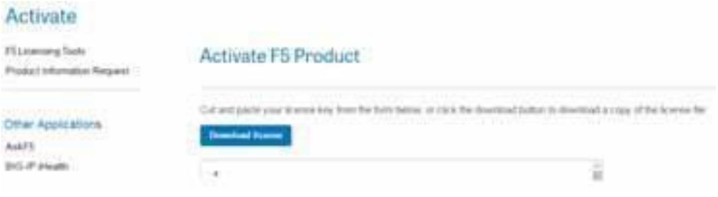

f. Save the file as License.txt to a writable disk location.

8 Go back to the F5 Setup Utility > > License General Properties page.

9 Continue with the License setup.

a. Click Choose File.

b. Select the License.txt file saved in step 7.

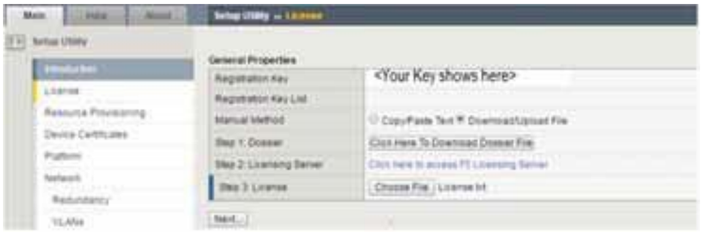

General Properties

Registration Key	<Your Key shows here>
Registration Key List	
Manual Method	○ Copy/Paste Text ● Download/Upload File
Step 1: Dossier	Click Here To Download Dossier File
Step 2: Licensing Server	Click here to access F5 Licensing Server
Step 3: License	Choose File License.txt

c. Click Next.

A status message is displayed indicating that the configuration data is loading followed by a message that the license configuration changes are in progress.

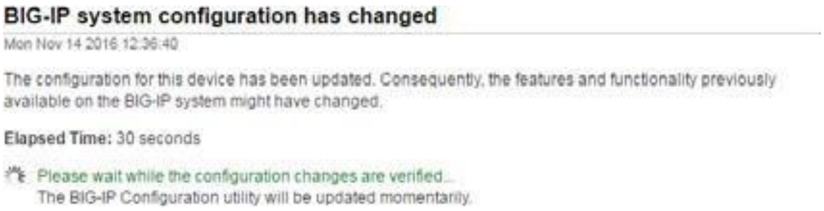

BIG-IP system configuration has changed

Mon Nov 14 2016 12:36:40

The configuration for this device has been updated. Consequently, the features and functionality previously available on the BIG-IP system might have changed.

Elapsed Time: 30 seconds

Please wait while the configuration changes are verified...
The BIG-IP Configuration utility will be updated momentarily.

d. Wait until a process completed message is displayed.

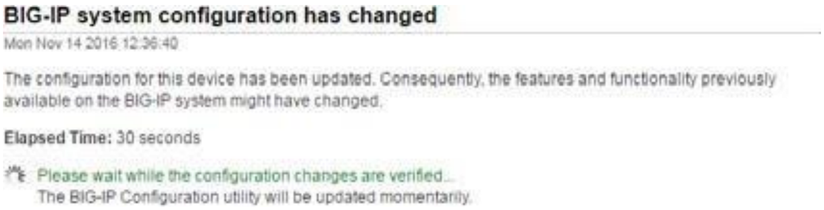

BIG-IP system configuration has changed

Mon Nov 14 2016 12:36:40

The configuration for this device has been updated. Consequently, the features and functionality previously available on the BIG-IP system might have changed.

Elapsed Time: 30 seconds

Please wait while the configuration changes are verified...
The BIG-IP Configuration utility will be updated momentarily.

10 Verify the activated license by selecting the Setup Utility, License tab on the F5 BIGIP page.

For example: Applying an Evaluation AFM (Firewall) license.

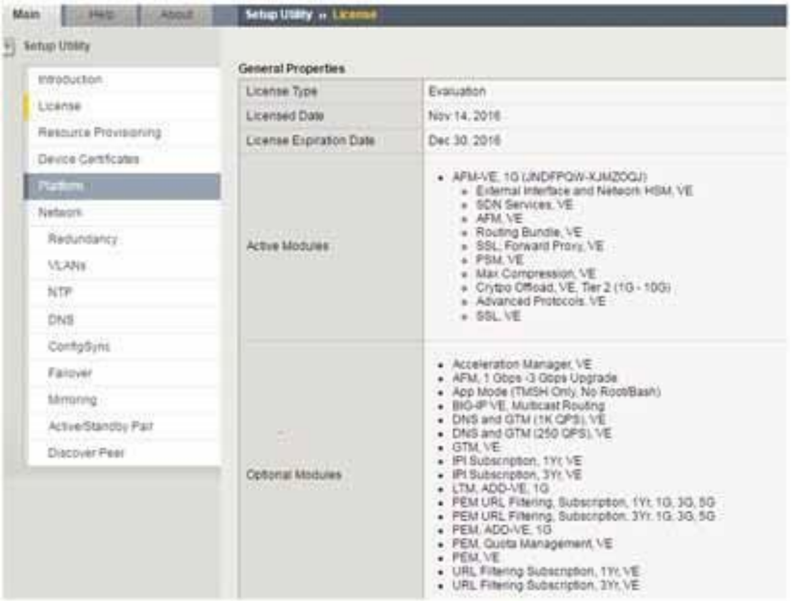

Setup Utility » License

General Properties

License Type	Evaluation
Licensed Date	Nov 14, 2016
License Expiration Date	Dec 30, 2016
Active Modules	AFM-VE, 1G (JNDFPQW-XJMZOQJ) External Interface and Network HSM, VE SDN Services, VE AFM, VE Routing Bundle, VE SSL, Forward Proxy, VE PSM, VE Max Compression, VE Crytpo Offload, VE, Tier 2 (1G - 10G) Advanced Protocols, VE SSL, VE
Optional Modules	Acceleration Manager, VE AFM, 1 Gbps -3 Gbps Upgrade App Mode (TMSH Only, No Root/Bash) BIG-IP VE, Multicast Routing DNS and GTM (1K QPS), VE DNS and GTM (250 QPS), VE GTM, VE IPI Subscription, 1Yr, VE IPI Subscription, 3Yr, VE LTM, ADD-VE, 1G PEM URL Filtering, Subscription, 1Yr, 1G, 3G, 5G PEM URL Filtering, Subscription, 3Yr, 1G, 3G, 5G PEM, ADD-VE, 1G PEM, Quota Management, VE PEM, VE URL Filtering Subscription, 1Yr, VE URL Filtering Subscription, 3Yr, VE

11 Continue on through the rest of the Setup Wizard.

The Resource Provisioning page should show which modules are nowlicensed, such as an AFM (Firewall) evaluation license.

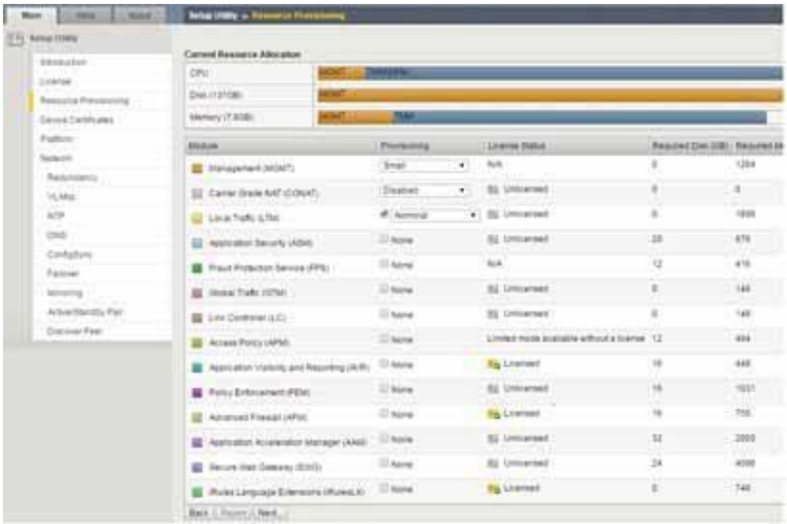

Current Resource Allocation

CPU	MGMT / TMM/PPM
Disk (121GB)	MGMT
Memory (7.8GB)	MGMT / TMM

Module	Provisioning	License Status	Required Disk (GB)	Required Me
Management (MGMT)	Small	N/A	0	1264
Carrier Grade NAT (CGNAT)	Disabled	Unlicensed	0	0
Local Traffic (LTM)	✓ Nominal	Unlicensed	0	1699
Application Security (ASM)	□ None	Unlicensed	20	676
Fraud Protection Service (FPS)	□ None	N/A	12	416
Global Traffic (GTM)	□ None	Unlicensed	0	144
Link Controller (LC)	□ None	Unlicensed	0	148
Access Policy (APM)	□ None	Limited mode available without a license	12	494
Application Visibility and Reporting (AVR)	□ None	Licensed	16	448
Policy Enforcement (PEM)	□ None	Unlicensed	16	1031
Advanced Firewall (AFM)	□ None	Licensed	16	755
Application Acceleration Manager (AAM)	□ None	Unlicensed	32	2000
Secure Web Gateway (SWG)	□ None	Unlicensed	24	4096
iRules Language Extensions (iRulesLX)	□ None	Licensed	0	748

Consult F5’s product documentation or contact F5 Support team directly for further instructions on configuring the BIGIP VE VM following license activation.

Setting Up the Juniper Firefly VNF Package

Thissection explains the scope, functions, and processesassociated with the Virtualized Network Function (VNF) onboarding and instantiation:

- Understanding the Firefly VNF Package
- Preparing an Image

Understanding the Firefly VNF Package

The Juniper Networks Firefly SRX model is a virtual appliance that provides security and networking services at theedge in avirtualized private or public cloud environment. Th is appliance runs on a virtual machine (VM) on a server.

The Cruz Firefly Virtualized Network Function(VNF) packageprovidessample VNFdescriptors and associated artifacts used to deploy the Juniper Firefly SRX model described in this section.

Juniper Networksprovides images and licensesneeded to deployand enable the Juniper Firefly SRX features.

The following list provides the VNF package vendor and validation information:

Vendor and Model Information	
Vendor	Juniper Networks
Product Family	Firefly
Models	SRX
Operating Systems	Junos$^{\circledR}$ OS
Version	12.1
Validated Environment	
VIM	OpenStack
Version	Liberty
Installer	Mirantis
Version	8
Network Configurations	VXLAN, VLAN Segmentation
Cruz **MANO (Management and Operation) Validated Functions**	
MANO	NSD, VNFD, NSR, VNFR Link LCEs
VNFM	Generic (MANO Lite) Independent/Specific
NFVO	MANO Lite
Cruz **RMO (Resource Management and Operation) Validated Functions**	
Inventory	Resource Discovery and Inventory Resource Creation Method (NFV-Model)
Monitoring	SNMP Monitoring CPU & Memory KPIs TFA (flows)
Config	ACLI

NetRestore File Management	NR Backup NR Restore NR Firmware (not validated)

This section covers the:

- Firefly VNF Package Contents
- Cruz VNF Package - Install Component Files
- License Requirements
- Known Issues
- References

Once you have a good understanding of the VNF package, you can get started preparing an image.

Firefly VNF Package Contents

The Firefly VNFpackage containsartifacts necessaryto installand deploy Virtualized Network Functions (VNFs) for the SRX Virtualization Deployment Unit (VDU) offered as part of the Juniper Firefly product family.

Multiple sample VNF descriptors were created with basic flavors to deploythe VNF VDUs, either as part of a standalone VN F record, or as part of a higher level, end-to-end Network Service record:

- NFV VNF Descriptors
- Network Service Descriptor

NFV VNF Descriptors

Thissectiondescribes thefollowing Cruz Network Functions Virtualization (NFV) VNF (virtualized network functions) descriptors:

- Juniper Firefly SRX Cluster
- Juniper Firefly Virtual Appliance

Juniper Firefly SRX Cluster. The sample Juniper Firefly SRX cluster descriptor is located in the following file:

`owareapps/nfv-vnfm-firefly/db/vnfd.juniper-firefly-cluster.xml`

This Virtualized Network Function (VNF) descript or defines a standalone network as a varied composition of Juniper Firefly Virtualization Deployment Units (VDUs), based on the assigned VNF favor.

Here are the **Juniper Firefly SRX** descriptor deploy parameters used:

Field Name	Data Type	Description	Default Value	Valid Values	Required Values
Management-Network	IP Address	Address of the Management Network to be created on the OpenStack system	172.191.0.0/24	Valid CIDR Address	Y
LAN1 - Network	IP Address	Address of the LAN1 Network to be created on the OpenStack system	172.191.1.0/24	Valid CIDR Address	Y
LAN2 – Network	IP Address	Address of the LAN2 Network to be created on the OpenStack system	172.191.2.0/24	Valid CIDR Address	Y

Field Name	Data Type	Description	Default Value	Valid Values	Required Values
Cluster-Name	String	Name of the Cluster identifying the group of networks and interconnected VDUs	None	Alpha-numeric String	Y

Here are the **Juniper Firefly SRX** descriptor flavors used:

Flavor Name	VDU Name	VDU Count	CPU	Disk (GB)	Memory (MB)
Basic	Firefly-SRX	1	2	2	2048

Viewthe VNF descriptor's composition details from the Cruz Onboarding page > VNF Descriptor portlet.

Juniper Firefly Virtual Appliance. The sample Juniper Firefly Virtual Appliance descriptor is located in the following file:

```
owareapps/nfv-vnfm-firefly/db/vnfd.juniper-firefly-virtual-appliance.xml
```

Thissample Virtualized Network Function(VNF) descriptor defines packaging and behaviorofone or more Firefly SRX Virtualization Deployment Units (VDUs) and can be included as part of a Network Descriptor.

Here are the **Juniper Firefly Virtual Appliance** descriptor flavors:

Flavor Name	VDU Name	VDU Count	CPU	Disk (GB)	Memory (MB)
Firefly Basic	Firefly Virtual Appliance	1	2	2	2048

Viewthe VNF descriptor's composition details from the Cruz Onboarding page > VNF Descriptor portlet.

Network Service Descriptor. The sample Network Service descriptor (NSD) provided shows how to define a Network Service and utilize any combination of the Virtualized Network Function (VNF) descriptors provided in this package.

View the Network Service descriptor contents from the Cruz Network Service Descriptors portlet.

NOTE:

TheNetworkService Descriptors portletlocationis dependent onyourcompany's Cruz NFVinstallation and configuration. By default, the Network Service Descriptors portlet is located on the On-Boarding page and the Network Services page.

This sample NSD is located in the following file:

```
owareapps/nfv-vnfm-firefly/db/ns.juniper-firefly-sample.xml
```

Here are the **Network Service** descriptor flavors used:

NS Flavor Name	VNF Flavor Names	VDU Name	VDU Count	CPU	Disk (GB)	Memory (MB)
Standard-Basic	Firefly Basic	Firefly Virtual Appliance	1	2	2	2048

Cruz VNF Package - Install Component Files

The following Virtualized Network Function (VNF) packageinstallation componentsare required for the Juniper Firefly product:

VNF (OCP)	nfv-vnfm-firefly.ocp
Device Driver (DDP)	juniper.ddp

The Cruz VNF package descriptors for the Juniper Firefly product include:

NVF VNF Package			Resource Management		
VIM Software Image Records	VNF Descriptors	Sample NSD	Discovery Authentication Records	Discovery Profiles	Device Driver
Y	Y	Y	Y	Y	Y

The VNF package descriptors include:

- Discovery Authentication Profiles
- Discovery Profiles
- NFV VIM Software Image Descriptors

Discovery Authentication Profiles

The Juniper Firefly Discovery Authentication profiles are located in the following file:

```
owareapps/nfv-vnfm-firefly/db/rc.disc-auths-firefly.xml
```

This file contains the following discovery authentication profiles:

Authentication Record Name	Description
Juniper Firefly SNMP	SNMP v2 Credentials
Juniper Firefly CLI	SSH v2 Credentials

Discovery Profiles

The Juniper Firefly Discovery profiles are located in the following file:

```
owareapps/nfv-vnfm-firefly/db/rcdisc-profile-firefly.xml
```

This filecontains the Juniper Firefly Discovery profile that automatically discoversand populates resource records during Virtualized Network Functionality (VNF) record instantiation. The discovery profile name is JuniperFirefly_Discovery.

The following tasks run once the target VM in the OpenStack environment is fully booted and can respond to SNMP and CLI requests from the management system:

1. Resync
2. DataCollection ForGroupOfDevices
3. Discovery Links for a Group of Devices
4. Refresh Monitor Targets
5. Juniper JUNOS SNMP Trap Forward to Cruz

NFV VIM Software Image Descriptors

See Preparing an Image for image customization required to produce and use this snapshot.

Name	Scope/Description	Container Format	Disk Format	Min RAM (MB)	MIN Disk (GB)	Resource Discovery Profile
Juniper-Firefly-snapshot-2	Juniper Firefly Snapshot 2 - Customized for Cruz Resource Management - created from original Juniper image junos-vsrx-12.1X46-D25.7-domestic.jva	BARE	QCOW2	2048	2	Y

License Requirements

Contact Juniper Networksregardingpricingand logisticsforobtainingand installing licensefor the Juniper Firefly (SRX).

Known Issues

There are no known issues currently.

References

This section provides references to:

- Vendor Documentation (URLs)
- VNF Configuration Methods

Vendor Documentation (URLs)

Referto the Juniper Networkswebsite (http://www.juniper.net) forallvendor-specific documentation.

For the Juniper Networks documentation and release not es, refer to the following Juniper websites page:

`http://www.juniper.net/support/downloads/?p=vsrx#docs`

`http://www.juniper.net/support/downloads/?p=firefly`

VNF Configuration Methods

A combination of image (snapshot) preparation, along with Virtualization Deployment Unit (VDU) Post Instantiate Configuration Lifecycle Events (LCEs), is implemented in this VNF Package to ensure that the Juniper Firefly VDUs are automatically brought under management.

Customized Image Preparation. The base Firefly SRXmodelfamilyimagesprovidedby Juniper requires additional manual configuration before producing VMs in the OpenStack environment that can be discovered and managed using the Resource Management applications.

See Preparing an Image on page 784 for a description of how to create an updated image (snapshot) for the Firefly SRX model series to use to boot up VMs that the Cruz application manages. Use the snapshot produced to create Firefly VMs in the OpenStack environment, so that no manual intervention is required to apply the management system communication settings.

Cruz Resource Configuration Tasks. Once the target VM in the OpenStack environment is fully booted and can respond to SNMP and CLI requests from the management system, post discovery tasks are run.

For more details about the JuniperFirefly_Discovery profile and the tasks that run, see Discovery Profiles on page 782.

Preparing an Image

Before you can discover and manage a Virtualized Network Functionality (VNF) resource from the application, you need to:

1. Set up your OpenStack environment by preparing the Firefly VIM image with the necessary configuration.
2. Create a software image record from the Cruz application.

The junos-vsrx-12.1X46-D25.7-domestic.jva file provided by Juniper Networks must be converted to a qcow2 format before configuration can be performed.

Setting Up Your OpenStack Environment

To set up your OpenStack environment, you need to:

1. Convert the Firefly self-extracting package file (.jva) to a qcow2 file format.
2. Create a project.
3. Deploy the vendor's base image junos-vsrx-domestic.jva file by uploading it to a target Bare Metal OpenStack controller and then converting it to a qcow2 format.
4. Create a flavor representing the smallest supported flavor to use (Figure 16-5). The Juniper vendor provides the following JUNOS VSRX Version 12.x.x flavor:

Juniper-Firefly-snapshot-2	Flavor Name: cpu-2-ram-2048-disk-2	CPU – 2, Memory – 2 GB (2048 MB), Disk – 2 GB

5. Make configuration changes to the snapshot image.
6. Create a VM snapshot that is used to create and export a Virtualized Infrastructure Manager (VIM) software image descriptor within the application.

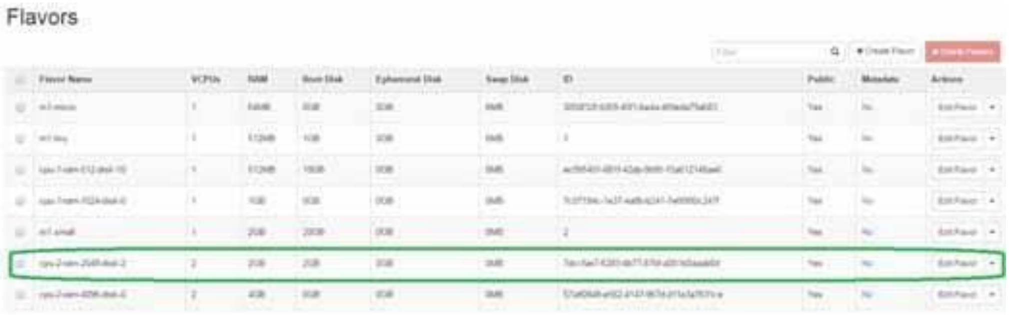
Flavors

	Flavor Name	VCPUs	RAM	Root Disk	Ephemeral Disk	Swap Disk	ID	Public	Metadata	Actions
□	m1.micro	1	64MB	0GB	0GB	0MB	[illegible]	Yes	No	Edit Flavor ▾
□	m1.tiny	1	512MB	1GB	0GB	0MB	1	Yes	No	Edit Flavor ▾
□	cpu-1-ram-512-disk-10	1	512MB	10GB	0GB	0MB	[illegible]	Yes	No	Edit Flavor ▾
□	cpu-1-ram-1024-disk-0	1	1GB	0GB	0GB	0MB	[illegible]	Yes	No	Edit Flavor ▾
□	m1.small	1	2GB	20GB	0GB	0MB	2	Yes	No	Edit Flavor ▾
□	cpu-2-ram-2048-disk-2	2	2GB	2GB	0GB	0MB	[illegible]	Yes	No	Edit Flavor ▾
□	cpu-2-ram-4096-disk-0	2	4GB	0GB	0GB	0MB	[illegible]	Yes	No	Edit Flavor ▾

Figure 16-5. Juniper Firefly Flavors Example

Converting JVA to QCOW2

Convert the Firefly self-extracting package file (.jv a) to a qcow2 file format from the command line as follows.

1. Upload junos-vsrx-domestic.jva file you received from Juniper to a target Bare Metal OpenStack controller.
2. Access the OpenStack controller using the `ssh` command.
3. Navigate to the var/tmp directory.
4. Verify that the .jva file is located in the junos-vsrx-domestic-NUMBER directory.
5. Enter the following command:

   ```
   glance image-create --name "Juniper-Firefly-InitialConversion" --disk-
     format qcow2 --container-format bare --is-public True --file junos-vsrx-
     domestic.img
   ```

   This creates the qcow2 file.
6. Verify that the qcow2 image was created by entering `glance image-list`.

Now that you have converted the .jva file to a qcow2 format, you are ready to create a project.

Creating a Project

Create a project in an OpenStack VIM (Virtualized Infrastructure Manager) as follows.

1. Select Identity > Project.
2. Click Create Project.

   The Create Project window is displayed.
3. Enter a name and detailed descriptions.

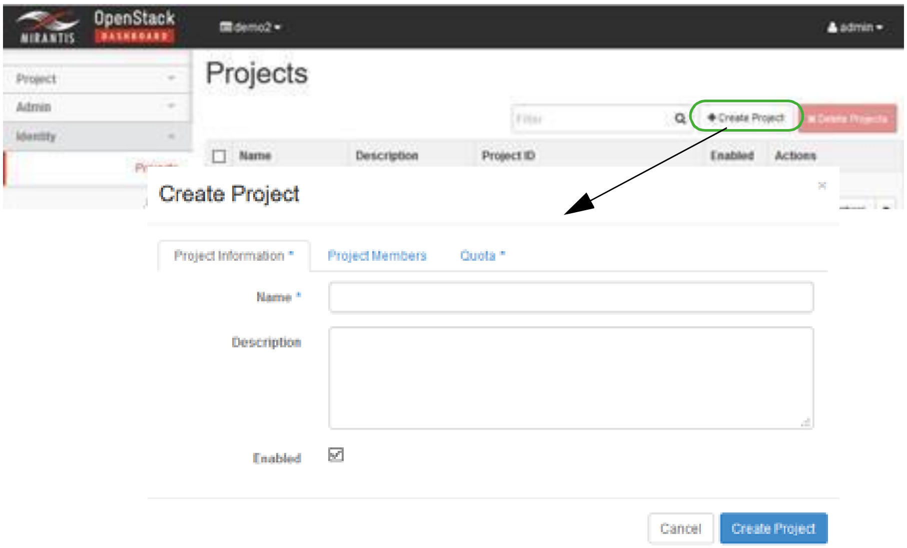

4 Assign users/permissions.

a. Click Project Members.

b. Add the following users to the Project Members list.

- heat-cfn
- heat
- admin

c. Set the member permissions to admin.

5 Click Create Project.

The new project is saved to the Projects list.

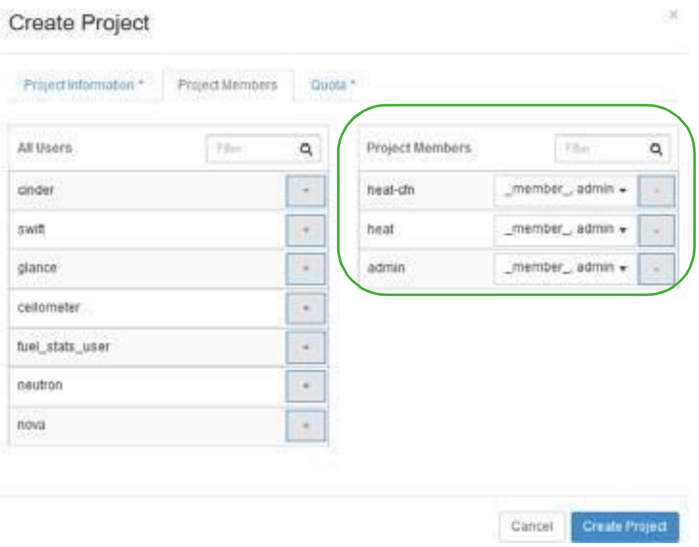

Create Project

Project Information * | Project Members | Quota *

All Users	
cinder	+
swift	+
glance	+
ceilometer	+
fuel_stats_user	+
neutron	+
nova	+

Project Members		
heat-cfn	_member_, admin	-
heat	_member_, admin	-
admin	_member_, admin	-

6 Create project networking in an OpenStack VIM (KILO or later release).

a. Create a project router with a gateway interface.

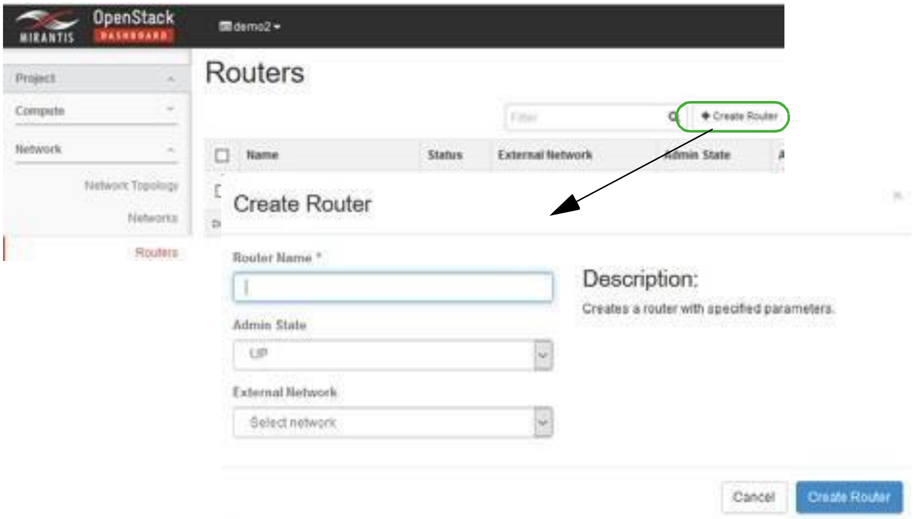

Routers

Create Router

Router Name *

Admin State: UP

External Network: Select network

Description:

Creates a router with specified parameters.

b. Make sure that the router's gateway IP is reachable from the LAN (your PC) following project router creation.

c. Create and attach the following networks to the project router:

- net-mgmt
- net-data1
- net-data2

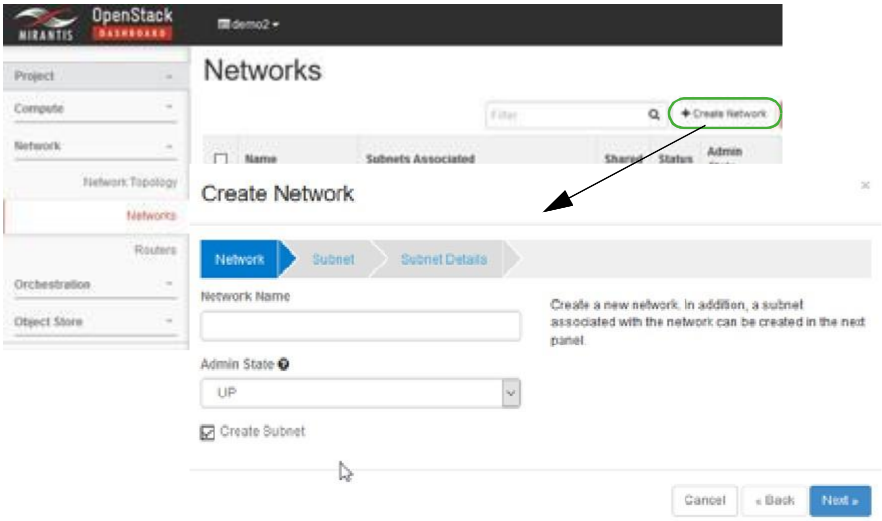

Create a new network. In addition, a subnet associated with the network can be created in the next panel.

Now you can instantiate a VM from the vendor-created image.

Instantiating a VM from a Vender-Converted Image

Instantiate a VM (virtual machine in the new OpenStack project from the Juniper-Firefly-snapshot-2 image created in the Creating a VM Snapshot section as follows.

1 Boot up the Juniper-Firefly-snapshot-2 image with the following network connections.

Linux interface name	Firefly Interface Name	Network
eth0	mgmt	net-mgmt
<None > *	1.1	net-data-1
<None > *	1.2	net-data-2
* The < None> int erface names require CLI configuration along with a pro per license before the interface can be configured/enabled and pass traffic. This is done as post-bootup and post device discovery configuration applied by the system (not part of the base image snapshot).		

2 Assign a floating IP address to the eth0/mgmt connection.

3 Go to the OpenStack dashboard.

4 Select Project > Compute > Instances.

The instances list isdisplayed.

5 Verify that the deployed instance is listed.

Create flavors before creating the VM snapshots.

Creating a Flavor

Create the following Firefly flavor from the OpenStack dashboard.

Juniper-Firefly-snapshot-2	Flavor Name: cpu-2-ram-2048-disk-2	CPU – 2, Memory – 2 GB (2048 MB), Disk – 2 GB

1 Select Admin > System > Flavors.

2 Click Create Flavor.

The Create Flavor window is displayed.

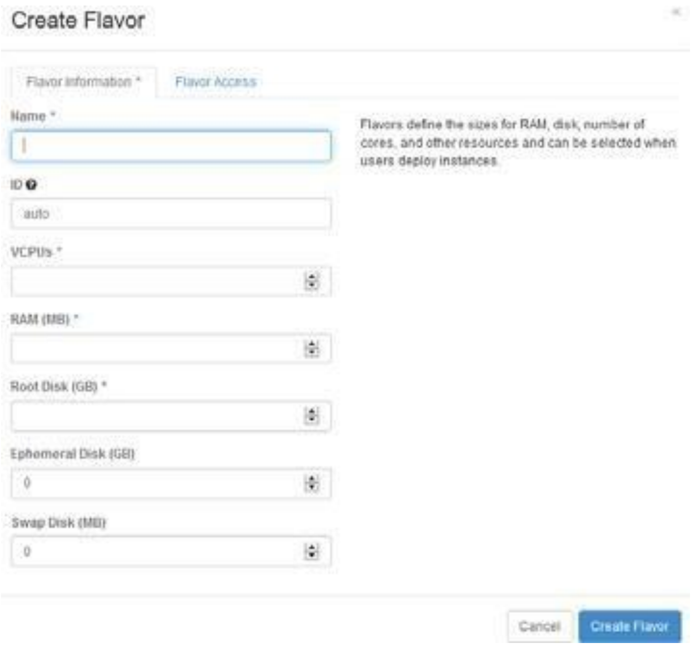

Create Flavor

Flavors define the sizes for RAM, disk, number of cores, and other resources and can be selected when users deploy instances.

ID: auto

Ephemeral Disk (GB): 0

Swap Disk (MB): 0

3 Enter a name, such as cpu-2-ram-2048-disk-2.

4 Set VCPUs, memory, and disk values.

5 Specify which projects to apply the flavor.

If you do not specify a project, the flavor is available to all projects.

6 Click Create Flavor.

Yourflavoriscreated.

Now it is time to manually configure the converted image.

Manually Configuring the Snapshot

The Juniper-Firefly-Initial conversion image must now be deployed and configured manually to create a snapshot that functions with the application.

Manually configure the snapshot from the OpenStack dashboard as follows.

1 Select Admin > System > Images.

2 Launch a Juniper-Firefly-InitialConversion instance from Images.

3 Makesurethat youselectthecorrectflavor(cpu-2-ram-2048-disk-2) and thecreated networks are included.

4 Wait until instantiation completes.

5 Enter the instance overview.

6 Navigate to the instance console.

7 Login as root.

8 Configure the admin user (admin/D0rado!!) as follows:

```
logout

Amnesiac (ttyv0)

Login: root

--- JUNOS 12.1X16-D25.7 built 2012-09-06 01:40:34 UTC
root@% cli
root> edit private
warning: uncommitted changes will be discarded on exit
Entering configuration mode

[edit]
root# edit system

[edit system]
root# set root-authentication plain-text-password
New password:
Retype new password:

root# edit loginuser admin

[edit system login user admin]
root# set uid 2001

[edit system login user admin]
root# set class super-user

[edit system login user admin]
root# set pl
syntax error.
root# set authentication plain-text-password
New password:
Retype new password:

[edit system login user admin]
root# top

[edit]
root# commit
commit complete

[edit]
root#
```

9 Login as the admin user.

10 Configure the interfaces for each network connection asfollows:

```
ge-0/0/0: net-mgmt network port
ge-0/0/1: net-data-1 network port
ge-0/0/2: net-data-2 network port
```

For example, configure the ge-0/0/0 port (repeat this process for ge-0/0/1 and ge-0/0/2) “set unit 0 family inet dhcp” as follows:

```
[edit]
admin@vSRX1# show interfaces ge-0/0/0
unit 0 {
  family inet {
      dhcp;
  }
}
```

11 Configure security zones.

a. Delete ge-0/0/0 interface from security-zone untrust (default settings).

b. Add the ge-0/0/0 interface to security-zone trust and add host-inbound-traffic settings as follows:

```
[edit]
admin@vSRX1# show security zones
security-zone trust {
                 tcp-rst;
                 interfaces {
                     ge-0/0/0.0 {
                         host-inbound-traffic {
                             system-services {
                                 all;
                             }
                             protocols {
                                 all;
                             }
                         }
                     }
                 }
}
security-zone untrust {
                 screen untrust-screen;
}

[edit]
admin@vSRX1#
```

Next you need to create a VM snapshot.

Creating a VM Snapshot

The VM snapshot is us ed to create and export a Virtualized Infrastructure Manager (VIM) software image descriptor from the application and then the VIM software image descriptor is used to manage, deploy, or migrate images.

Gracefully shutdown the VM manually in the OpenStack environment and t hen take a snapshot produced as a qcow2 image file.

Create a VM snapshot from the OpenStack dashboard as follows.

1 Shut off the instance from.

a. Select Project > Compute > Instances.

A list of instances is displayed.

b. Select the Shut Off Instance action for the appropriate instance.

A confirmation message is displayed.

c. Click shut Off Instance.

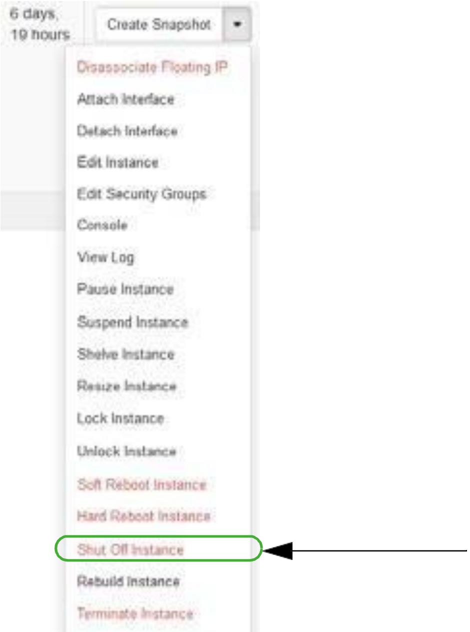

If you prefer to shutdown Firefly VM processes from the Firefly SRX command-line interface (CLI) shell, use the `request system power-off` command or the `request system halt` command. Refer to your Juniper firef ly documentation for more details.

2 Create a snapshot image.

a. Select Admin > System > Instances.

A list of instances is displayed.

b. Select the Create Snapshot action for the instance.

The Create a Snapshot window is displayed.

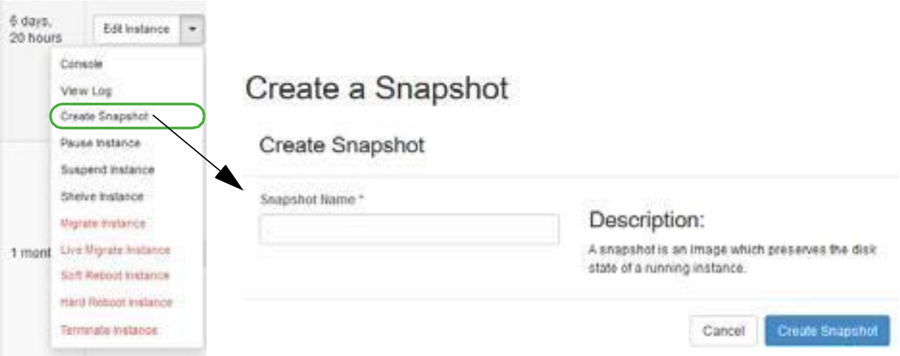

Create a Snapshot

Create Snapshot

Snapshot Name *

Description:

A snapshot is an image which preserves the disk state of a running instance.

Cancel | Create Snapshot

c. Enter a name, such as Juniper-Firefly-Snapshot-2.

d. Click Create Snapshot.

3 Assign OpenStack image properties to the snapshot image created in step 2.

a. Select Admin > System > Images.

b. Select the Edit Image action for the instance.

The Update Image properties are displayed.

c. Set the Disk to 2 and the Memory requirements to 2048.

d. Select Public.

This is required so that the Cruz system can access and manage the image later on.

e. Click Update Image

Your changesare saved.

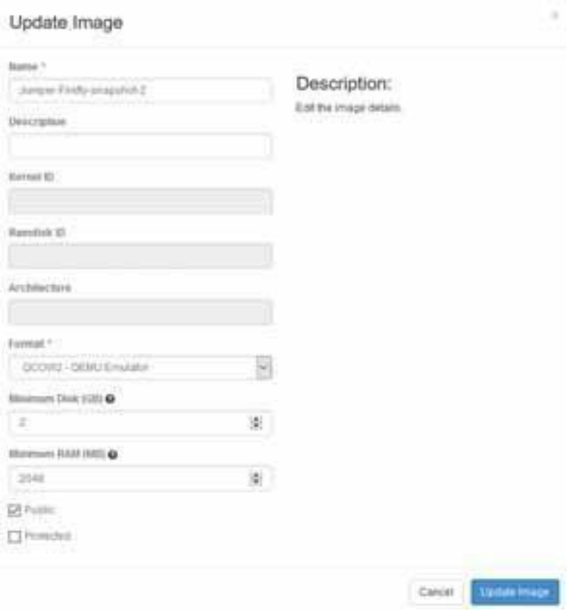

Update Image

Name *
Juniper-Firefly-snapshot-2

Description

Kernel ID

Ramdisk ID

Architecture

Format *
QCOW2 - QEMU Emulator

Minimum Disk (GB)
2

Minimum RAM (MB)
2048

☑ Public

□ Protected

Description:

Edit the image details.

Cancel | Update Image

Now you are ready to create a software image descriptor from the Cruz application.

Creating a Software Image Descriptor

The Network Functions Virtualization (NFV) feature uses the software image descriptor to manage, deploy, or migrate images. Create a software image descriptor record from the Cruz application by:

1. Creating a Software Image Descriptor for Snapshot Image
2. Modifying the VIM Software Image Descriptor

The steps provided in this section assume that you have already created a VIM that houses the snapshot image in the OpenStack environment. If a VIM that houses the snapshot image does not exist, see Setting Up Your OpenStack Environment on page 784 before continuing.

Creating a Software Image Descriptor for Snapshot Image

Create a software image descriptor for a snapshot image from the Cruz application as follows.

1. Navigate to the Software Images portlet.

   The portlet's location depends on your company's configuration.

2. Right-click an image.
3. Select Download.

   The Save VIM Image to Disk Parameters input window is displayed.

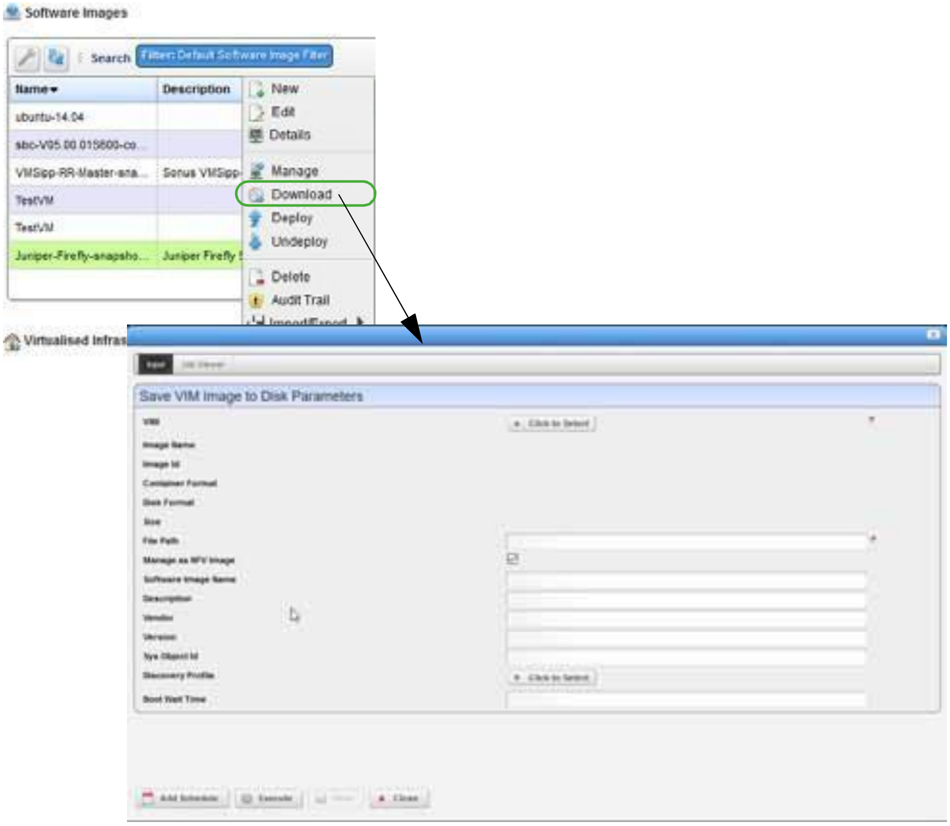

4. Select the Virtualized Infrastructure Manager (VIM) housing thesnapshot qcow2 image created earlier.
5. Select the snapshot image to download.

For example: `Juniper-Firefly-snapshot-2.qcow2`

The fields are populated based on your selection.

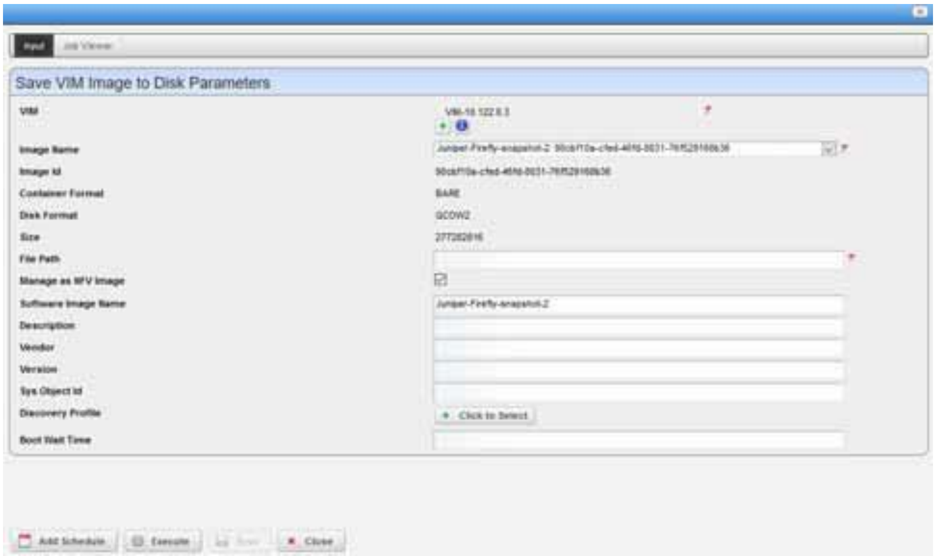

Save VIM Image to Disk Parameters

VIM	VM-10.122.0.3
Image Name	Juniper-Firefly-snapshot-2 90cbf10a-cfed-46f6-8031-76f529168636
Image Id	90cbf10a-cfed-46f6-8031-76f529168636
Container Format	BARE
Disk Format	QCOW2
Size	277282816
File Path	
Manage as NFV Image	☑
Software Image Name	Juniper-Firefly-snapshot-2
Description	
Vendor	
Version	
Sys Object Id	
Discovery Profile	Click to Select
Boot Wait Time	

6 Click Execute.

7 Copy the image downloaded from your local disk to your company’s centralized file server (such as an FTP Server) to redistribute the new snapshot image across additional OpenStack VIMs.

Before the image descriptor record is complete and ready to export, you need to edit the image descriptor.

Modifying the VIM Software Image Descriptor

These instructions assume that you have already created a Virtualized Infrastructure Manager (VIM) image descriptor. If you have not created the image descriptor, see Creating a Software Image Descriptor for Snapshot Image on page 793 before continuing.

Modify the VIM software image descriptor from the application as follows.

1 Navigate to the Software Images portlet.

The portlet’s location depends on your company’s configuration.

2 Right-click an image.

3 Select Manage.

The Manage new Software Image from VIM Parameters window is displayed.

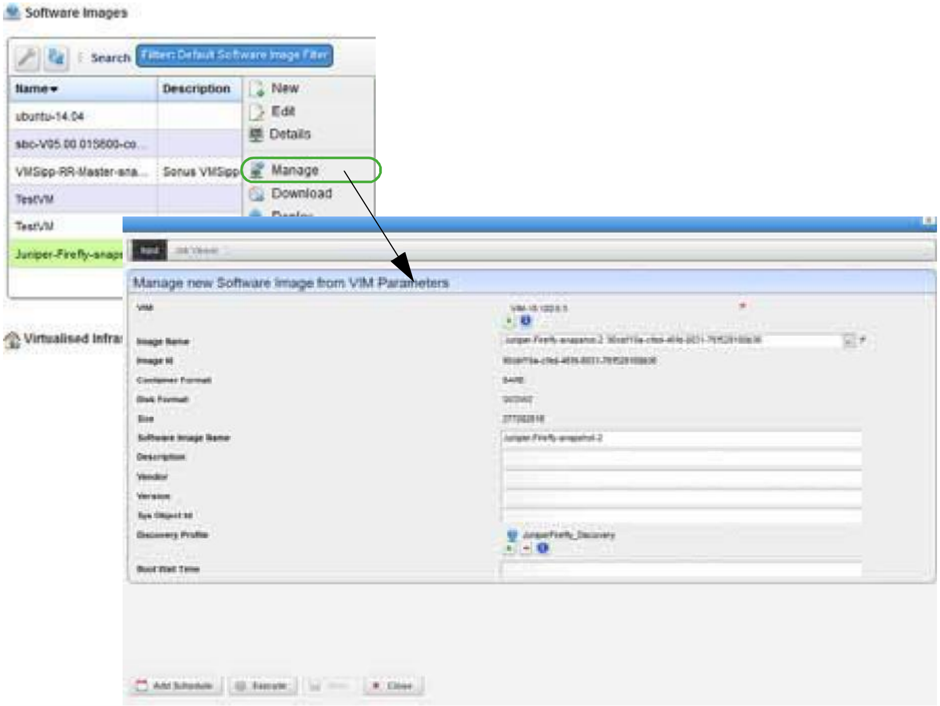

Manage new Software Image from VIM Parameters

VIM	VM-10.1[illegible]
Image Name	Juniper-Firefly-snapshot-2: [illegible]
Image Id	[illegible]
Container Format	BARE
Disk Format	QCOW2
Size	277282816
Software Image Name	Juniper-Firefly-snapshot-2
Description	
Vendor	
Version	
Sys Object Id	
Discovery Profile	JuniperFirefly_Discovery
Boot Wait Time	

4 Select the VIM that houses the snapshot image.

5 Select the snapshot image name.

For example: `Juniper-Firefly-snapshot-2:` *`uniqueImageID`*

The image ID, container format, disk format, size, and software image name fields are populated.

6 Select a discovery profile.

7 Click Execute.

A newsoftware image record is created matching the image name in your OpenStack environment.

8 Right-click an image.

9 Select Edit.

The Editing Software Image window is displayed.

10 Select the snapshot image from the Image Name list.

The Editing Software Image window is displayed.

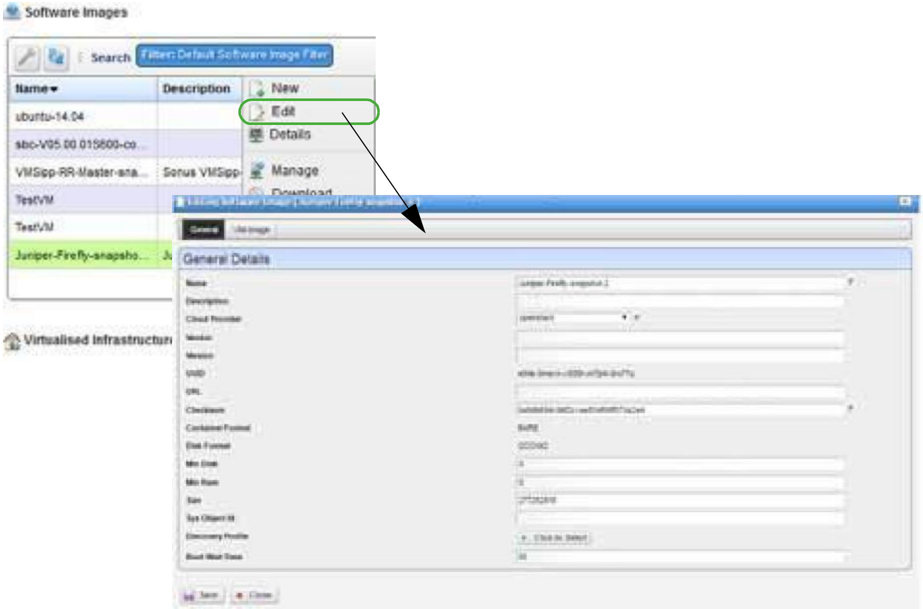

11 Enter the missing information as follows.

a. Add a meaningful description.

b. Specify the vendor, such as Juniper.

c. Add a meaningful version number, such as 12.1.

d. Enter a valid URL that points to your local environment's file server or disk location where the qcow2 file is stored and accessible. For example:

```
ftp://serverIP/firefly/Juniper-Firefly-snapshot-2.qcow2
```

e. Enter the SNMP SysObject ID for the Ju niper Firefly SRX instances, which is:

```
1.3.6.1.4.1.2636.1.1.1.2.96
```

f. Select the discovery profile containing the required ssh and SNMP authentication objects.

Discovery profiles are created by your system administrator or someone within your organization with administrative permissions.

g. Leave the Boot Wait Time field blank.

The Cruz system assigns a default integer value of 0 after you save the profile.

h. Review your inputs to ensureaccuracy.

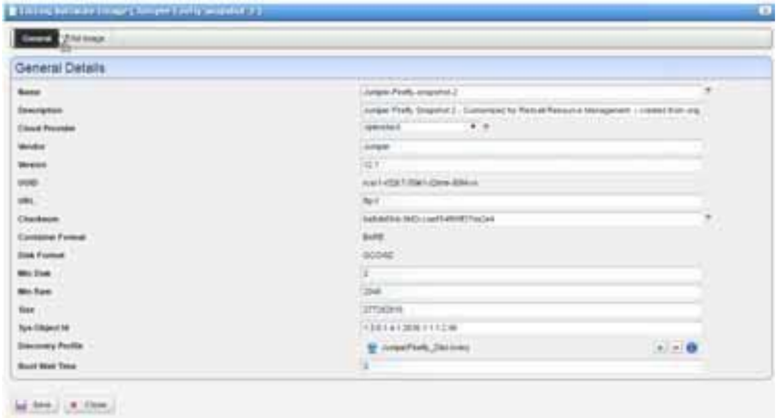

General Details

Name	Juniper-Firefly-snapshot-2
Description	Juniper Firefly Snapshot 2 - Customized by Redcell Resource Management - created from org
Cloud Provider	openstack
Vendor	Juniper
Version	12.1
UUID	[illegible]
URL	ftp://
Checksum	[illegible]
Container Format	BARE
Disk Format	QCOW2
Min Disk	2
Min Ram	2048
Size	[illegible]
Sys Object Id	[illegible]
Discovery Profile	JuniperFirefly_[illegible]
Boot Wait Time	[illegible]

12 Save the software image record.

13 Locate and Select the newly created Software Im age Record and export it to an archive disk location for backup purposes.

Understanding the Sonus VSBC VNF Package

Thissection explains the scope, functions, and processesassociated with the Virtualized Network Functionality (VNF) onboarding and instantiation:

- Understanding the Firefly VNF Package
- Firefly VNF Package Contents
- Cruz VNF Package - Install Component Files
- License Requirements
- Known Issues
- References

Sonus SBC VNFPackage

The Sonus Network Session Border Controller (SBC) Software Edition (S We) product is virtualized and fit s into a service provider's or enterprise's Network Functions Virtualization (NFV) strategy – optimizing resource utilization based on their specific needs.

The SBC SWe product, also referred to as VSBC, utilizes the same code base and firmware as the Sonus SBC 5000 and 7000 series and provides the same carrier-class redundancy to ensure service continuity.

A popular use of the SBC S We as a Virtualized Network Function (VNF) is in support of voice communications, including the use of SI P, H.323, and over 80 signaling variants providing Call Admission Control (CAC).

The VNF package provides sample VNF descriptors and associated artifacts to deploy the SBC SWe (VSBC) as a VNF in the OpenStack environment. This VNF package also includes:

- Asample standalone Sonus VSBC Cluster VNF – to demonstrate a SIP communication deployment scenario with scaling operations (elasticity) based upon call volume
- Asimpler Sonus VSBCVirtual Appliance VNF descriptor for SBC SWe standalone usage– or for insertion into the Network Service record, alongside other VNFs, that can be defined and customized based on your specific network topology and intended use

Sonus Networks provides qcow2 images and licenses that are readyfor the OpenStack environment. You need these files to deploy and enable the VSBC VNF features.

The following list provides the VNF package vendor and validation information:

Vendor and Model Information	
Vendor	Sonus Networks
Product Family	SBC
Models	VSBC (Virtual SBC)
Operating Systems	SBC SWe (Software Edition)
Version	5.01
Validated Environment	
VIM	OpenStack
Version	Liberty
Installer	Mirantis
Version	8
Network Configurations	VXLAN, VLAN Segmentation
Cruz **MANO (Management and Operation) Validated Functions**	
MANO	NSD, VNFD, NSR, VNFR Link LCEs
VNFM	Generic (MANO Lite) Independent/Specific (N-N/A)
NFVO	MANO Lite
Cruz **RMO (Resource Management and Operation) Validated Functions**	
Inventory	Resource Discovery and Inventory Resource Creation Method (NFV-Model)
Monitoring	SNMP Monitoring
Config	ACLI

Sonus SBC VNF Package Contents

TheSonus SBCVNFpackagecontainsartifactsnecessaryto installanddeployVirtualized Network Functions (VNFs) for Sonus VSBC Virtualization Deployment Units (VDUs) offered.

The following sample VNF descriptors were created with basic flavors to deploy the VNF VDUs, either as part of a standalone VNF record, or as part of a higher level, end-to-end Network Service record.

- Juniper Firefly SRX Cluster
- Juniper Firefly Virtual Appliance

Sonus VSBC VirtualAppliance

The sample Sonus VSBC Virtual Appliance descriptor is located in the following file:

`owareapps/nfv-vnfm-sonus/db/vnfd.sonus-virtual-appliance.xml`

This is a standalone Sonus VSBC that you can deploy without a Network Service.

You can copy (branch) this basic Virtualized Network Function (VNF) housing a single VSBC SWe Virtualization Deployment Units (VDU) using the Cruz VNFD portlet, Branch action and easily convert it to a non-standalone VNF to incorporate a s a VNF member of a higher-level Network Service descriptor.

The Sonus VSBC Virtual Appliance VNF descriptor definesa standalone VNF that deploysasingle VSBC VDU along with the required project networks and security groups in the OpenStack environment.

The Sonus VSBC Virtual Appliance VNF creates the following artifacts during deployment:

- Tenant Router — OpenStack Neutron Router
- Private Tenant Networks required for the VSBC to connect with to function:

Network Name	Description
MGT0	Management Network with Public (Floating) IP Address Assignment
HA	High Availability Network for Internal SBC VM Communication
PKT0	Internal Packet Network
PKT1	External Packet Network

- Security Groups controlling TCP, UDP, and IC MP port access to/from the VSBC's network ports.

Here are the **Sonus VSBC Virtual Appliance** descriptor deploy parameters used:

Field Name	Data Type	Description	Default Value	Valid Values	Require Values
Cluster Name	String	Cluster Name	ExampleCluster	Any String	N

Here are the **Sonus VSBC Virtual Appliance** descriptor flavors used:

Flavor Name	VDU Name	VDU Count	CPU	Disk (GB)	Memory (MB)
Basic	VSBC VA	1	2	65	10240

View the VNF descriptor's composition details from the Cruz Onboarding page > VNF Descriptor portlet.

Sonus VSBC Cluster

The sample Sonus VSBC Cluster descriptor is located in the following file:

`owareapps/nfv-vnfm-sonus/db/vnfd.sonus-vsbc-cluster-v1.xml`

The Sonus VSBC Cluster Virtualized Network Function (VNF) descriptor defines a standalone VNF that fully automates the deployment and monitoring of multiple types of Servers/VMs alongsideofoneortwo VSBC Virtualization Deployment Units (VDUs). This VNFisdeployedasa cluster of VDUs and is isolated to an OpenStack project.

The Sonus VSBC Cluster VNF creates the following artifacts during deployment:

- TenantRouter—OpenStack Neutron Router
- Private Tenant Networks interconnecting the VSBC with a central Dynamic Name Service (DNS) in addition to a SIP client and server VM:

Network Name	Description
MGT0	Management Network with Public (Floating) IP Address Assignment
HA	High Availability Network for Internal SBC VM Communication
PKT0	Internal Packet Network
PKT1	External Packet Network

- Security Groupscontrolling TCP,UDP,and ICMP portaccessto/from each VSBC's network ports.
- The following VDU types that get instantiated as VMs in the target OpenStack project:

VDU Type	Description
DNS	Used to perform load balancing across multiple VSBC VDU instances.
VSBC	Performs SIP Call Admission and Control for this scenario.
SIP Client	Used to originate SIP calls with the SIP Server VDU and utilized to test and demonstrate SIP call connectivity and generate call volume using the VSBCs.
SIP Server	Used to establish SIP calls with the SIP Client and utilized to test and demonstrate SIP call connectivity and generate call volume using the VSBCs.

Here are the **Sonus VSBC Cluster** descriptor deploy parameters used:

Field Name	Data Type	Description	Default Value	Valid Values	Required Values
Cluster Name	String	Cluster Name	ExampleCluster	Any String	N
DNS IP	String	IP address of DNS Server	8.8.8.8	Valid IP Address	N
DNS Zone	String	DNS zone	example.com	Any String	N
VSBC DNS Host Name	String	Host Name in the Domain to use for the Updated DNS Entry	vsbc	Any String	N

Here are the **Sonus VSBC Cluster** descriptor flavors used:

Flavor Name	VDU Name	VDU Count	CPU	Disk (GB)	Memory (MB)
Basic	C-VSBC	1	2	65	10240
	Ubuntu DNS	1	1	10	512
	SIPP Server	1	1	20	2048
	SIPP Client	1	1	20	2048
Advanced	C-VSBC	2	4	130	20480
	Ubuntu DNS	1	1	10	512
	SIPP Server	1	1	20	2048
	SIPP Client	1	1	20	2048

View the VNF descriptor’s composition details from the Cruz Onboarding page > VNF Descriptor portlet.

Cruz VNF Package - Install Component Files

The following Cruz installation components are required for the Sonus VSBC product:

VNF (OCP)	nfv-vnfm-sonus.ocp
Device Driver (DDP)	ddsonus.ddp

The Cruz VNF package descriptors for the **Sonus VSBC** VNF product include:

NVF VNF Package			Resource Management		
VIM Software Image Records	VNF Descriptors	Sample NSD	Discovery Authentication Records	Discovery Profiles	Device Driver
Y	Y	N	Y	Y	Y

The VNF package descriptors include:

- Discovery Authentication Profiles
- Discovery Profiles
- NFV VIM Software Image Descriptors

Discovery Authentication Profiles

The Sonus VSBC Discovery Authentication profiles are located in the following file:

`owareapps/nfv-vnfm-sonus/db/rc.disc-auths-vsbc.xml`

This file contains the following discovery authentication profiles:

Authentication Record Name	Description
sonus-cli	SSH v2 Credentials
snmp-sonus-vsbc	SNMP v2 Credentials

Discovery Profiles

The Sonus VSBC Discovery profiles are located in the following file:

`owareapps/nfv-vnfm-sonus/db/rc.disc-profile-vsbc-v50x.xml`

This file contains the Sonus VSBC Discovery profile (Sonus_VSBC_v5.0.x) that automatically discovers and populates resource records during Virtualized Network Function (VNF) record instantiation.

The following tasks are run once the target VM in the OpenStack environment is fully booted and can respond to SNMP and command-line interface (CLI) requests from the management system:

1. Sonus SBC Reset CLI Password
2. Resync
3. Refresh Monitor Targets
4. Sonus VSBC Set Account Max Session
5. Sonus vSBC SNMP Trap Forwarding to Cruz

NFV VIM Software Image Descriptors

Here is alistofthe Network Functions Virtualization (NFV) Virtualized Infrastructure Manager (VIM) software image descriptors.

Name	Scope/Description	Container Format	Disk Format	Min RAM (MB)	Min Disk (GB)	Resource Discovery Profile
sbc-V05.00.01S600-connexip-os_03.00.02-S600_amd64-GA	Base image for the SBC SWe (VSBC) - version 5.0.1 GA Release. This is the original qcow2 file supplied by Sonus.	BARE	QCOW 2	10240	65	Y
ubuntu-14.04	Open Source version of the Ubuntu OS utilized to create the DNS VDU in the Sonus VSBC Cluster VNF.	BARE	QCOW 2	0	0	N
VMSipp-RR-Master-snapshot-1	Open Linux Image with cloud-in it and SIP testing software installed. Utilized for the SIP Client and Server VDUs in the Sonus VSBC Cluster VNF.	BARE	QCOW 2	0	20	N

License Requirements

Sonus Networks does require and provide licenses to operate an SBC SWe (VSBC) as a fully functional Virtual Machine in the OpenStack environment.

Contact the Sonus Networks support team regarding license requirements and logistics.

Known Issues

The Sonus VSBC instances require that the CLI user account's password change during the initial login session. A task was implemented in the Sonus Device Driver so that it applies a password change via CLI communication with the VSBC during the initial SSH session established to a newly deployed VSBC Virtualized Network Function's (VNF's) Virtualization Deployment Unit (VDU).

Contact the Sonus Networks support team for further details regarding known issues or caveats utilizing the SBC SWe VNF.

References

This section provides references to:

- Vendor Documentation (URLs)
- VNF Configuration Methods

Vendor Documentation (URLs)

Refer to the Sonus Networks website (http://ww w.sonus.net) for all vendor-specific documentation.

For the Sonus VSBC documentation and release notes, refer to the following Sonus Networks websites page:

`http://www.sonus.net/products/session-border-controllers/sbc-swe`

VNF Configuration Methods

A combination of image (snapshot) preparation, along with Virtualization Deployment Unit (VDU) Post Instantiate Configuration Lifecycle Events (LCEs), is implemented in this Virtualized Network Function (VNF) package to ensure that the Sonus VSBC VDUs are automatically brought under management.

Customized Image Preparation

For the VSBC VDUs:

- The original qcow2 image was used to create this VNF package, without a need to customize or take a snapshot for the Sonus VSBC VNF VDUs.
- Configuration changes are applied automatically through discovery actions to the VSBC running in the OpenStack system post boot-up in order to have Cruz monitor, configure and manage the VSBC.

For the SIP Client and Server VDUs:

- Cloud-init was added by first using the original image, VMSipp-RR-Master, and then installing cloud-init on a running instance boot ed from the original image. A snapshot image was created with cloud-init included.
- The SIP Client and Server VDUs are then configured with instance specific IP address information using the OpenStack HEAT Server Resource metadata facility in conjunction with cloud-init during execution of the VNF VDU's instantiate lifecycle task.

For the DNS VDU:

- A cloud enabled Ubuntu version 14.04 image is used that has cloud-init included in the original image.
- The DNS service is installed and configured using the OpenStack HEAT Server Resource metadata facility in conjunction with cloud-init during execution of the VNF VDU's instantiate Lifecycle task.

Cruz Resource Configuration Tasks

A specialized action, Sonus SBC Rest CLI Password, is specified as one of the execution actions in the sample Resource Discovery profile included in this VNF package. This action is required to perform the initial password reset required by the SBC when logging in the very first time to a new VSBC VM instance using the CLI user account.

Thepassword reset action is automaticallyexecuted, prior to performing resource discovery/resync and post configuration tasks.

Once the target VM in the OpenStack environment is fully booted and can respond to SNMP and CLI requests fromthe management system, post discoverytasks are run. For more details about the Sonus VSBC profile and the tasks that run, see Discovery Profiles on page 782.

Managing NFVI/VIM Resources

The management of Network Functions Virtualization infrastructure (NFVI) orvirtualized infrastructure manager (VIM) resources from the application is no different from managing physical devices you discovered or created on your network. This section provides some basic instructions to manage NFVI/VIM resources. For a detailed description of the user interface, instructions, and other resource management options available, Resource Management Portlets and Editors on page 160.

Monitor VIM details, alarms, and history to determine what actions are required, such as VNF discovery, VIM resync, VIM delete, and so on.

- Managing VIMs
- Discovering VNF Records
- Managing Software Images
- Maintaining Resource Monitors
- Managing SDN Controllers

Managing VIMs

AVirtualized Infrastructure Managersportlet is where youdeploynetworkservices, physical network functions (PNFs) and virtualized network functions (VNFs).

This section provides example instructions for:

- Creating VIMs
- Modify a VIM
- Deleting a VIM
- Resyncing a VIM

Creating VIMs

Create VIMs from the Virtualized Infrastructure Manager portlet as follows.

1 Right-click > New.

The Editing VIM window is displayed.

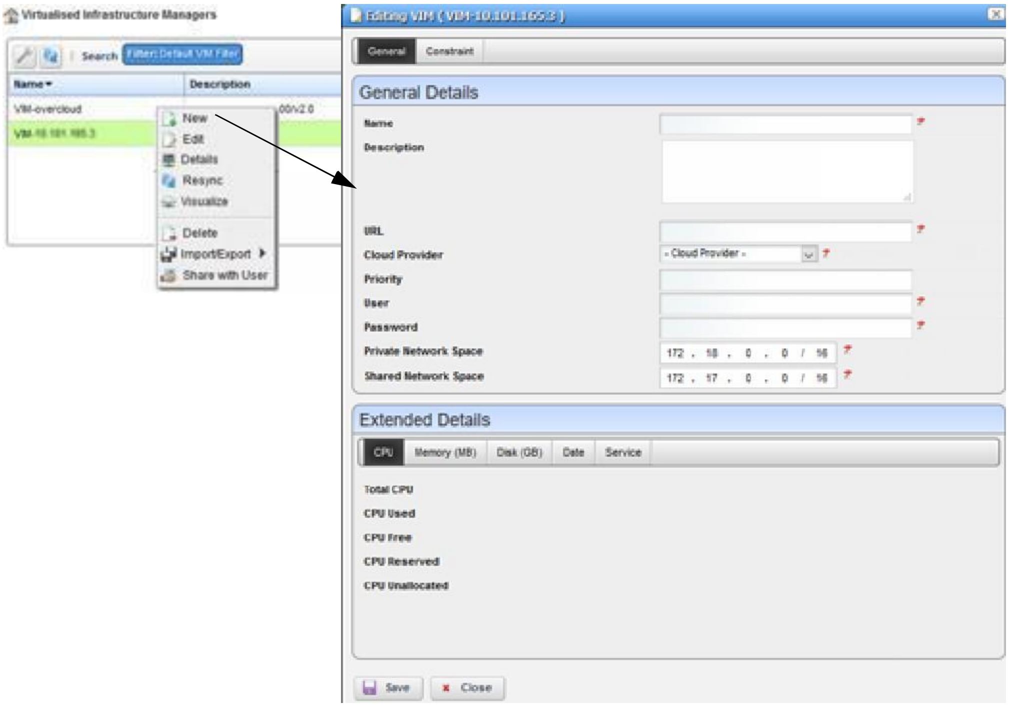

Private Network Space 172 . 18 . 0 . 0 / 16

Shared Network Space 172 . 17 . 0 . 0 / 16

2 Enter the required information.

Optionally, provide a more detailed description and apriority.

3 Specify any constraints.

Constraints specify where to place a network service or VNF if you do not specify a VIM during the staging process.

4 Save the VIM.

This brings the VIM under management.

Once you create the VIM, it is a good idea to do a resync to ensure that the VIM has the latest capacity information.

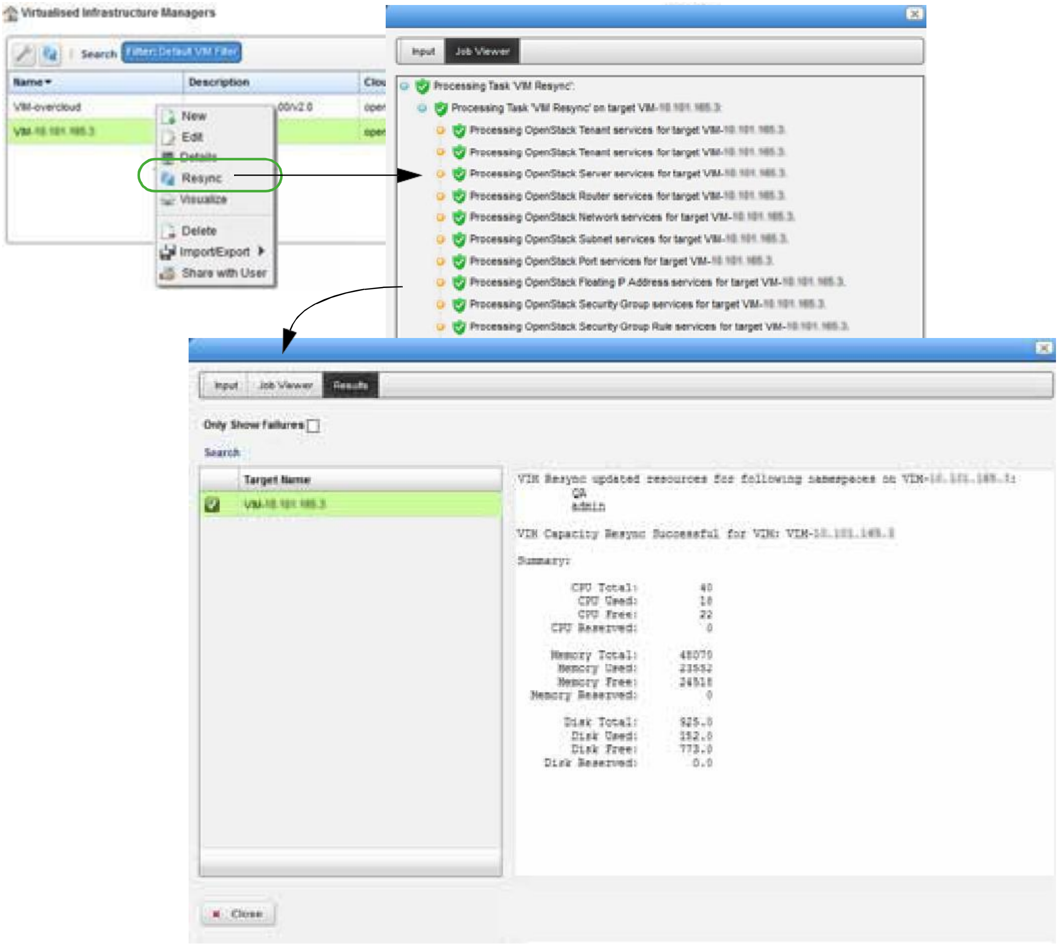

```
Processing Task 'VM Resync':
  Processing Task 'VM Resync' on target VM-[redacted]:
    Processing OpenStack Tenant services for target VM-[redacted].
    Processing OpenStack Tenant services for target VM-[redacted].
    Processing OpenStack Server services for target VM-[redacted].
    Processing OpenStack Router services for target VM-[redacted].
    Processing OpenStack Network services for target VM-[redacted].
    Processing OpenStack Subnet services for target VM-[redacted].
    Processing OpenStack Port services for target VM-[redacted].
    Processing OpenStack Floating P Address services for target VM-[redacted].
    Processing OpenStack Security Group services for target VM-[redacted].
    Processing OpenStack Security Group Rule services for target VM-[redacted].
```

```
VIM Resync updated resources for following namespaces on VIM-[redacted]:
        QA
        admin

VIM Capacity Resync Successful for VIM: VIM-[redacted]

Summary:

        CPU Total:          40
         CPU Used:          18
         CPU Free:          22
     CPU Reserved:           0

     Memory Total:       48070
      Memory Used:       23552
      Memory Free:       24518
  Memory Reserved:           0

       Disk Total:       925.0
        Disk Used:       152.0
        Disk Free:       773.0
    Disk Reserved:         0.0
```

5 View VIM details to see everything related to the VIM, such as:

- Capacity resources (CPU, memory, disk)
- Defined constraints
- Hypervisors
- VIM Images
- Reference details (IP addresses, network, port, router, capacity, and so on)

6 Repeat steps 1 through 6 for each new VIM.

Modify a VIM

Modify a VIM from the Virtualized Infrastructure Manager portlet as follows.

1 Select a VIM.

2 Right-click > Edit.

The Editing VIM window is displayed.

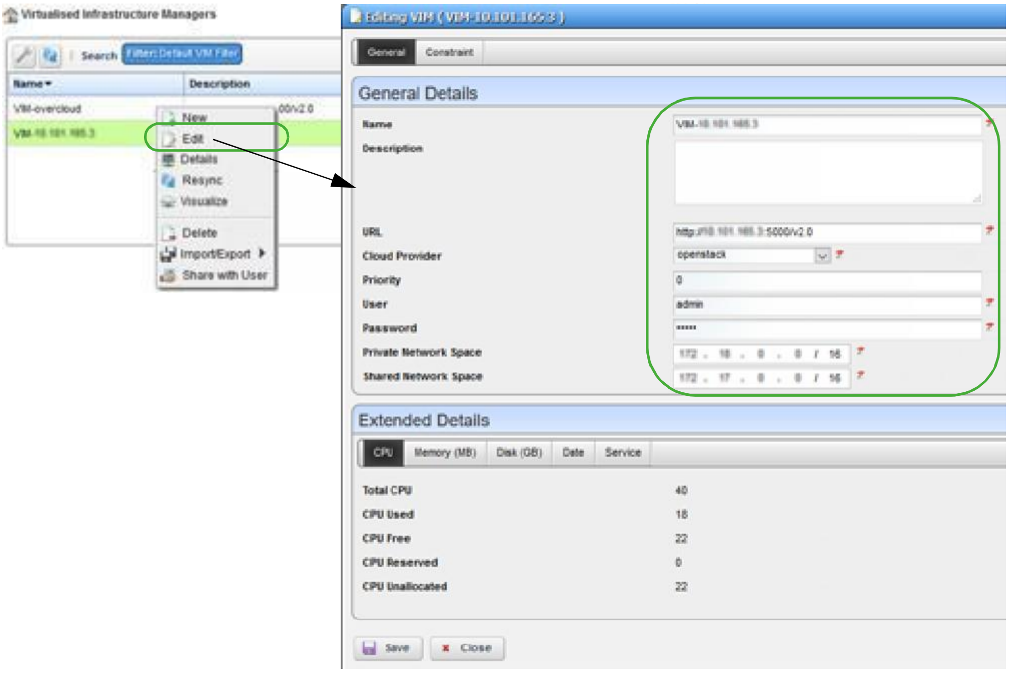

Name: VIM-10.101.165.3

Description:

URL: http://10.101.165.3:5000/v2.0

Cloud Provider: openstack

Priority: 0

User: admin

Password: *****

Private Network Space: 172 . 16 . 0 . 0 / 16

Shared Network Space: 172 . 17 . 0 . 0 / 16

Total CPU	40
CPU Used	18
CPU Free	22
CPU Reserved	0
CPU Unallocated	22

3 Modify the general information as needed.

4 Modify the constraints as needed.

5 Save your changes.

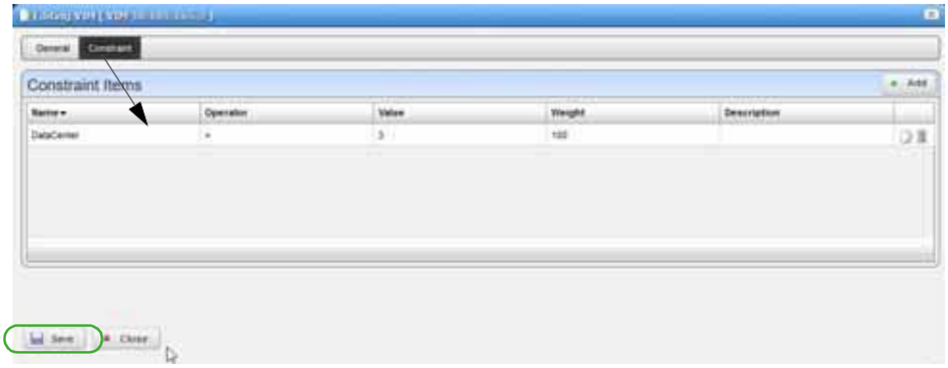

Name	Operation	Value	Weight	Description
DataCenter	=	3	100	

Deleting a VIM

Delete a VIM from the Virtualized Infrastructure Manager portlet as follows.

1 Select a VIM.

2 Right-click > Delete.

A confirmation message is displayed.

3 Click Delete.

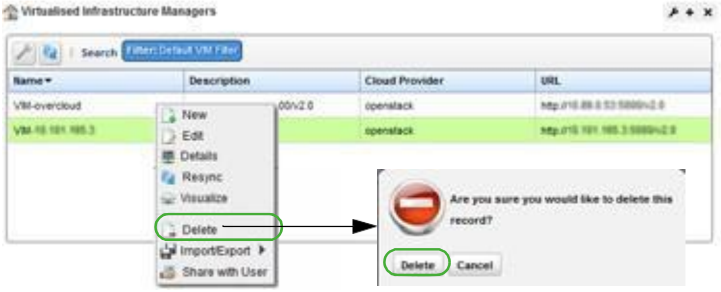

Virtualised Infrastructure Managers

Name	Description	Cloud Provider	URL
VIM-overcloud	...00/v2.0	openstack	[redacted]
[redacted]		openstack	[redacted]

Are you sure you would like to delete this record?

4 Verify that the VIM no longer exists.

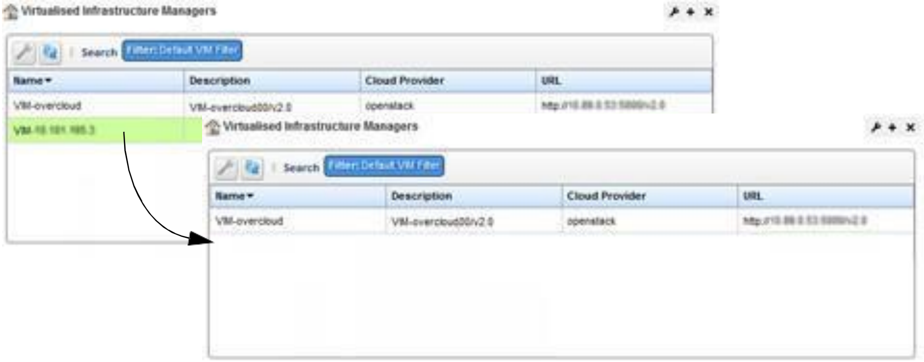

Virtualised Infrastructure Managers

Name	Description	Cloud Provider	URL
VIM-overcloud	VIM-overcloud00/v2.0	openstack	[redacted]
[redacted]			

Virtualised Infrastructure Managers

Name	Description	Cloud Provider	URL
VIM-overcloud	VIM-overcloud00/v2.0	openstack	[redacted]

Resyncing a VIM

Theresyncaction refreshes the selected virtualized infrastructure manager (VIM) with the latest capacity information.

Resync a VIM from the Virtualized Infrastructure Manager portlet byright clicking a VIM and then selecting Resync, or from the Actions portlet asfollows.

1 Navigate to the VIM Resync action.

2 Right-click VIMResync > Execute.

The Executing Action (VIM Resync) window displays a list of VIM entities.

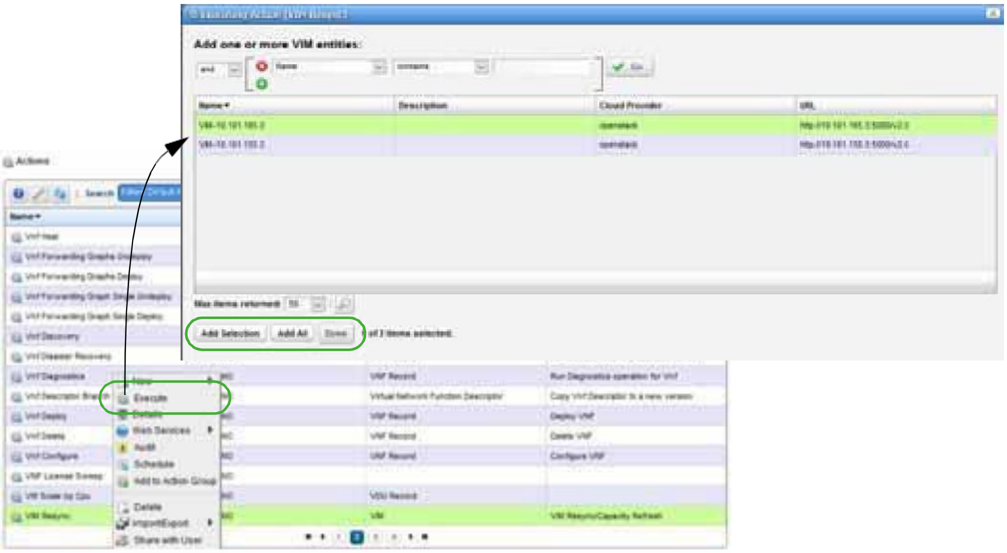

3 Select a VIM.

4 Click Add Selection and then Done.

The Executing Action (VIM Resync) Info window is displayed.

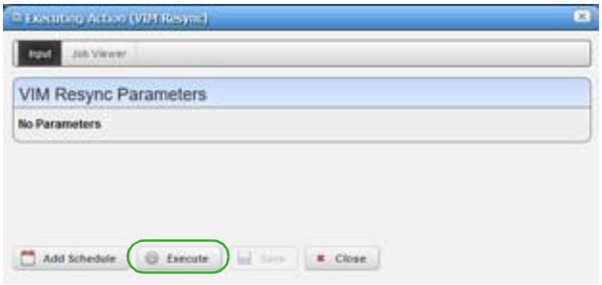

5 Click Execute.

The Executing Action (VIM Resync) Job Viewer displaystheprogressand uponcompletion, the results are displayed.

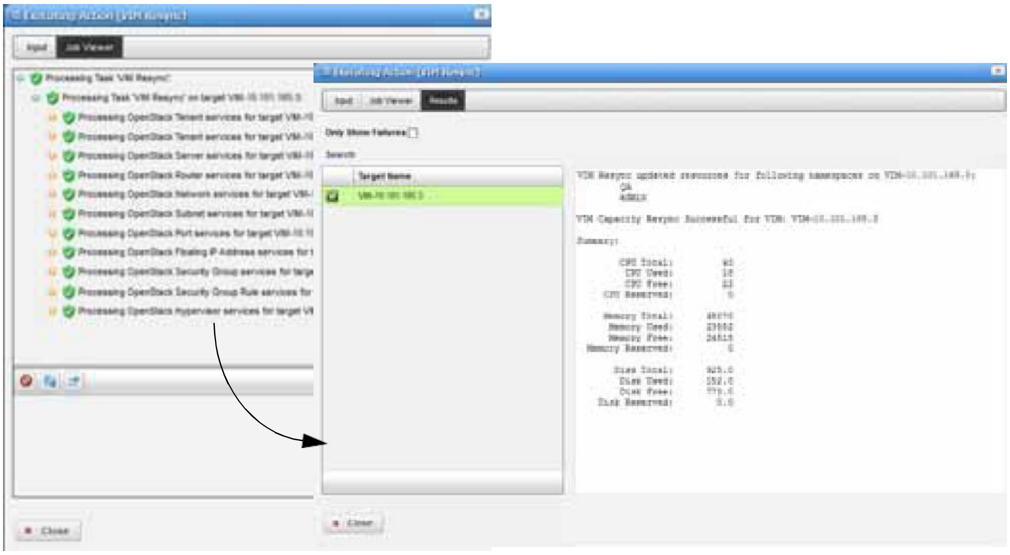

Processing Task 'VIM Resync'
Processing Task 'VIM Resync' on target VIM-[illegible]
Processing OpenStack Tenant services for target VIM-[illegible]
Processing OpenStack Tenant services for target VIM-[illegible]
Processing OpenStack Server services for target VIM-[illegible]
Processing OpenStack Router services for target VIM-[illegible]
Processing OpenStack Network services for target VIM-[illegible]
Processing OpenStack Subnet services for target VIM-[illegible]
Processing OpenStack Port services for target VIM-[illegible]
Processing OpenStack Floating IP Address services for t[illegible]
Processing OpenStack Security Group services for targe[illegible]
Processing OpenStack Security Group Rule services for [illegible]
Processing OpenStack Hypervisor services for target VI[illegible]

```
VIM Resync updated resources for following namespaces on VIM-[illegible]:
        QA
        ADMIN

VIM Capacity Resync Successful for VIM: VIM-[illegible]

Summary:

     CPU Total:        40
      CPU Used:        18
      CPU Free:        22
  CPU Reserved:         0

  Memory Total:     48670
   Memory Used:     23552
   Memory Free:     24818
Memory Reserved:        0

    Disk Total:     925.0
     Disk Used:     152.0
     Disk Free:     773.0
 Disk Reserved:       0.0
```

Discovering VNF Records

The discovery VNF action finds VNF records existing on the specified VIM. Youcan discover VNF records from the Virtual Network Function Records portlet or the Actions portlet. The instructions provided in this section show how to discover VNFs from the Actions portlet.

Discover VNF records from the Actions portlet as follows.

1 Navigate to the Vnf Discovery action.

2 Right-click Vnf Discovery> Execute.

The VIM entities list is displayed.

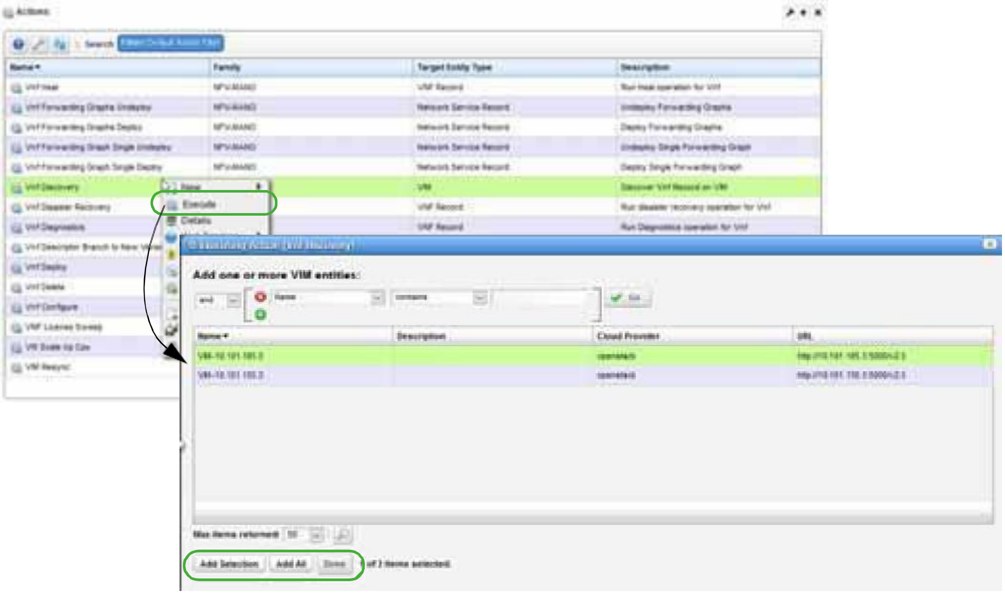

Name	Family	Target Entity Type	Description
Vnf Heal	NFV-MANO	VNF Record	Run Heal operation for Vnf
Vnf Forwarding Graphs Undeploy	NFV-MANO	Network Service Record	Undeploy Forwarding Graphs
Vnf Forwarding Graphs Deploy	NFV-MANO	Network Service Record	Deploy Forwarding Graphs
Vnf Forwarding Graph Single Undeploy	NFV-MANO	Network Service Record	Undeploy Single Forwarding Graph
Vnf Forwarding Graph Single Deploy	NFV-MANO	Network Service Record	Deploy Single Forwarding Graph
Vnf Discovery		VIM	Discover Vnf Record on VIM
Vnf Disaster Recovery		VNF Record	Run disaster recovery operation for Vnf
Vnf Diagnostics		VNF Record	Run Diagnostics operation for Vnf
Vnf Descriptor Branch to New Ve[illegible]			
Vnf Deploy			
Vnf Delete			
Vnf Configure			
VNF License Sweep			
VM Scale Up Cpu			
VIM Resync			

Name	Description	Cloud Provider	URL
VIM-10.101.185.3		openstack	http://10.101.185.3:5000/v2.0
VIM-10.101.150.2		openstack	http://10.101.150.2:5000/v2.0

3 Select a VIM

4 Click Add Selection and the Done.

The Vnf Discovery Parameters are displayed.

5 Specify the namespaces.

6 Click Execute.

The Executing Action(Vnf Discovery) Job Viewerdisplaystheprogressanduponcompletion, the results are displayed.

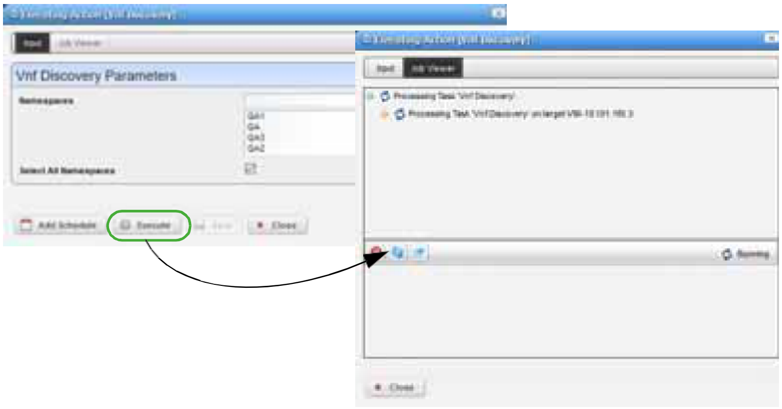

Processing Task 'Vnf Discovery'
Processing Task 'VnfDiscovery' on target VM-10.131.165.3

7 Verify that the target names listed are found in the Virtual Network Function Records portlet.

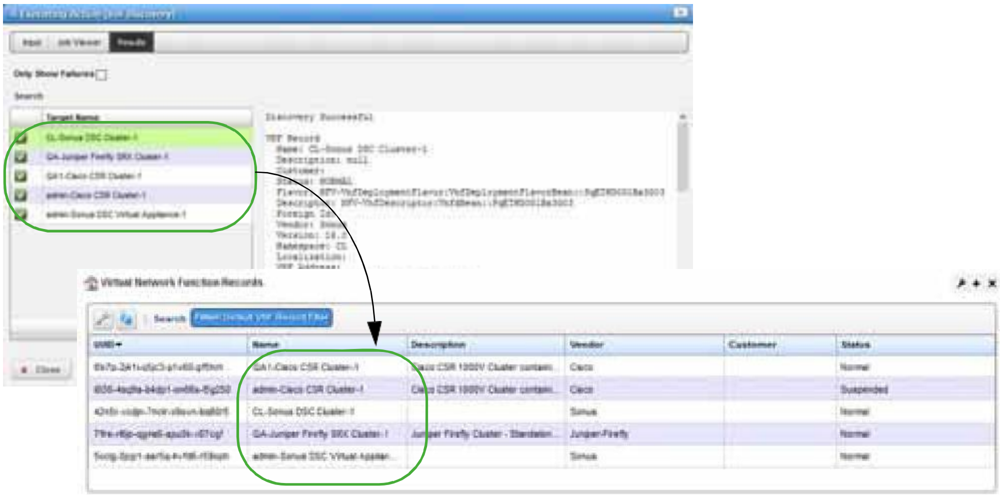
Executing Action (Vnf Discovery)
Input
Job Viewer
Results
Only Show Failures
Search
Close
Virtual Network Function Records
Search
Filter: Default VNF Record Filter

	Target Name
☑	CL-Sonus DSC Cluster-1
☑	QA-Juniper Firefly SRX Cluster-1
☑	QA1-Cisco CSR Cluster-1
☑	admin-Cisco CSR Cluster-1
☑	admin-Sonus DSC Virtual Appliance-1

```
Discovery Successful

VNF Record
  Name: CL-Sonus DSC Cluster-1
  Description: null
  Customer:
  Status: NORMAL
  Flavor: NFV-VnfDeploymentFlavor:VnfDeploymentFlavorBean::[illegible]
  Descriptor: NFV-VnfDescriptor:VnfdBean::[illegible]
  Foreign Id:
  Vendor: Sonus
  Version: 18.2
  Namespace: CL
  Localization:
  VNF Address:
```

UUID	Name	Description	Vendor	Customer	Status
[illegible]	QA1-Cisco CSR Cluster-1	Cisco CSR 1000V Cluster contain...	Cisco		Normal
[illegible]	admin-Cisco CSR Cluster-1	Cisco CSR 1000V Cluster contain...	Cisco		Suspended
[illegible]	CL-Sonus DSC Cluster-1		Sonus		Normal
[illegible]	QA-Juniper Firefly SRX Cluster-1	Juniper Firefly Cluster - Standalon...	Juniper-Firefly		Normal
[illegible]	admin-Sonus DSC Virtual Applian...		Sonus		Normal

Managing Software Images

A software image represents an image that can be deployed to a virtualized infrastructure manager (VIM) and then turned up (activated and provisioned) foraspecificvirtualization deployment unit (VDU) within avirtualized networkfunction (VNF). Onceasoftwareimage exists, youcan modify, deploy, undeploy, view details, and so on.

This section provides example instructions for:

- Creating a Software Image
- Deploying a Software Image

Creating a SoftwareImage

You have the option to create a software image from the Software Images portlet using either the New or Manage options. For simplicity, the example provided here uses the New option.

Create a software image from the Software Images portlet as follows.

1 Right-click > New.

The Creating new Software Image window is displayed.

2 Enter the required information and any other optional information.

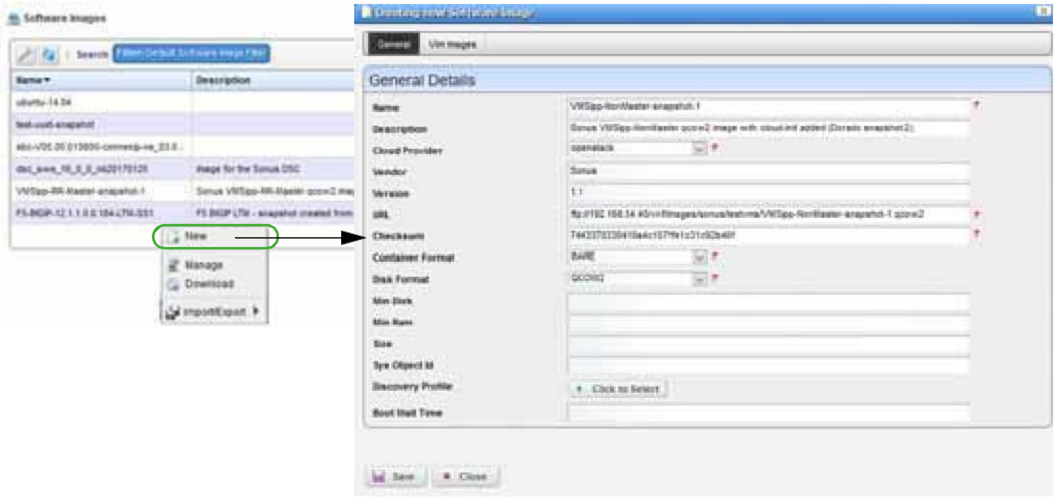

Name	Description
ubuntu-14.04	
test-cust-snapshot	
sbc-V05.00.01S800-connexip-os_03.0...	
dsc_swe_16_0_0_b420170125	image for the Sonus DSC
VWSpp-RR-Master-snapshot-1	Sonus VWSpp-RR-Master qcow2 ima...
F5-BIGIP-12.1.1.0.0.184-LTM-SS1	F5 BIGIP LTM - snapshot created from...

Name	VWSpp-NonMaster-snapshot-1
Description	Sonus VWSpp-NonMaster qcow2 image with cloud-init added (Dorado snapshot2)
Cloud Provider	openstack
Vendor	Sonus
Version	1.1
URL	ftp://192.168.34.40/nfvimages/sonus/[illegible]/VWSpp-NonMaster-snapshot-1.qcow2
Checksum	[illegible]
Container Format	BARE
Disk Format	QCOW2
Min Disk	
Min Ram	
Size	
Sys Object Id	
Discovery Profile	Click to Select
Boot Wait Time	

3 Click Save.

The Audit Trail Viewer shows that an inventory item was created for the image and the Software Images portlet displays the new image.

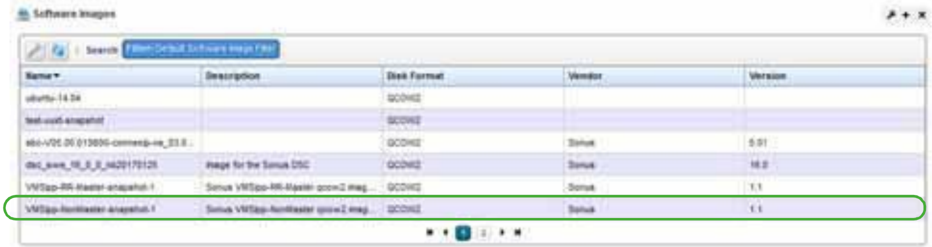

Name	Description	Disk Format	Vendor	Version
ubuntu-14.04		QCOW2		
test-cust-snapshot		QCOW2		
sbc-V05.00.01S800-connexip-os_03.0...		QCOW2	Sonus	5.01
dsc_swe_16_0_0_b420170125	image for the Sonus DSC	QCOW2	Sonus	16.0
VWSpp-RR-Master-snapshot-1	Sonus VWSpp-RR-Master qcow2 imag...	QCOW2	Sonus	1.1
VWSpp-NonMaster-snapshot-1	Sonus VWSpp-NonMaster qcow2 imag...	QCOW2	Sonus	1.1

Before the image is available for use, you need to deploy it to one or more VIMs.

Deploying a SoftwareImage

Successfully deploying a software image makes the image available from the selected virtualized infrastructure managers (VIMs), whichallows you to activate and provisionspecific virtualization deployment unit (VDU) within avirtualized network function (VNF). Deployed imagesare added to the VIM Images list.

Deploy a software image from the Software Images portlet as follows.

1 Select an image.

2 Right-click > Deploy.

The Deploy Image to VIMParametersare displayed.

3 Select the VIMs to which you want to deploy the image.

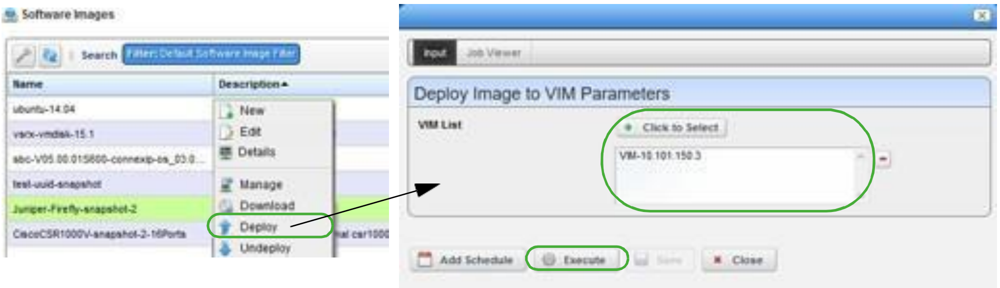

4 Click Execute.

The Job Viewer displays the processing progress.

5 Wait until the deploy process successfully completed.

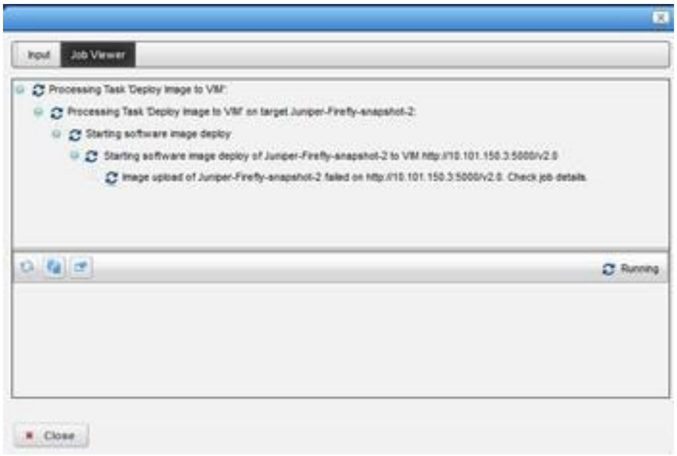

```
Processing Task 'Deploy Image to VIM':
  Processing Task 'Deploy Image to VIM' on target Juniper-Firefly-snapshot-2:
    Starting software image deploy
      Starting software image deploy of Juniper-Firefly-snapshot-2 to VIM http://10.101.150.3:5000/v2.0
        Image upload of Juniper-Firefly-snapshot-2 failed on http://10.101.150.3:5000/v2.0. Check job details.
```

6 Verify that the image was added to the VIM Images list.

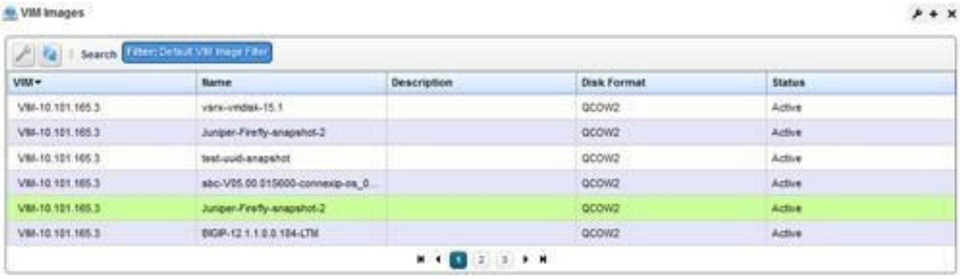

VIM	Name	Description	Disk Format	Status
VIM-10.101.165.3	vsrx-vmdisk-15.1		QCOW2	Active
VIM-10.101.165.3	Juniper-Firefly-snapshot-2		QCOW2	Active
VIM-10.101.165.3	test-uuid-snapshot		QCOW2	Active
VIM-10.101.165.3	sbc-V05.00.015600-connexip-os_0...		QCOW2	Active
VIM-10.101.165.3	Juniper-Firefly-snapshot-2		QCOW2	Active
VIM-10.101.165.3	BIGIP-12.1.1.0.0.184-LTM		QCOW2	Active

Maintaining Resource Monitors

The options available and the steps you take to create a resource monitor varies depending on the type of monitor you choose to create. The example steps provided here are for creating, modifying, and deleting a key metrics monitor.

In addition to these basic steps, there are many other settings to set up for monitoring, such as calculated metrics, thresholds, inventory mappings, and conditions. See Resource Monitors on page 361 for a detailed description of the Resource Monitors portlet, features, and options.

Creating a Monitor

Create a Key Metrics Monitor from the Resource Monitors portlet as follows.

1 Right-click > New Monitor > Key Metrics.

The Creating New Monitor window displays the General parameters.

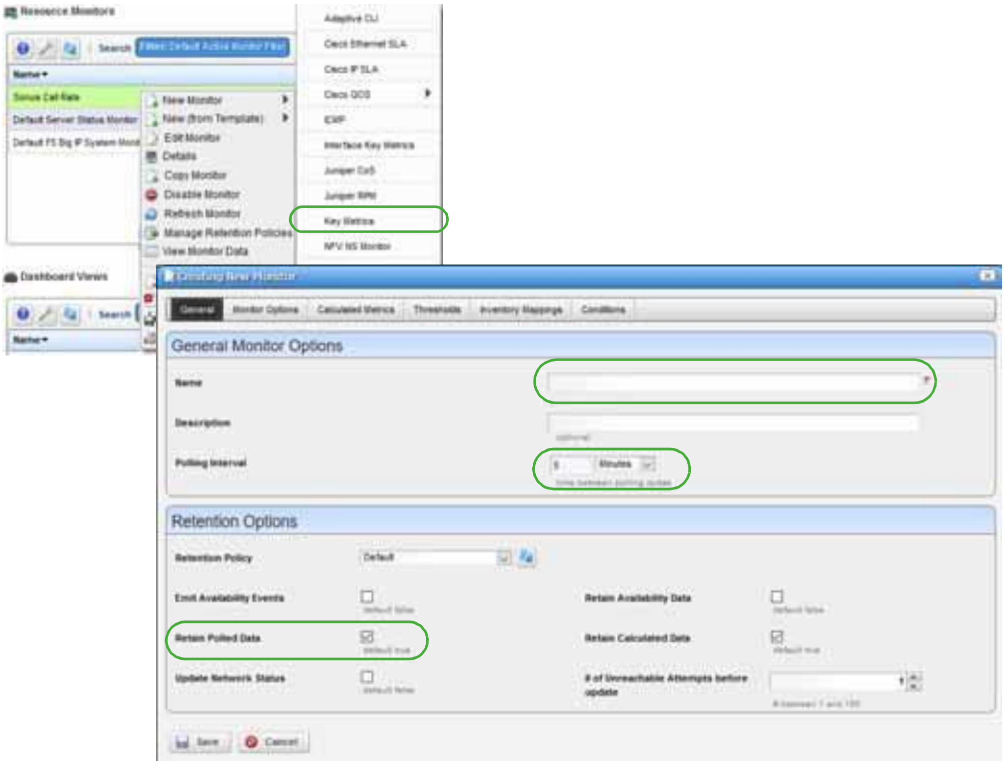

General Monitor Options

Name

Description — optional

Polling Interval: 5 Minutes — time between polling cycles

Retention Options

Retention Policy: Default

Emit Availability Events ☐ default false

Retain Availability Data ☐ default false

Retain Polled Data ☑ default true

Retain Calculated Data ☑ default true

Update Network Status ☐ default false

# of Unreachable Attempts before update — # between 1 and 100

2 Enter aname.

For example, enter Test Key Metrics Monitor.

3 Select a polling interval.

The default is 5 minutes.

4 Make sure that Retain Polled Data is selected.

5 Accept all other default retention options.

6 Set monitor options.

a. Select Monitor Options.

b. Click Add for Monitor Entities.

An entities list is displayed.

c. Select devices on which you want to apply key metrics.

d. Click Add Selection.

e. Click Done.

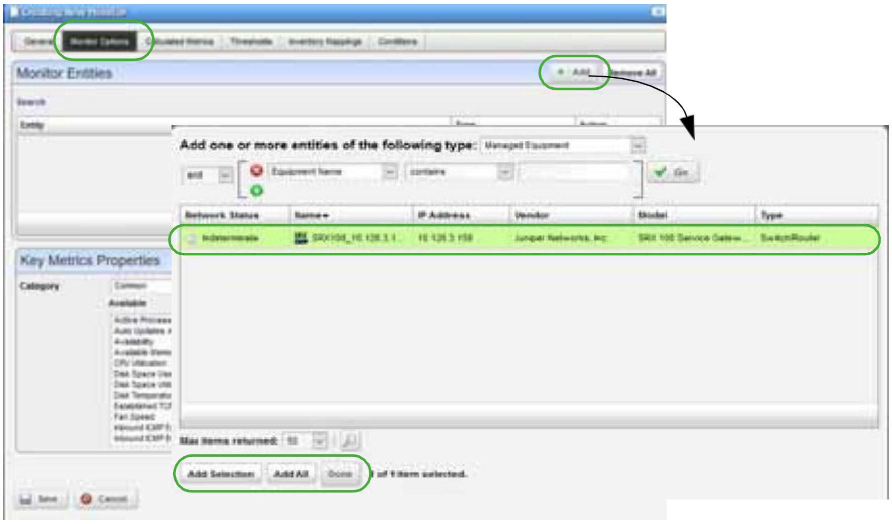

f. Select key metrics Category.

g. Move the appropriate available metrics to the Selected list.

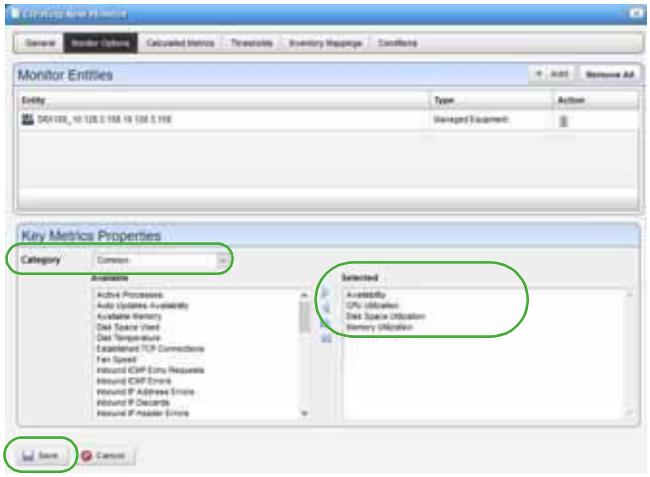

Monitor Entities

Entity	Type	Action
[illegible]	Managed Equipment	

Key Metrics Properties

Category: Common

Available:
Active Processes
Auto Updates Availability
Available Memory
Disk Space Used
Disk Temperature
Established TCP Connections
Fan Speed
Inbound ICMP Echo Requests
Inbound ICMP Errors
Inbound IP Address Errors
Inbound IP Discards
Inbound IP Header Errors

Selected:
Availability
CPU Utilization
Disk Space Utilization
Memory Utilization

7 Save your new monitor.

8 Verify that the Resource Monitors portlet lists your new monitor.

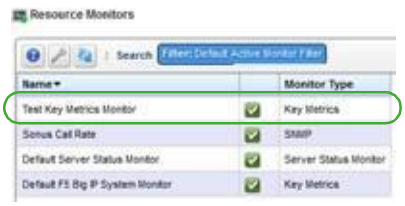

Name		Monitor Type
Test Key Metrics Monitor	☑	Key Metrics
Sonus Call Rate	☑	SNMP
Default Server Status Monitor	☑	Server Status Monitor
Default F5 Big IP System Monitor	☑	Key Metrics

Modifying a Monitor

1 Select the resource monitor you want to modify.

For example, the Test Key Metrics Monitor created earlier.

2 Right-click > Edit Monitor.

The Editor Monitor window displays the General parameters.

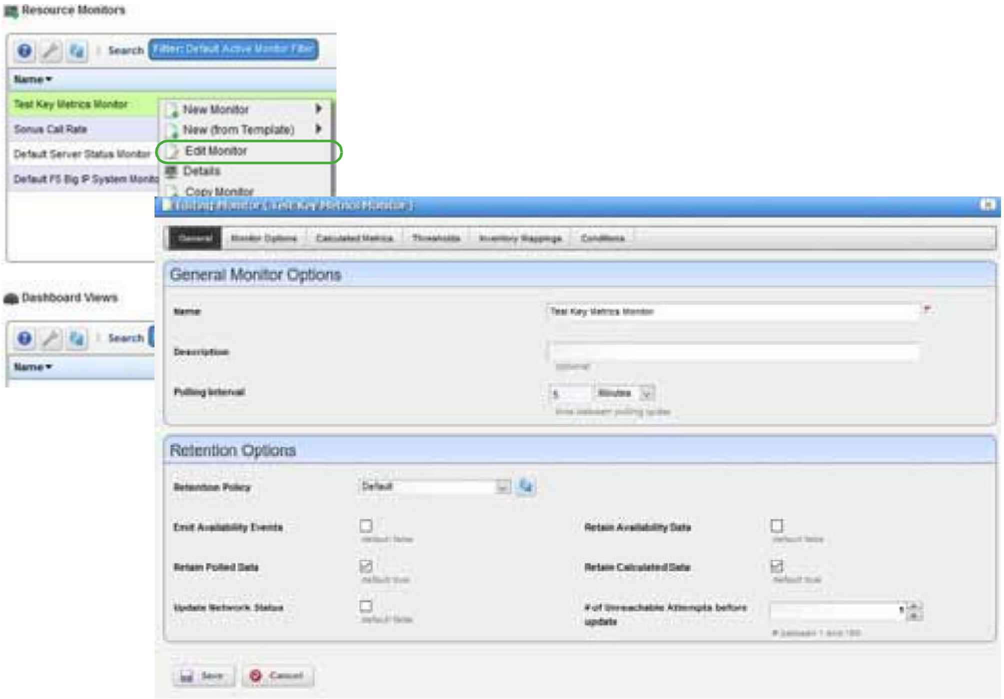

Name: Test Key Metrics Monitor

Description:

Polling Interval: 5 Minutes

Retention Policy: Default

Emit Availability Events: □

Retain Availability Data: □

Retain Polled Data: ☑

Retain Calculated Data: ☑

Update Network Status: □

# of Unreachable Attempts before update:

3 Make any needed changes to:

- Polling Interval or Retention Options
- Monitor Options
- Calculated Metrics
- Thresholds
- Inventory Mapping
- Conditions

4 Save your changes.

Deleting a Monitor

Delete a resource monitor from the Resource Monitors portlet as follows.

1 Select the resource monitor you want to delete.

For example, the Test Key Metrics Monitor created earlier.

2 Right-click > Delete.

A confirmation message is displayed.

3 Click Delete.

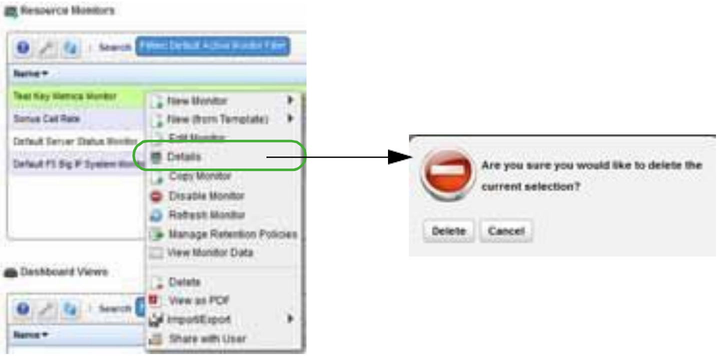

Are you sure you would like to delete the current selection?

4 Verify that the Resource Monitors portlet no longer lists the deleted monitor.

Notice that the Test Key Metrics Monitor created earlier no longer exists.

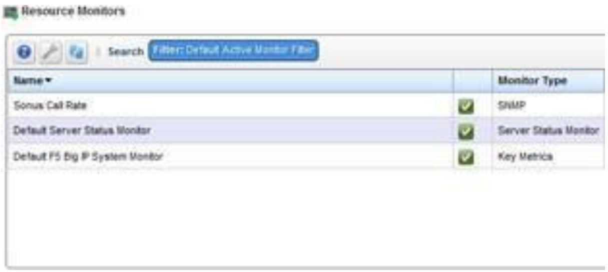

Name		Monitor Type
Sonus Call Rate	☑	SNMP
Default Server Status Monitor	☑	Server Status Monitor
Default F5 Big IP System Monitor	☑	Key Metrics

Managing SDN Controllers

A software-defined networking (SDN) controller is an application that manages flow control to enable intelligent networking. SDN controllers are based on protocols, such as OpenFlow, that allow servers to tell switches where to send packets.

In a conventional switch, packet forwarding (the data path) and high-level routing (the control path) occurs on the same device. An OpenFlow switch separates the data path from the control path.

This section provides example instructions for:

- Creating SDN Controllers
- Modifying an SDN Controller
- Deleting an SDN Controller

Creating SDN Controllers

Create SDN controllers from the SDN Controllers portlet as follows.

1 Right-click > New.

The Editing SDN Controller window is displayed.

2 Enter the required information.

3 Enter any constraints.

4 Click Save.

The SDN Controllers portlet now lists the new controller.

5 Repeat steps 1 through 4 for each new controller needed.

Modifying an SDNController

Modify an SDN controller from the SDN Controllers portlet as follows.

1 Right-click *controller* > Edit.

The Editing SDN Controller window is displayed.

2 Modify the information appropriately.

3 Click Save.

The SDN Controllers portlet now lists the new controller.

Deleting an SDN Controller

Delete an SDN controller from the SDN Controllers portlet as follows.

1 Right-click *controller* > Delete.

A confirmation message is displayed.

2 Click Delete.

3 Verify that the controller was removed from the list.

Managing Descriptors

Network function virtualization (NFV) descriptor management is done by users with administrative permissions, such as administrators, network designers, or engineers. This section covers the basic create, import, modify, and delete descriptor tasks. For a description of all the available options to monitor and maintain descriptors, see Network Virtualization Portlets on page 866.

- Maintaining Network Service Descriptors
- Maintaining VNF Descriptors
- Maintaining PNF Descriptors

Maintaining Network Service Descriptors

This section covers the following basic network service descriptor maintenance tasks:

- Creating a Network Service Descriptor
- Importing a Network Service Descriptor
- Modifying a Network Service Descriptor
- Deleting a Network Service Descriptor

For a description of all the available options to monitor and maintain descriptors, see Network Service Portlets on page 878.

Creating a Network Service Descriptor

Create a network service (NS) descriptor from the Network Service Descriptors portlet as follows.

1 Right-click > New.

The Editing Network Service Descriptor window is displayed.

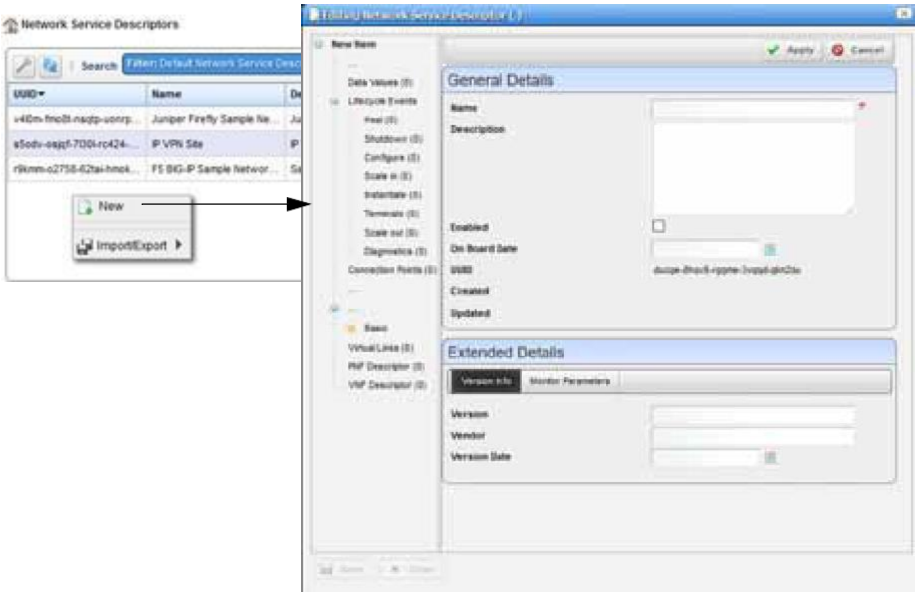

2 Enter a name and optional description.

For example, information for an F5 BIG-IP network service.

3 Enter an on-board date and the version information.

4 Click Apply.

The Editing Network Service Descriptor wind ow is updated based on the information provided so far.

5 Save the record.

6 Verify that the NS descriptors list displays the new descriptor.

Depending on the descriptor, you may need to define dependencies, data values, lifecycle events connection points, and so on. The remaining steps are only for connection points, deployment flavors, virtual links, and VNF descriptors.

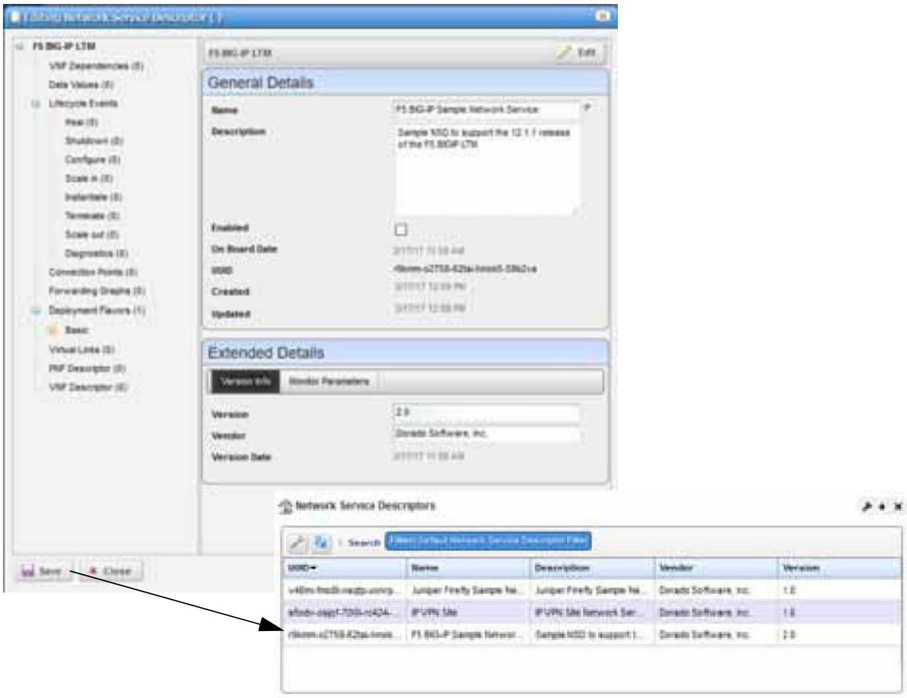

Name	F5 BIG-IP Sample Network Service
Description	Sample NSD to support the 12.1.1 release of the F5 BIGIP LTM
Enabled	□
On Board Date	2/17/17 11:58 AM
UUID	r9kmm-s2758-62tai-hmxk5-59b2va
Created	2/17/17 12:09 PM
Updated	2/17/17 12:09 PM
Version	2.0
Vendor	Dorado Software, Inc.
Version Date	2/17/17 11:58 AM

UUID	Name	Description	Vendor	Version
v40m-fmd9-ragtp-uonrp...	Juniper Firefly Sample Ne...	Juniper Firefly Sample Ne...	Dorado Software, Inc.	1.0
a5odv-oagpf-7309-r2424...	IP VPN Site	IP VPN Site Network Ser...	Dorado Software, Inc.	1.0
r9kmm-s2758-62tai-hmxk...	F5 BIG-IP Sample Networ...	Sample NSD to support t...	Dorado Software, Inc.	2.0

7 Define connection points, such as a NS connection point for an F5 BIG-IP VNF service.

a. Right-click *descriptor* > Edit.

b. Click Connection Points.

Notice that there are no connection points listed.

c. Click Add.

The general parameters are displayed.

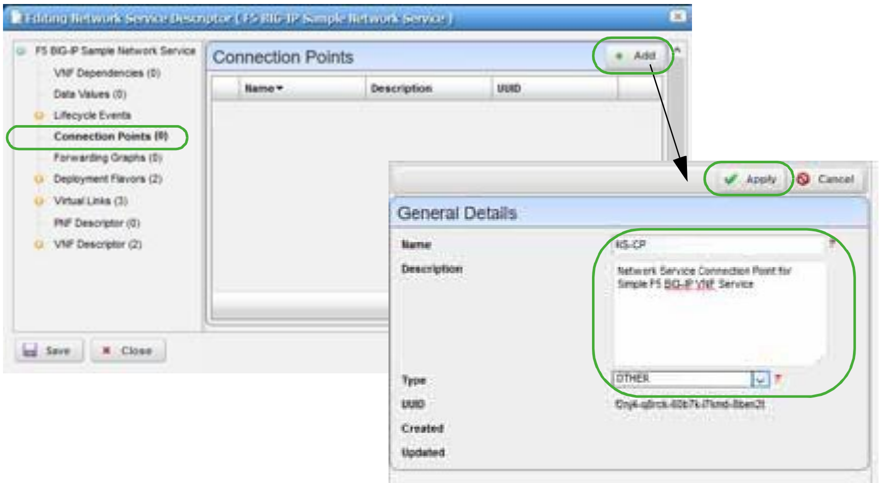

Name: NS-CP

Description: Network Service Connection Point for Simple F5 BIG-IP VNF Service

Type: OTHER

UUID: 0nyk-q8rck-60b7k-l7kmd-8ben2t

Created

Updated

d. Enter the required information.

e. Click Apply.

The new connection point definition is listed.

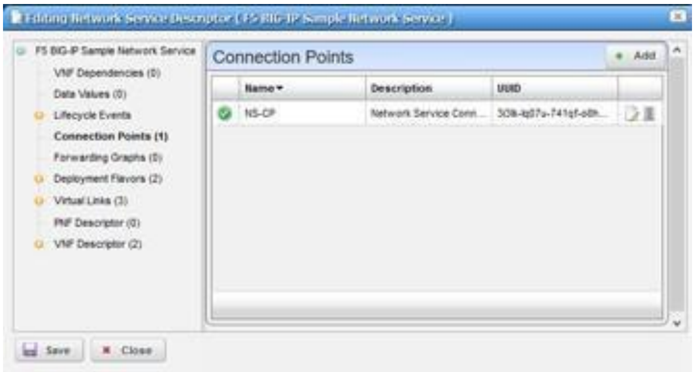

Name	Description	UUID
NS-CP	Network Service Conn...	308-4p07a-741qf-o8h...

8 Define NS deployment flavors, such as standard-small and standard-medium.

a. Click Deployment Flavors.

Notice that there are no deployment flavors listed.

b. Click Add.

c. Enter information for a small flavor.

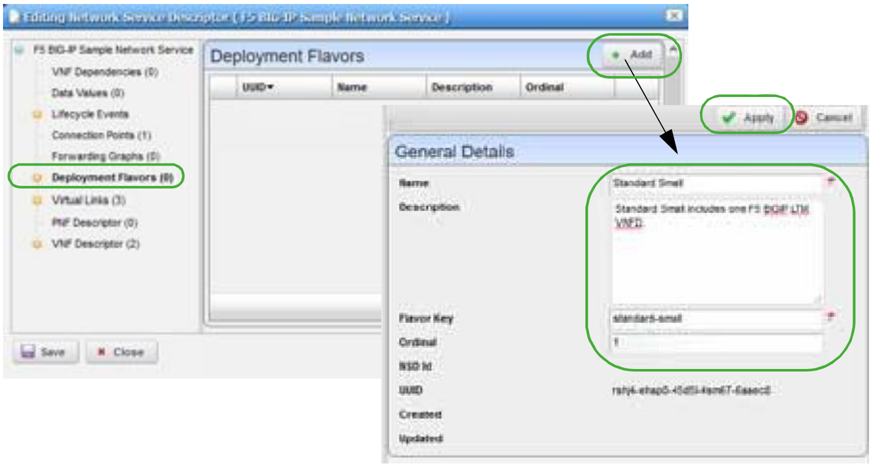

Name: Standard Small

Description: Standard Small includes one F5 BIGIP LTM VNFD.

Flavor Key: standard-small

Ordinal: 1

UUID: rshj4-ehap0-45d5l-4sm6T-6aaecd

d. Click Apply.

The flavor definition is added to the list.

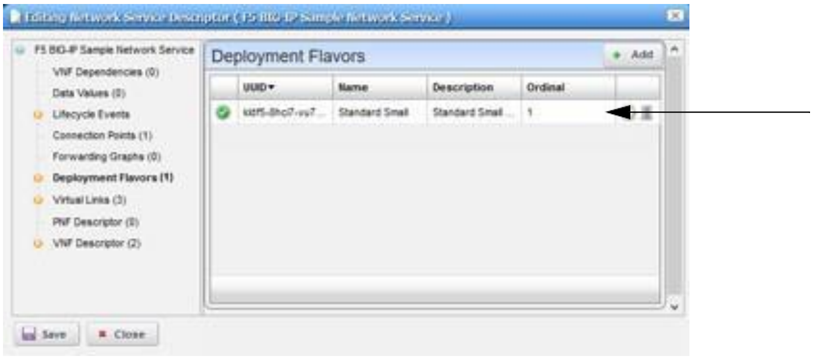

	UUID	Name	Description	Ordinal
✓	kitf5-8hci7-vu7...	Standard Small	Standard Small ...	1

e. Click Add.

f. Enter information for a medium flavor.

g. Click Apply.

The flavor definition is added to the list.

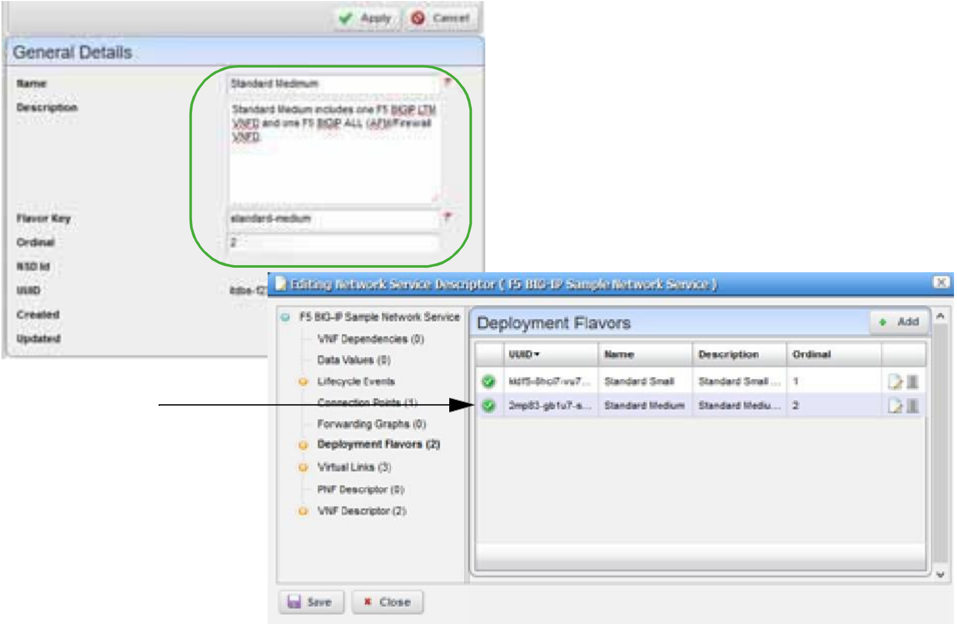

Name: Standard Medium

Description: Standard Medium includes one F5 BIGIP LTM VNFD and one F5 BIGIP ALL (AFM/Firewall VNFD.

Flavor Key: standard-medium

Ordinal: 2

NSD Id

UUID: itdbe-f2

Created

Updated

	UUID	Name	Description	Ordinal
✓	kldf5-8hcl7-vu7...	Standard Small	Standard Small ...	1
✓	2mp83-gb1u7-s...	Standard Medium	Standard Mediu...	2

9 Define NS virtual links, such as a management, internal, and external link.

a. Click Virtual Links.

Notice that there are no virtual links listed.

b. Click Add.

The general and extended parameters are displayed.

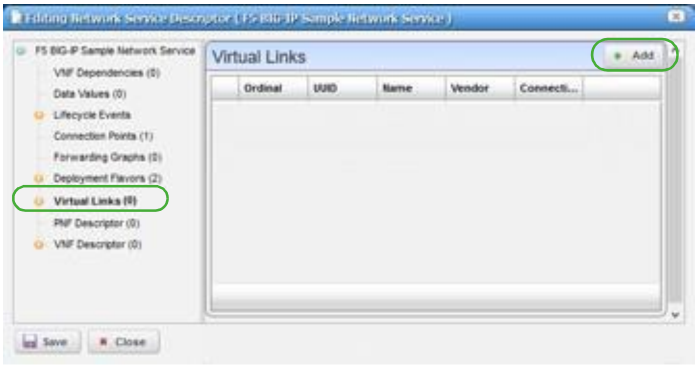

Ordinal	UUID	Name	Vendor	Connecti...

c. Enter the required information, such as name, vendor information, link management method, and connectivity type for a manager link.

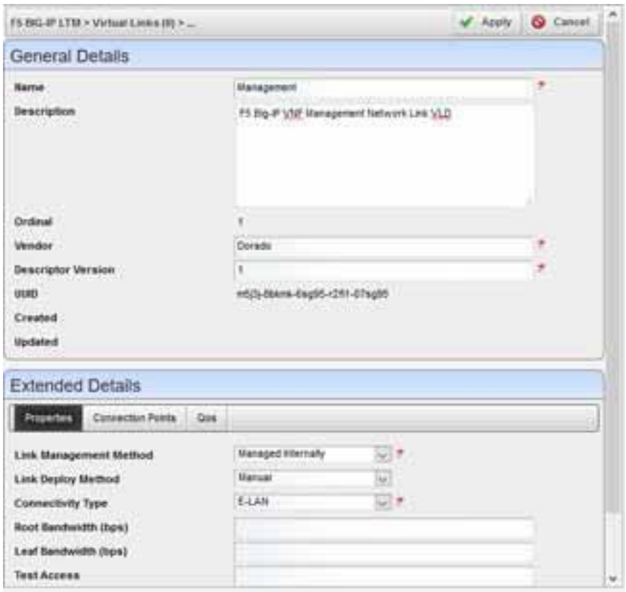

General Details

Name	Management
Description	F5 Big-IP VNF Management Network Link VLD
Ordinal	1
Vendor	Dorado
Descriptor Version	1
UUID	m5j3j-8bkm6-6ag95-r2f41-07ag85
Created	
Updated	

Extended Details

Link Management Method	Managed Internally
Link Deploy Method	Manual
Connectivity Type	E-LAN
Root Bandwidth (bps)	
Leaf Bandwidth (bps)	
Test Access	

d. Click Apply.

The new link is listed.

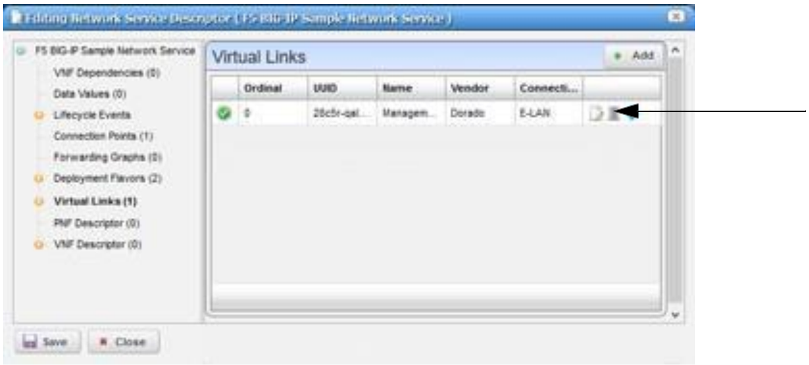

Virtual Links

	Ordinal	UUID	Name	Vendor	Connecti...
✓	0	28c5r-qal...	Managem...	Dorado	E-LAN

e. Click Add.

The general and extended parameters are displayed.

f. Enter the required information, such as name, vendor information, link management method, and connectivity type for an internallink.

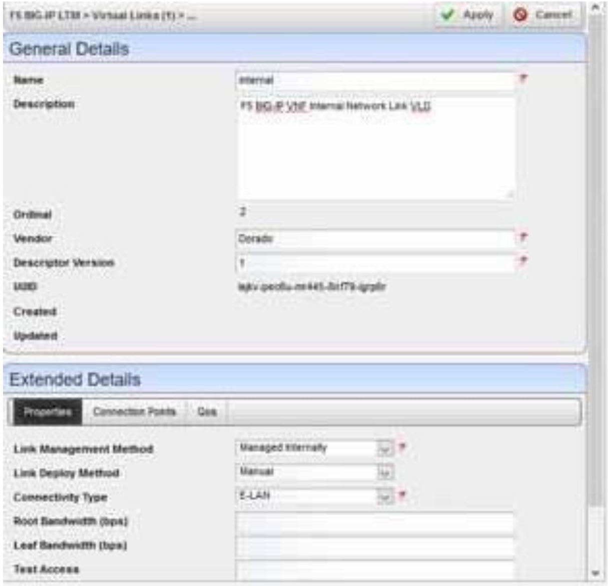

General Details

Name	internal
Description	F5 BIG-IP VNF Internal Network Link VLD
Ordinal	2
Vendor	Dorado
Descriptor Version	1
UUID	lqkv-peo8u-m445-8lf79-lgrplr
Created	
Updated	

Extended Details

Link Management Method	Managed Internally
Link Deploy Method	Manual
Connectivity Type	E-LAN
Root Bandwidth (bps)	
Leaf Bandwidth (bps)	
Test Access	

g. Click Apply.

The new link is listed.

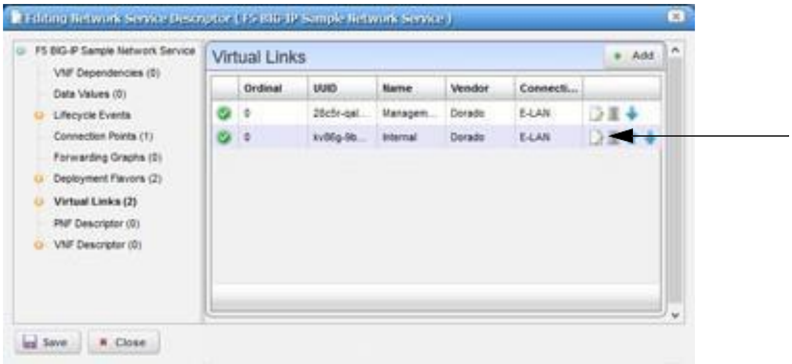

Virtual Links

	Ordinal	UUID	Name	Vendor	Connecti...
✓	0	28c5r-qal...	Managem...	Dorado	E-LAN
✓	0	kv86g-9b...	internal	Dorado	E-LAN

h. Click Add.

The general and extended parameters are displayed.

i. Enter the required information, such as name, vendor information, link management method, and connectivity type for an externallink.

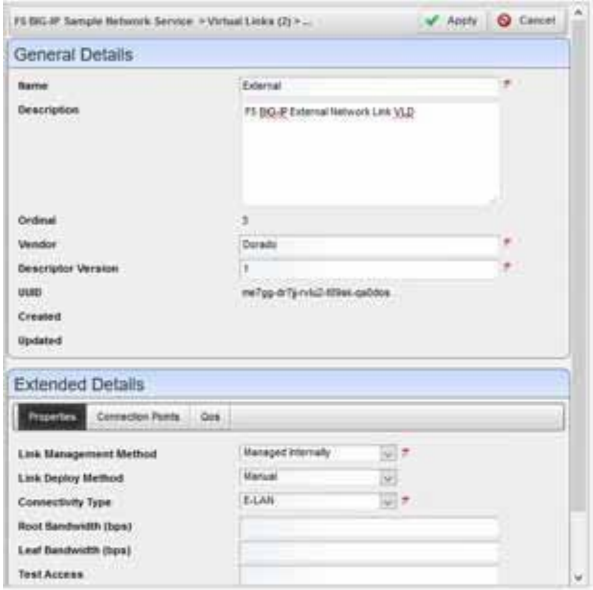

Name	External
Description	F5 BIG-IP External Network Link VLD
Ordinal	3
Vendor	Dorado
Descriptor Version	1
UUID	me7gg-drTj-rvlu2-t09ek-qa0dos
Created	
Updated	
Link Management Method	Managed Internally
Link Deploy Method	Manual
Connectivity Type	E-LAN
Root Bandwidth (bps)	
Leaf Bandwidth (bps)	
Test Access	

j. Click Apply.

The new link is listed.

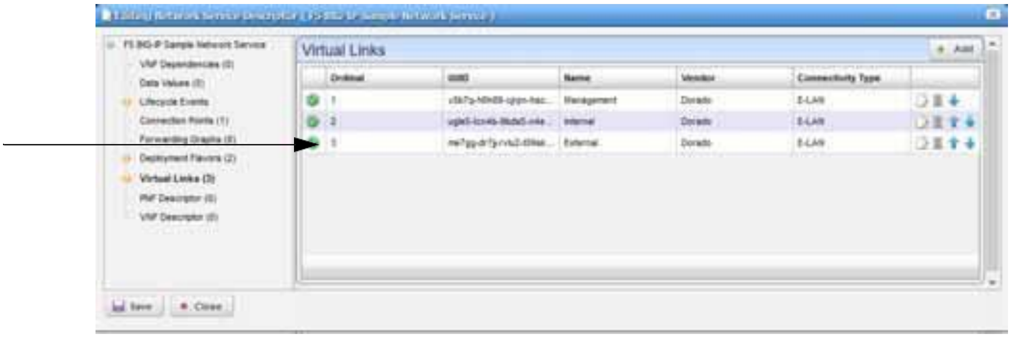

Ordinal	UUID	Name	Vendor	Connectivity Type
1	v5k7g-h8h89-qlgn-hac...	Management	Dorado	E-LAN
2	ugle5-lco4s-8kda5-n4e...	Internal	Dorado	E-LAN
3	me7gg-drTj-rvlu2-t09ek...	External	Dorado	E-LAN

10 Define VNF descriptors, such as F5 BIG-IP LTM and F5 BIG-IP ALL AFM virtual appliances.

a. Click VNF Descriptor.

Notice that there are no VNF descriptors listed.

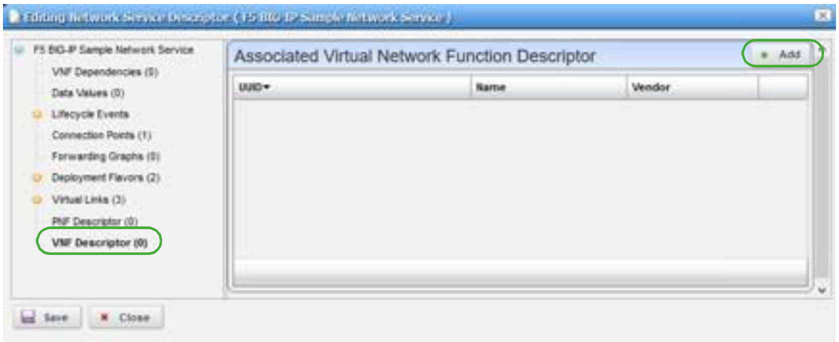

b. Click Add.

A list of descriptor entities are displayed.

c. Select a descriptor from which to create a VNF.

For example, F5 BIG-IP LTM Virtual Appliance.

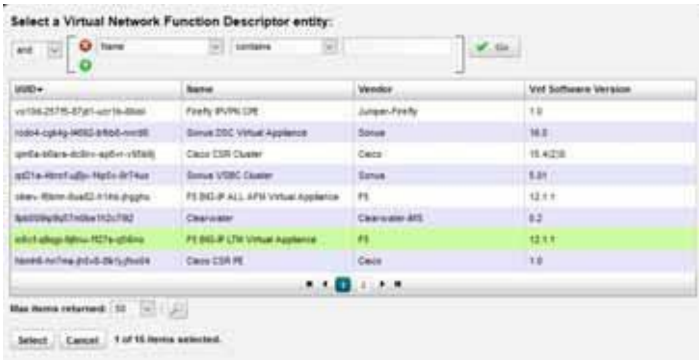

Select a Virtual Network Function Descriptor entity:

UUID	Name	Vendor	Vnf Software Version
[illegible]	Firefly IPVPN CPE	Juniper-Firefly	1.0
[illegible]	Sonus DSC Virtual Appliance	Sonus	16.0
[illegible]	Cisco CSR Cluster	Cisco	15.4(2)S
[illegible]	Sonus VSBC Cluster	Sonus	5.01
[illegible]	F5 BIG-IP ALL AFM Virtual Appliance	F5	12.1.1
[illegible]	Clearwater	Clearwater-IMS	0.2
ic6cf-q9qgi-9j6nu-f127a-q5i0no	F5 BIG-IP LTM Virtual Appliance	F5	12.1.1
[illegible]	Cisco CSR PE	Cisco	1.0

Max items returned: 50

1 of 15 items selected.

d. Click Select.

The selected descriptor is added to the associated VNF descriptor list.

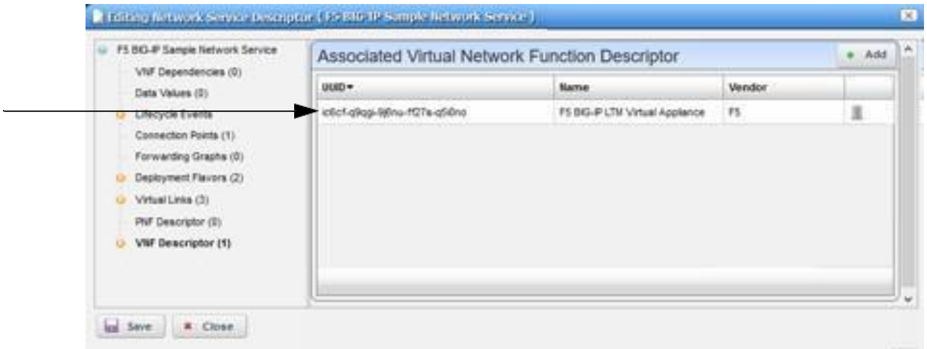

Editing Network Service Descriptor (F5 BIG-IP Sample Network Service)

Associated Virtual Network Function Descriptor

UUID	Name	Vendor
ic6cf-q9qgi-9j6nu-f127a-q5i0no	F5 BIG-IP LTM Virtual Appliance	F5

e. Click Add.

A list of descriptor entities are displayed.

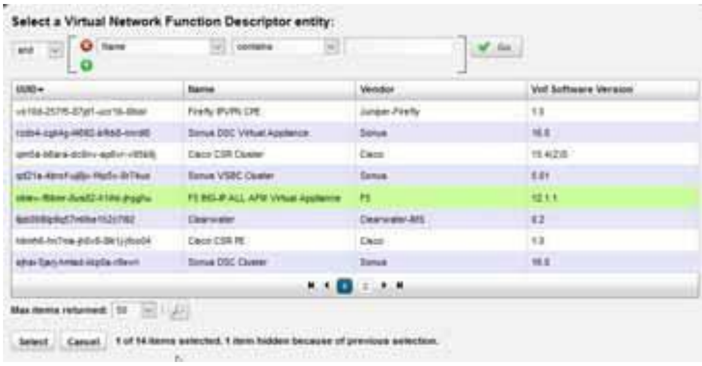

Select a Virtual Network Function Descriptor entity:

UUID	Name	Vendor	Vnf Software Version
[illegible]	Firefly IPVPN CPE	Juniper-Firefly	1.0
[illegible]	Sonus DSC Virtual Appliance	Sonus	16.0
[illegible]	Cisco CSR Cluster	Cisco	15.4(2)S
[illegible]	Sonus VSBC Cluster	Sonus	5.01
[illegible]	F5 BIG-IP ALL AFM Virtual Appliance	F5	12.1.1
[illegible]	Clearwater	Clearwater-IMS	0.2
[illegible]	Cisco CSR PE	Cisco	1.0
[illegible]	Sonus DSC Cluster	Sonus	16.0

Max items returned: 50

1 of 14 items selected. 1 item hidden because of previous selection.

f. Select a descriptor from which to create a VNF.

For example, F5 BIG-IP All AFM Virtual Appliance.

g. Click Select.

The selected descriptor is added to the associated VNF descriptor list.

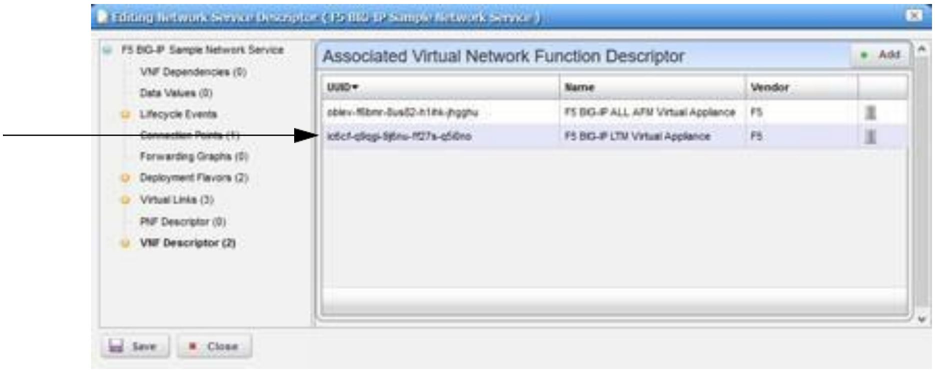

Associated Virtual Network Function Descriptor

UUID▾	Name	Vendor
oblev-f6bmr-8us82-h1ihk-jhgghu	F5 BIG-IP ALL AFM Virtual Appliance	F5
ic6cf-q9qgi-9jbnu-ff27s-q5i0no	F5 BIG-IP LTM Virtual Appliance	F5

11 Make this descriptor available.

a. Click the descriptor name, such as F5 BIG-IP Sample Network Service.

The general details are displayed.

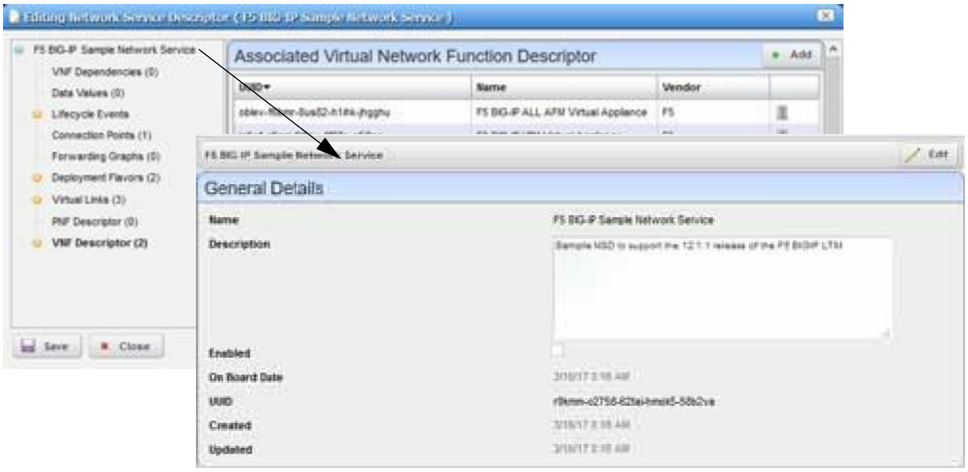
Editing Network Service Descriptor (F5 BIG-IP Sample Network Service)
F5 BIG-IP Sample Network Service
VNF Dependencies (0)
Data Values (0)
Lifecycle Events
Connection Points (1)
Forwarding Graphs (0)
Deployment Flavors (2)
Virtual Links (3)
PNF Descriptor (0)
VNF Descriptor (2)
Associated Virtual Network Function Descriptor
Add
UUID▾
Name
Vendor
oblev-f6bmr-8us82-h1ihk-jhgghu
F5 BIG-IP ALL AFM Virtual Appliance
F5
Save
Close
F5 BIG-IP Sample Network Service
Edit

General Details

Name: F5 BIG-IP Sample Network Service

Description: Sample NSD to support the 12.1.1 release of the F5 BIGIP LTM

Enabled: □

On Board Date: 3/16/17 8:16 AM

UUID: r9kmm-o2758-82tai-hmek5-58b2va

Created: 3/16/17 8:16 AM

Updated: 3/16/17 8:16 AM

b. Click Edit.

c. Select Enable.

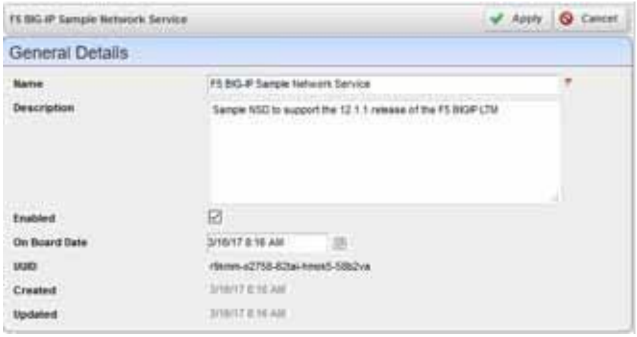

General Details

Name: F5 BIG-IP Sample Network Service

Description: Sample NSD to support the 12.1.1 release of the F5 BIGIP LTM

Enabled: ☑

On Board Date: 3/16/17 8:16 AM

UUID: r9kmm-o2758-82tai-hmek5-58b2va

Created: 3/16/17 8:16 AM

Updated: 3/16/17 8:16 AM

12 Save the NS descriptor.

Importing a Network Service Descriptor

Youcan import a network service (NS) descriptor from another system or from a backup that was exported to an XML file or you can import a descriptor from a URL. The example provided here imports a descriptor from an XML file.

Import an NS descriptor from the Network Service Descriptors portlet as follows.

1 Right-click > Import/Export. > Import.

The Import Network Services Descriptor(s) window is displayed.

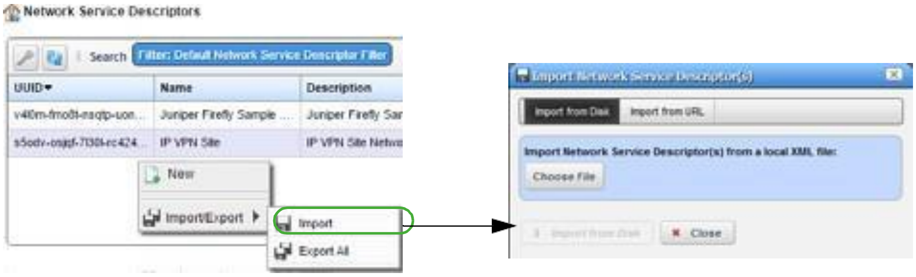
Network Service Descriptors
Search
Filter: Default Network Service Descriptor Filter
UUID
Name
Description
Juniper Firefly Sample ...
Juniper Firefly Sar
IP VPN Site
IP VPN Site Netwo
New
Import/Export
Import
Export All
Import Network Service Descriptor(s)
Import from Disk
Import from URL
Import Network Service Descriptor(s) from a local XML file:
Choose File
Import from Disk
Close

2 Click Choose File.

The navigation window is displayed.

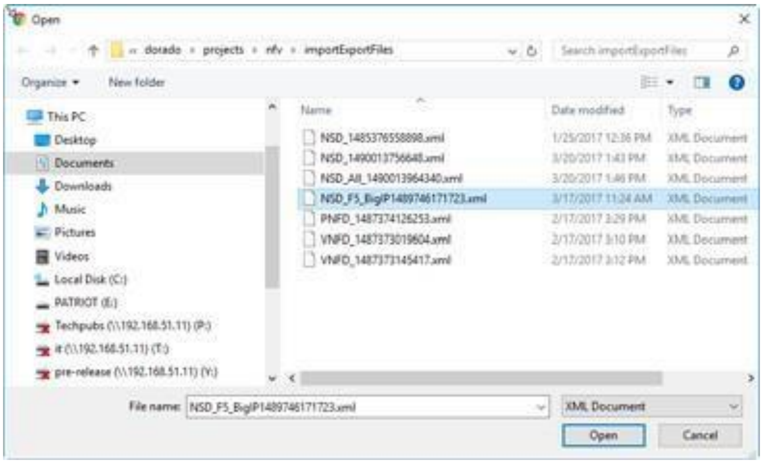

3 Navigate to the XML file.

4 Select the file to import.

For example: `NSD_F5_BigIP1485376558898.xml`

5 Click Open.

The Import Network Service Descriptor(s) Import from Disk button is activated.

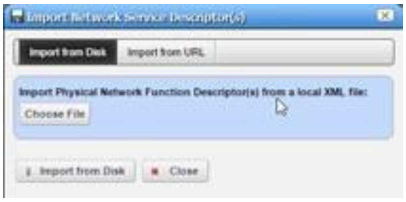
Import Network Service Descriptor(s)
Import from Disk
Import from URL
Import Physical Network Function Descriptor(s) from a local XML file:
Choose File
Import from Disk
Close

6 Click Import from Disk.

The selected descriptor is added to the PNF list.

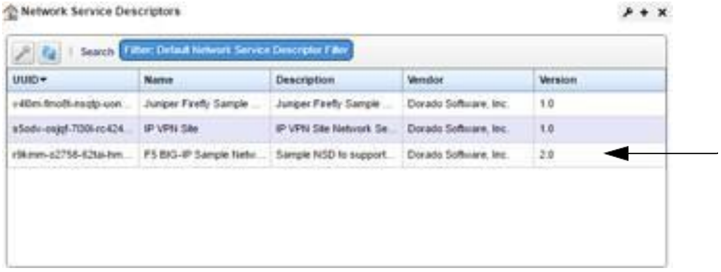
Network Service Descriptors

Search Filter: Default Network Service Descriptor Filter

UUID	Name	Description	Vendor	Version
v40m-fmo8t-nsqtp-uon...	Juniper Firefly Sample ...	Juniper Firefly Sample ...	Dorado Software, Inc.	1.0
a5odv-osjgf-7t30l-rc424...	IP VPN Site	IP VPN Site Network Se...	Dorado Software, Inc.	1.0
r9kmm-s2758-62tai-hm...	F5 BIG-IP Sample Netw...	Sample NSD to support...	Dorado Software, Inc.	2.0

Modifying a Network Service Descriptor

Modify a network service (NS) descriptor from the Network Service Descriptors portlet as follows.

1 Right-click > Edit.

The Editing Network Service Descriptor window is displayed.

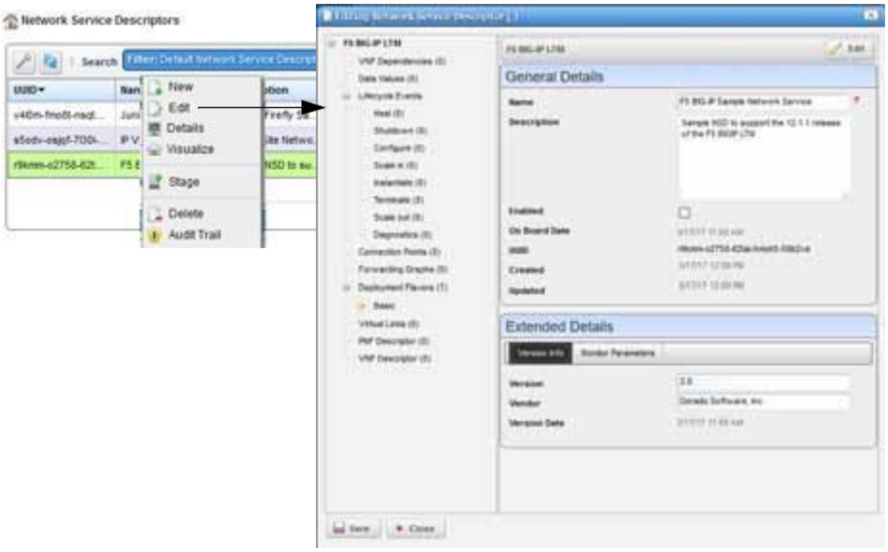

Name: F5 BIG-IP Sample Network Service

Description: Sample NSD to support the 12.1.1 release of the F5 BIGIP LTM

Enabled: □

Version: 2.0

Vendor: Dorado Software, Inc.

2 Click Edit to modify the general or extended details.

3 Modify any of the following needing changes:

- VNF Dependencies
- Data Values
- Lifecycle Events
- Connection Points
- Forwarding Graphic
- Deployment Flavors
- Virtual Links

- PNF Descriptor
- VNF Descriptor

4 Save your changes.

Deleting a Network Service Descriptor

Delete a network service (NS) descriptor from the Network Service Descriptors portlet as follows.

1 Right-click > Delete.

A confirmation message is displayed.

2 Click Delete.

The descriptor is removed.

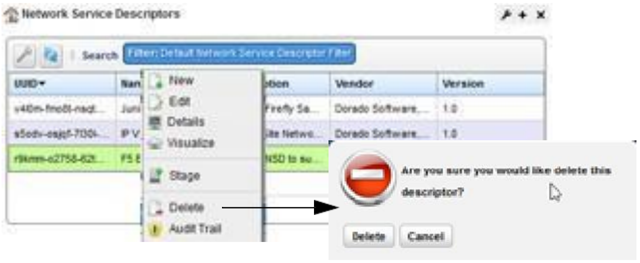
Network Service Descriptors
Search
Filter: Default Network Service Descriptor Filter
New
Edit
Details
Visualize
Stage
Delete
Audit Trail
Delete
Cancel

UUID▾	Nan	...tion	Vendor	Version
v40m-fmo8t-naqt...	Juni	Firefly Sa...	Dorado Software,...	1.0
a5odv-oajgf-7D0l-...	IP V	...ite Netwo...	Dorado Software,...	1.0
r9kmm-o2758-62t...	F5 E	NSD to su...		

Are you sure you would like delete this descriptor?

3 Verify that the descriptor is no longer listed.

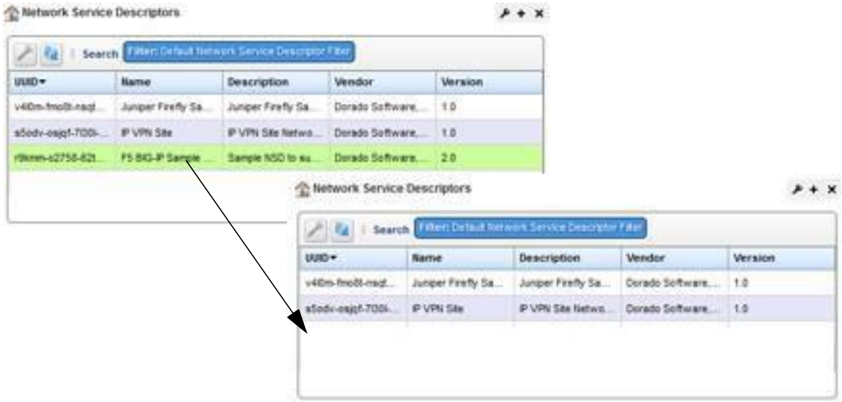
Network Service Descriptors
Search
Filter: Default Network Service Descriptor Filter
Network Service Descriptors
Search
Filter: Default Network Service Descriptor Filter

UUID▾	Name	Description	Vendor	Version
v40m-fmo8t-naqt...	Juniper Firefly Sa...	Juniper Firefly Sa...	Dorado Software,...	1.0
a5odv-oajgf-7D0l-...	IP VPN Site	IP VPN Site Netwo...	Dorado Software,...	1.0
r9kmm-o2758-62t...	F5 BIG-IP Sample ...	Sample NSD to su...	Dorado Software,...	2.0

UUID▾	Name	Description	Vendor	Version
v40m-fmo8t-naqt...	Juniper Firefly Sa...	Juniper Firefly Sa...	Dorado Software,...	1.0
a5odv-oajgf-7D0l-...	IP VPN Site	IP VPN Site Netwo...	Dorado Software,...	1.0

Maintaining VNF Descriptors

Thissectioncoversthefollowingbasicvirtualized networkfunction(VNF) descriptormaintenance tasks:

- Creating a VNF Descriptor
- Importing a VNF Descriptor
- Modifying a VNF Descriptor
- Deleting a VNF Descriptor

For a description of all the available options to monitor and maintain descriptors, see Network Service Portlets on page 878.

Creating a VNFDescriptor

Create a virtualized network function (VNF) from the Virtual Network Function portlet as follows.

1 Right-click > New.

The general and extended details are displayed.

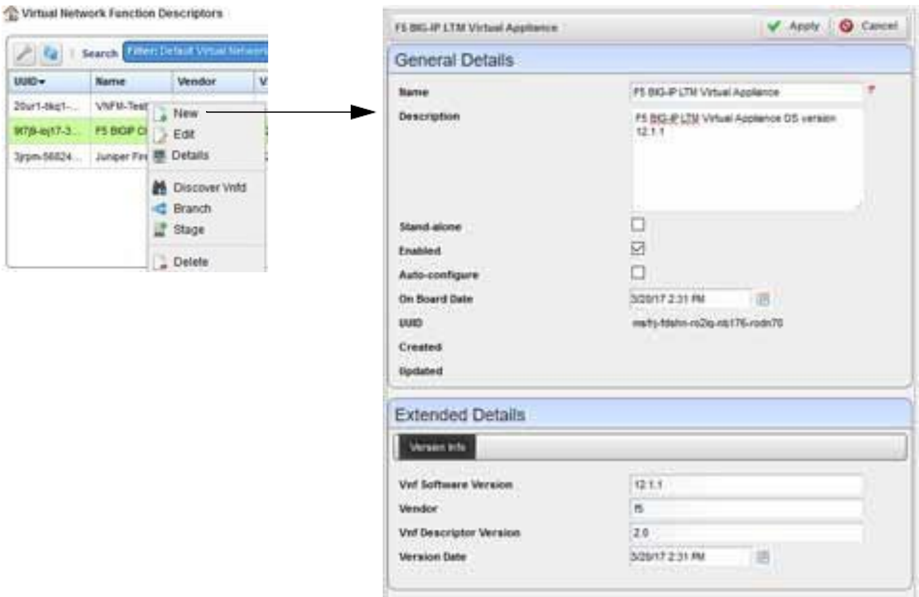

F5 BIG-IP LTM Virtual Appliance

General Details

Name	F5 BIG-IP LTM Virtual Appliance
Description	F5 BIG-IP LTM Virtual Appliance DS version 12.1.1
Stand-alone	□
Enabled	☑
Auto-configure	□
On Board Date	3/20/17 2:31 PM
UUID	msfrj-fdshn-ro2ig-nb176-rodn70
Created	
Updated	

Extended Details

Vnf Software Version	12.1.1
Vendor	f5
Vnf Descriptor Version	2.0
Version Date	3/20/17 2:31 PM

2 Enter the general and extended details, such as information for an F5 BIG-IP LTM Virtual Appliance VNF descriptor.

3 Click Apply.

The entered information is applied to the descriptor.

4 Define VNF deployment flavors, such as F5 BIGIP LTMVA Small.

a. Click Deployment Flavors.

b. Click Add.

c. Enter information for a small flavor.

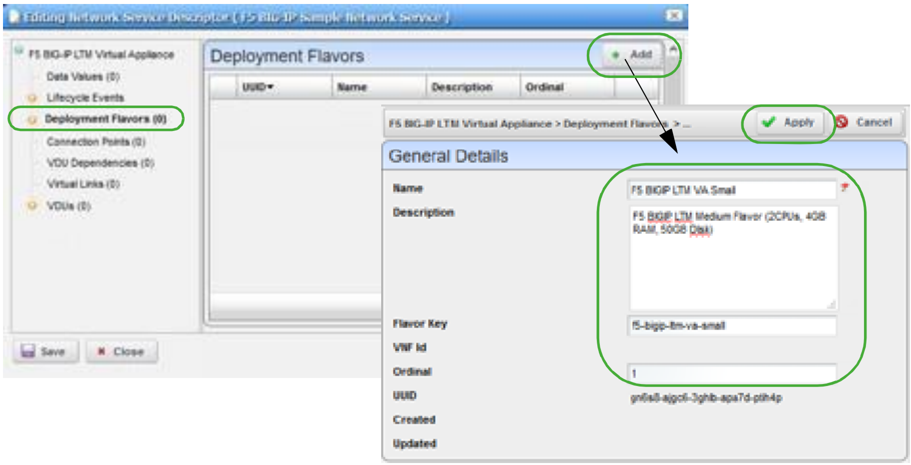

F5 BIG-IP LTM Virtual Appliance > Deployment Flavors > ...

General Details

Name	F5 BIGIP LTM VA Small
Description	F5 BIGIP LTM Medium Flavor (2CPUs, 4GB RAM, 50GB Disk)
Flavor Key	f5-bigip-ltm-va-small
VNF Id	
Ordinal	1
UUID	gn6s8-ajgcll-3ghb-apa7d-plh4p
Created	
Updated	

d. Click Apply.

The flavor definition is added to the list.

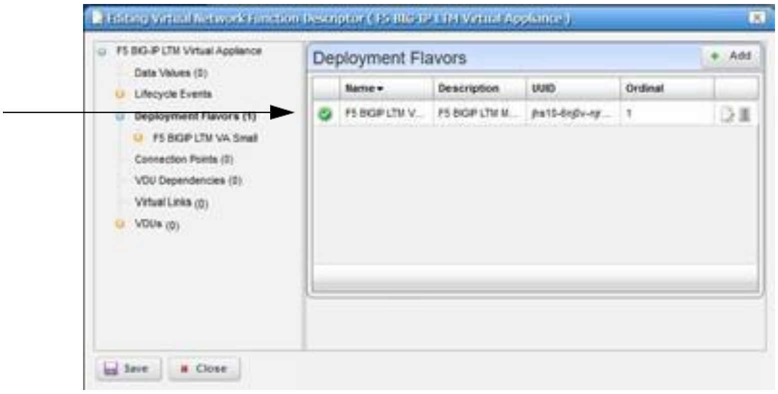

	Name	Description	UUID	Ordinal
✓	F5 BIGIP LTM V...	F5 BIGIP LTM M...	jhs10-6nj0v-nj...	1

5 Define connection points, such as management, internal, and external.

a. Click Connection Points.

The Connection Points list is displayed.

b. Click Add.

The general details are displayed.

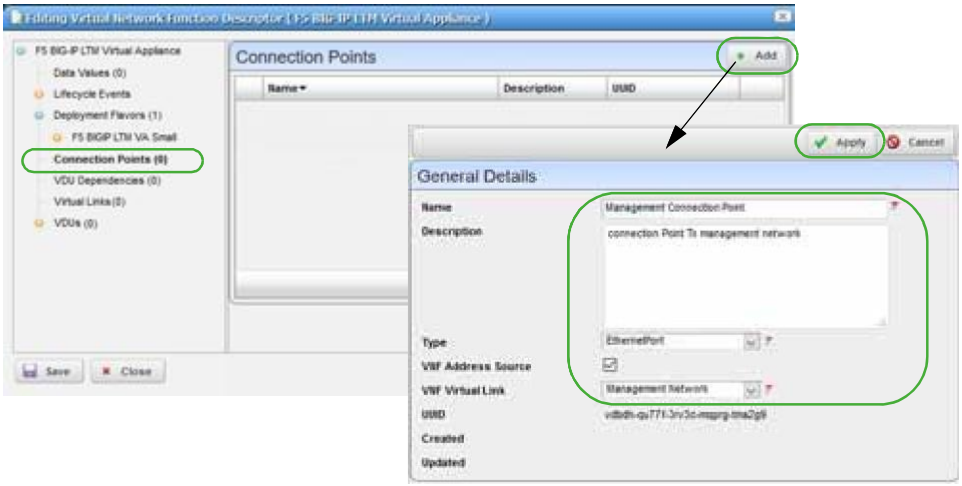

Name: Management Connection Point

Description: connection Point To management network

Type: EthernetPort

VNF Address Source: ☑

VNF Virtual Link: Management Network

UUID: vdbdh-gu771-3rv3c-mgprg-tma2g9

Created

Updated

c. Enter the required information, such as in formation for a management connection point.

d. Click Apply.

The new connection definition is listed.

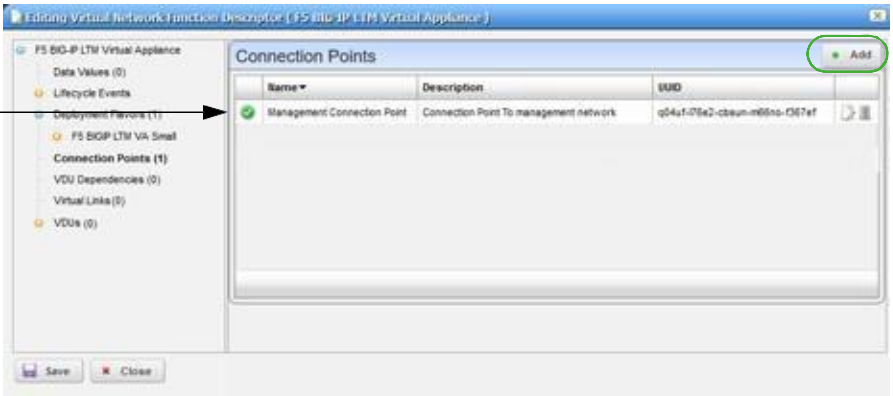

Connection Points

	Name	Description	UUID
✓	Management Connection Point	Connection Point To management network	q04uf-i76e2-cbsun-m66no-f367ef

e. Click Add.

The general details are displayed.

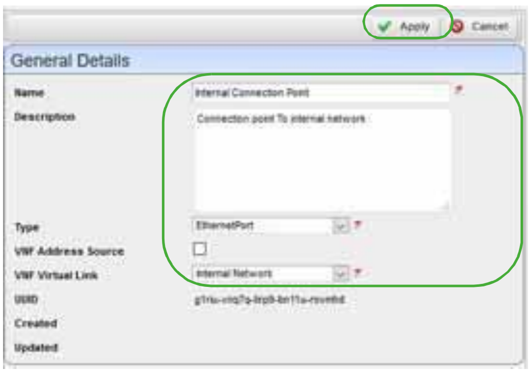

General Details

Name	Internal Connection Point
Description	Connection point To internal network
Type	EthernetPort
VNF Address Source	□
VNF Virtual Link	Internal Network
UUID	g1riu-viq7g-lirp9-bn11a-rsvmhd
Created	
Updated	

f. Enter the required information, such as in formation for an intern al connection point.

g. Click Apply.

The new connection definition is listed.

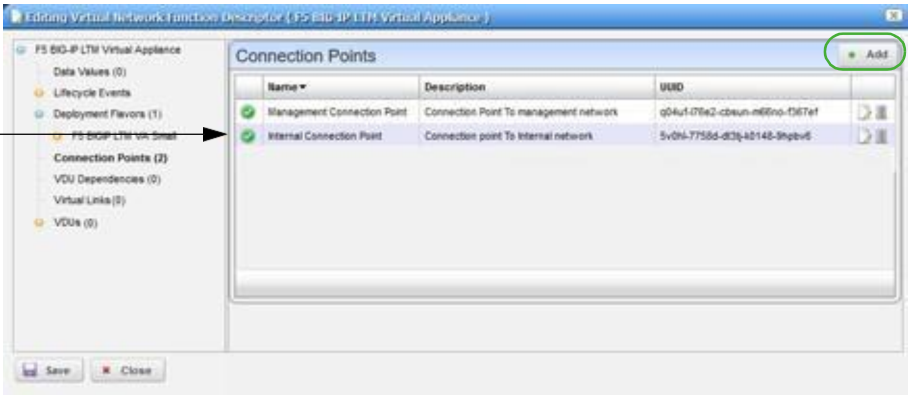

Connection Points

	Name	Description	UUID
✓	Management Connection Point	Connection Point To management network	q04uf-i76e2-cbsun-m66no-f367ef
✓	Internal Connection Point	Connection point To internal network	5v0hl-7758d-dl3tj-k8148-9hpbv6

h. Click Add.

The general details are displayed.

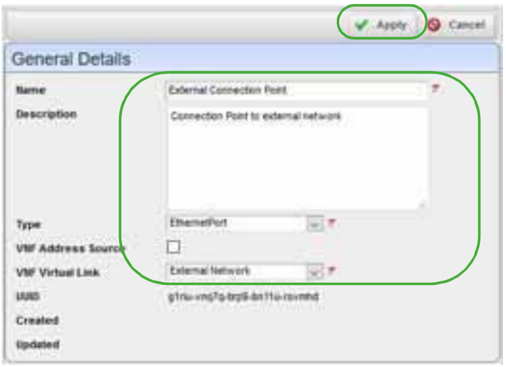

Name	External Connection Point
Description	Connection Point to external network
Type	EthernetPort
VNF Address Source	□
VNF Virtual Link	External Network
UUID	g1riu-vnq7g-brg9-bn11u-rovmhd
Created	
Updated	

i. Enter the required information, such as in formation for an extern al connection point.

j. Click Apply.

The new connection definition is listed.

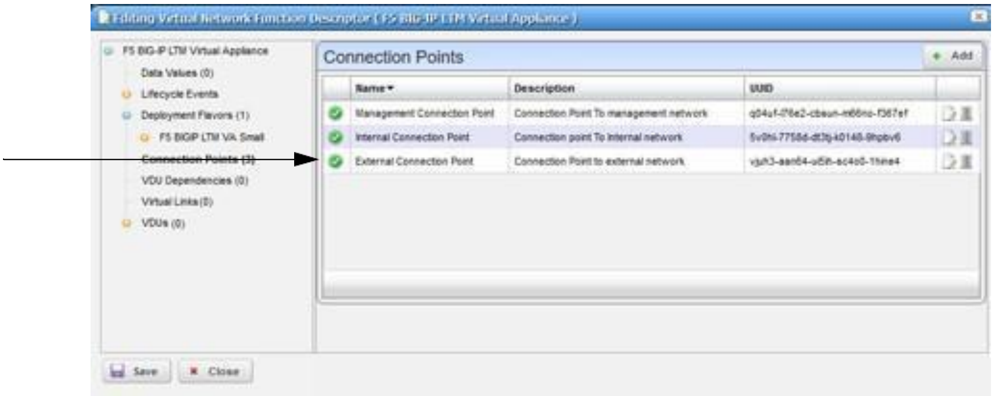

	Name	Description	UUID
✓	Management Connection Point	Connection Point To management network	q04uf-i76e2-cbeun-m66no-f367ef
✓	Internal Connection Point	Connection point To internal network	5v0hi-7758d-dt3g-k0148-9hpbv6
✓	External Connection Point	Connection Point to external network	vjuh3-aan64-u5th-ac4o0-1hme4

6 Define virtual links, such as management, internal, and external.

a. Click Virtual Links.

Notice no links are listed.

b. Click Add.

The general and extended parameters are displayed.

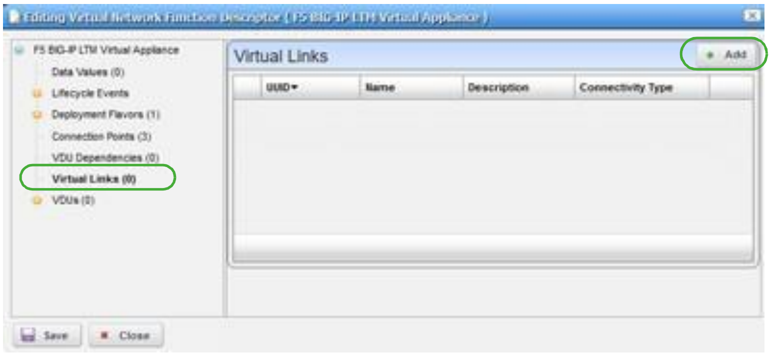

	UUID	Name	Description	Connectivity Type

c. Enter the required information, such as name, vendor information, link management method, and connectivity type for a manager link.

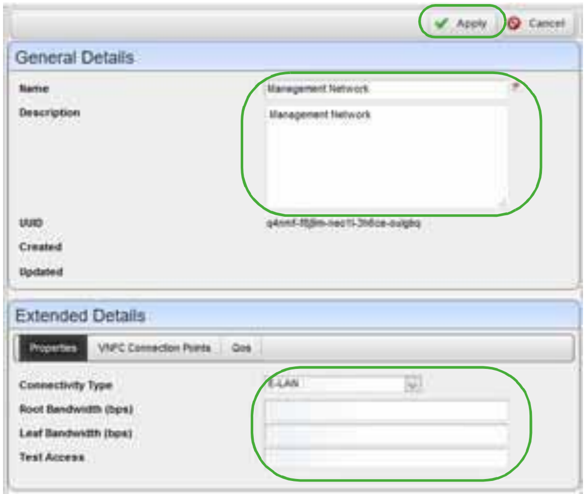

Name	Management Network
Description	Management Network
UUID	q4nmf-f8j8m-nec11-3h6ce-ouigbq
Created	
Updated	
Connectivity Type	E-LAN
Root Bandwidth (bps)	
Leaf Bandwidth (bps)	
Test Access	

d. Click Apply.

The new link is listed.

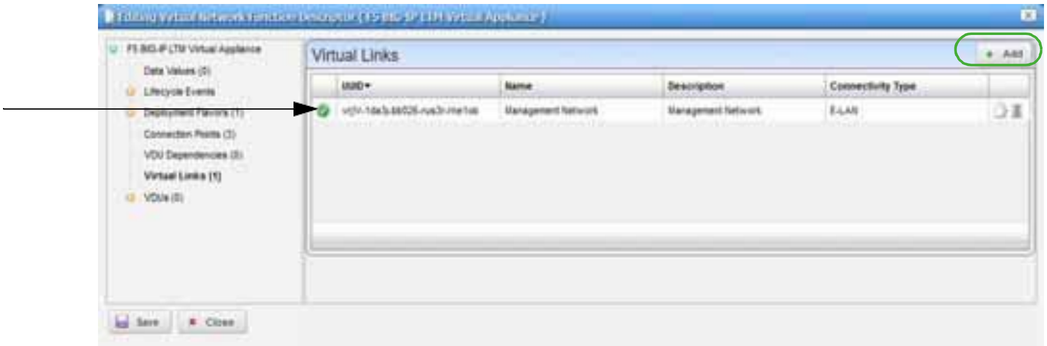

UUID	Name	Description	Connectivity Type
[illegible]	Management Network	Management Network	E-LAN

e. Click Add.

The general and extended parameters are displayed.

f. Enter the required information, such as name, vendor information, link management method, and connectivity type for an externallink.

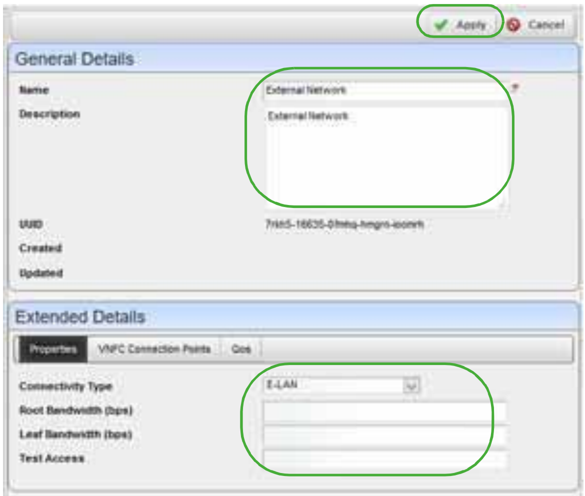

General Details

Name	External Network *
Description	External Network
UUID	7rkh5-16635-0fmhq-hmgrn-ioomrh
Created	
Updated	

Extended Details

Connectivity Type	E-LAN
Root Bandwidth (bps)	
Leaf Bandwidth (bps)	
Test Access	

g. Click Apply.

The new link is listed.

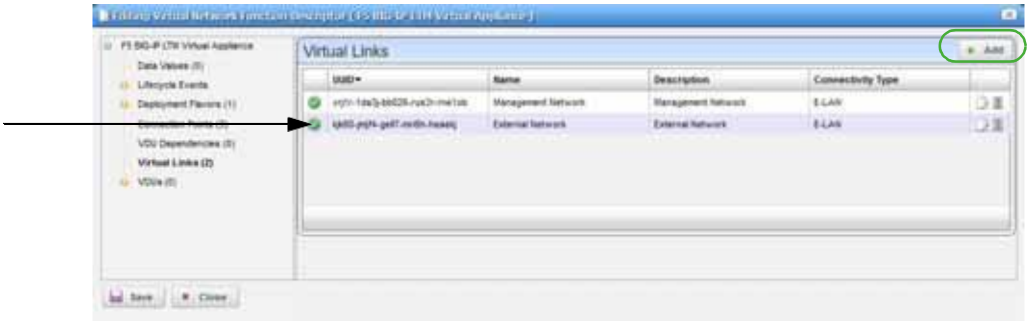

Virtual Links

	UUID▾	Name	Description	Connectivity Type
✓	yrjh-1da3j-bb028-rux3r-me1ob	Management Network	Management Network	E-LAN
✓	ijk60-yrjhk-ge87-nn6h-haaej	External Network	External Network	E-LAN

h. Click Add.

The general and extended parameters are displayed.

i. Enter the required information, such as name, vendor information, link management method, and connectivity type for an internallink.

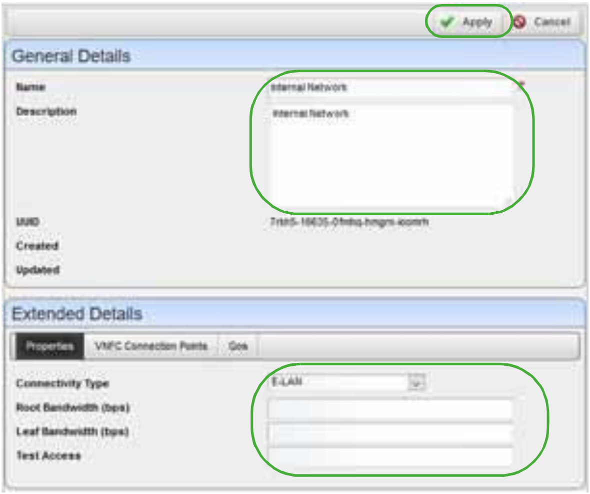

Name	Internal Network
Description	Internal Network
UUID	7rkh5-18635-0fnhq-hmgrn-iconrh
Created	
Updated	

Connectivity Type	E-LAN
Root Bandwidth (bps)	
Leaf Bandwidth (bps)	
Test Access	

j. Click Apply.

The new link is listed.

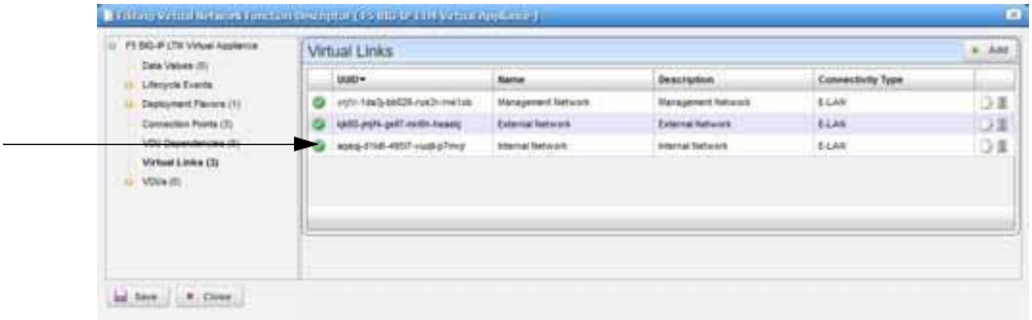

	UUID▾	Name	Description	Connectivity Type
✓	vrjhr-1da3j-bb028-rua3r-me1ob	Management Network	Management Network	E-LAN
✓	kjk00-jmjhk-ge87-mr8h-haaaq	External Network	External Network	E-LAN
✓	apaq-d1k8-4957-vud8-p7mqr	Internal Network	Internal Network	E-LAN

7 Specify VDUs, such as BIGIP LTM VA.

a. Click VDUs.

Notice that no VDUs are listed.

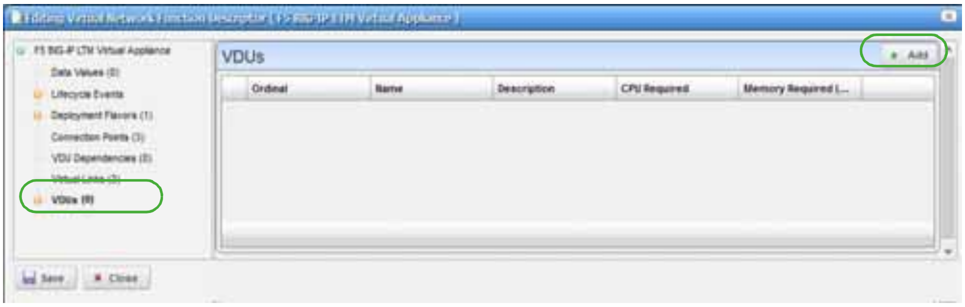

Ordinal	Name	Description	CPU Required	Memory Required (...

b. Click Add.

The general and extended parameters are displayed.

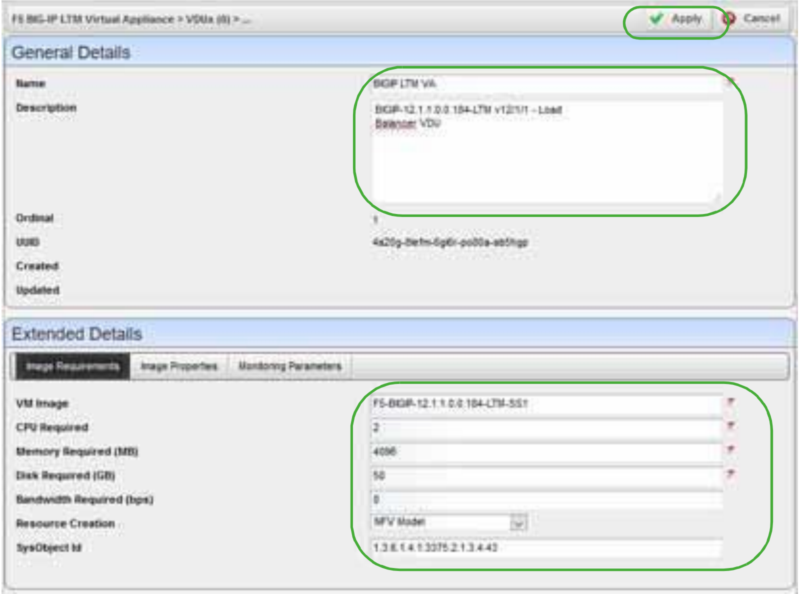

General Details

Name	BIGIP LTM VA
Description	BIGIP-12.1.1.0.0.184-LTM v12/1/1 - Load Balancer VDU
Ordinal	1
UUID	4a20g-8efm-6g6r-po00a-ab5hgp
Created	
Updated	

Extended Details

VM Image	F5-BIGIP-12.1.1.0.0.184-LTM-SS1
CPU Required	2
Memory Required (MB)	4096
Disk Required (GB)	50
Bandwidth Required (bps)	0
Resource Creation	NFV Model
SysObject Id	1.3.6.1.4.1.3375.2.1.3.4.43

c. Enter the required information, such as name and image requirements for F5 BIG-IP LTM.

d. Click Apply.

The new VDU is listed.

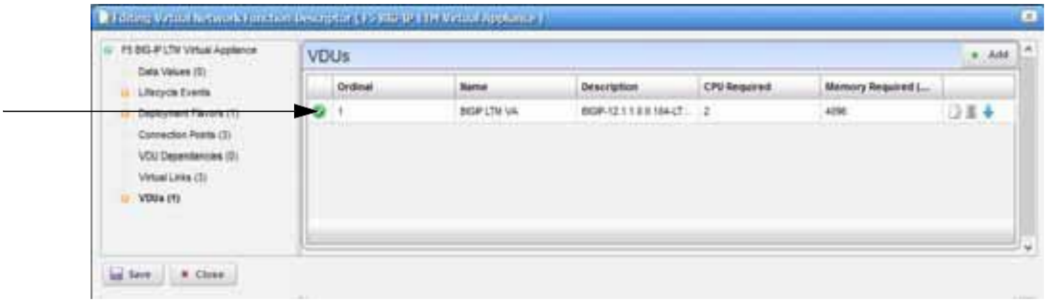

VDUs

	Ordinal	Name	Description	CPU Required	Memory Required (...
✓	1	BIGIP LTM VA	BIGIP-12.1.1.0.0.184-LT...	2	4096

8 Save your new VNF.

Importing a VNF Descriptor

Youcan import a virtualized network function (V NF) descriptor from another system or from a backup that was exported to a XML file or you can import a descriptor from a URL. The example provided here imports a descriptor from an XMLfile.

Import a virtualized network function from the Virtual Network Function portlet as follows.

1 Right-click > Import/Export. > Import.

The Import Virtualized Network Function Descriptor(s) window is displayed.

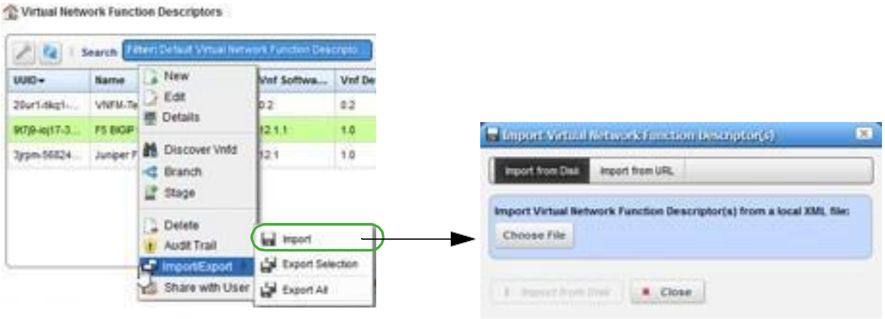

2 Click Choose File.

3 Navigate to the XML file.

4 Select the file to import.

For example: `VNFD_1487373019604.xml`

5 Verify that the descriptor was added to the VNF descriptors list.

Modifying a VNF Descriptor

Modify a virtualized network function (VNF) from the Virtual Network Function portlet as follows.

1 Right-click > Edit.

The Editing Network Service Descriptor window is displayed.

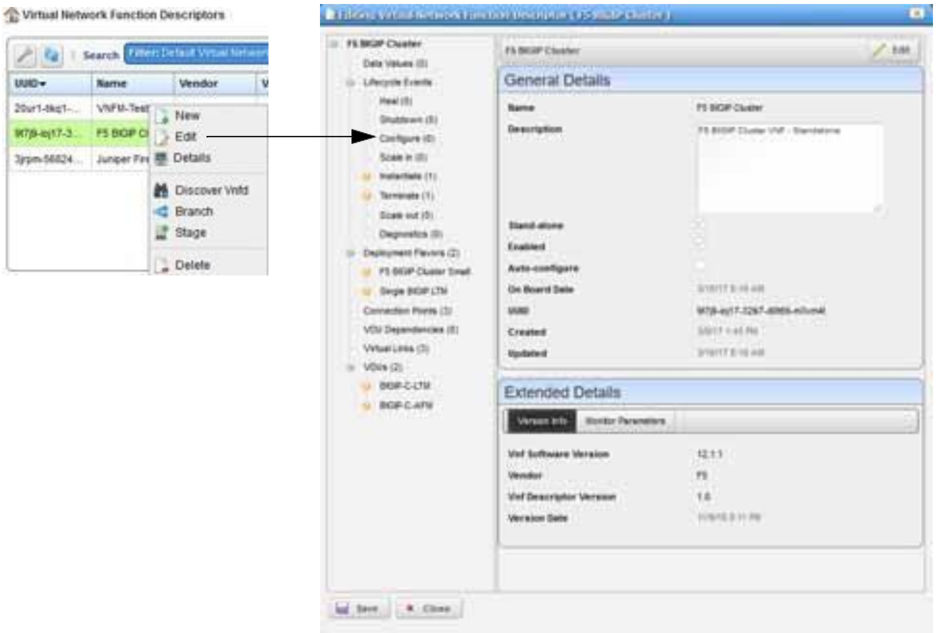

General Details

Name F5 BIGIP Cluster

Description F5 BIGIP Cluster VNF - Standalone

Stand-alone

Enabled

Auto-configure

Extended Details

Vnf Software Version 12.1.1

Vendor F5

Vnf Descriptor Version 1.0

2 Click Edit to modify the general or extended details.

3 Modify any of the following needing changes:

- D ata Values
- Lifecy cle Events
- Dep loyment Flavors
- Connect ion Points
- V DU Dependencies
- Virtual Links
- V DUs

4 Save your changes.

Deleting a VNF Descriptor

Delete a virtualized network function (VNF) from the Virtual Network Function portlet as follows.

1 Right-click > Delete.

A confirmation message is displayed.

2 Click Delete.

The descriptor is removed.

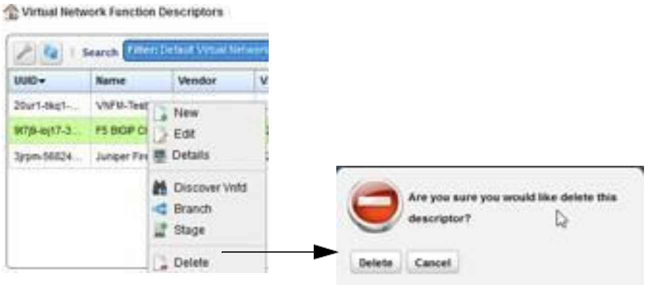

Virtual Network Function Descriptors
Search
Filter: Default Virtual Network
New
Edit
Details
Discover Vnfd
Branch
Stage
Delete
Delete
Cancel

UUID▾	Name	Vendor
20ur1-tkq1-...	VNFM-Test	
9t7j9-iej17-3...	F5 BIGP Cl	
3jrpm-56824...	Juniper Fir	

Are you sure you would like delete this descriptor?

3 Verify that the descriptor is no longer listed.

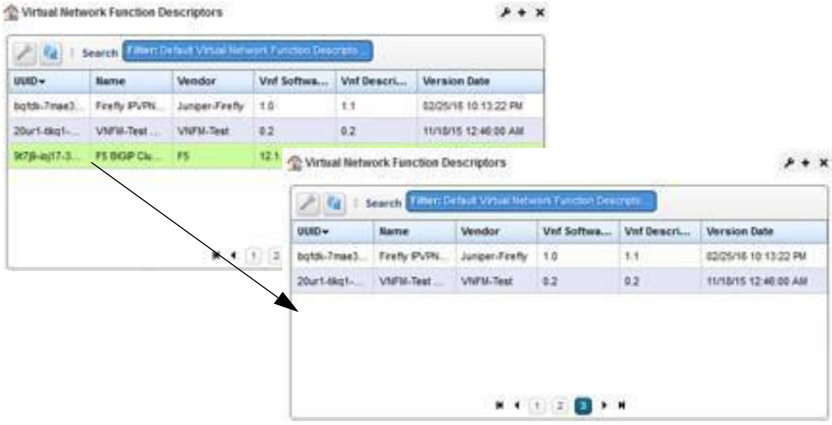

Virtual Network Function Descriptors
Search
Filter: Default Virtual Network Function Descripto...
Virtual Network Function Descriptors
Search
Filter: Default Virtual Network Function Descripto...

UUID▾	Name	Vendor	Vnf Softwa...	Vnf Descri...	Version Date
bqtdk-7mae3...	Firefly IPVPN...	Juniper-Firefly	1.0	1.1	02/25/16 10:13:22 PM
20ur1-tkq1-...	VNFM-Test ...	VNFM-Test	0.2	0.2	11/18/15 12:46:00 AM
9t7j9-iej17-3...	F5 BIGP Clu...	F5	12.1		

UUID▾	Name	Vendor	Vnf Softwa...	Vnf Descri...	Version Date
bqtdk-7mae3...	Firefly IPVPN...	Juniper-Firefly	1.0	1.1	02/25/16 10:13:22 PM
20ur1-tkq1-...	VNFM-Test ...	VNFM-Test	0.2	0.2	11/18/15 12:46:00 AM

Maintaining PNF Descriptors

This section covers the following basic physical network function (PNF) descriptor maintenance tasks:

- Creating a PNF Descriptor
- Importing a PNF Descriptor
- Modifying a PNF Descriptor
- Deleting a PNF Descriptor

For a description of all the available options to monitor and maintain descriptors, see Physical Network Function Portlets on page 898.

Creating a PNFDescriptor

Create a physical network function (PNF) from the Physical Network Function portlet as follows.

1 Right-click > New.

The Editing Physical Network Function Descriptor window is displayed.

2 Enter a name.

3 Select Enabled.

4 Click Apply.

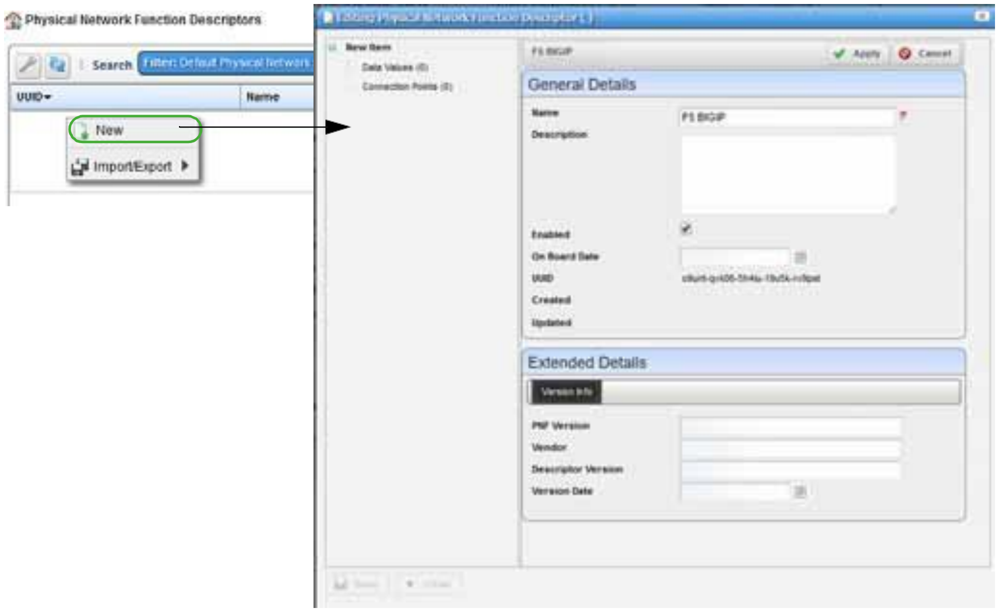

F5 BIGIP

Name: F5 BIGIP

Description:

Enabled: ☑

On Board Date:

UUID: c8urt-qn406-5h4a-1bd5k-n9pat

Created:

Updated:

PNF Version:

Vendor:

Descriptor Version:

Version Date:

5 Click Connection Points.

Notice that no items are listed.

6 Click Add.

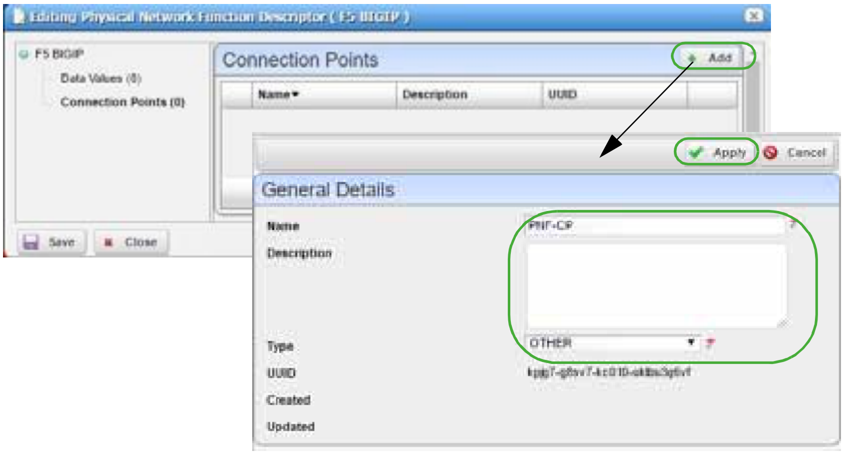

Name: PNF-CP

Description:

Type: OTHER

UUID: kpjg7-g8sv7-kc010-okbu3g6vf

Created

Updated

7 Enter a name and a type.

8 Click Apply.

The connection point is added to the list.

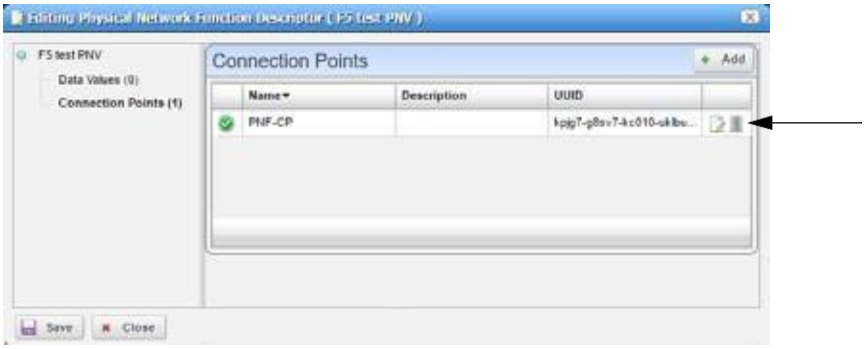

	Name	Description	UUID	
✓	PNF-CP		kpjg7-g8sv7-kc010-uklbu...	

9 Save your new descriptor.

Importing a PNF Descriptor

You can import a physical network function (PNF) descriptor from another system or from a backup that was exported to a XML file or you can import a descriptor from a URL. The example provided here imports a descriptor from an XMLfile.

Import a physical network function (PNF) from the Physical Network Function portlet as follows.

1 Right-click > Import/Export. > Import.

The Import Physical Network Descriptor(s) window is displayed.

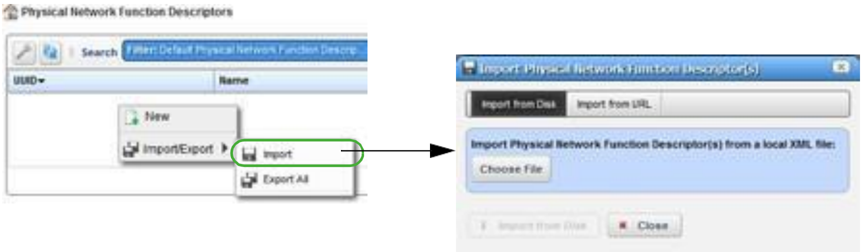
Physical Network Function Descriptors
Search
Filter: Default Physical Network Function Descrip...
UUID
Name
New
Import/Export
Import
Export All
Import Physical Network Function Descriptor(s)
Import from Disk
Import from URL
Import Physical Network Function Descriptor(s) from a local XML file:
Choose File
Import from Disk
Close

2 Click Choose File.

The navigation window is displayed.

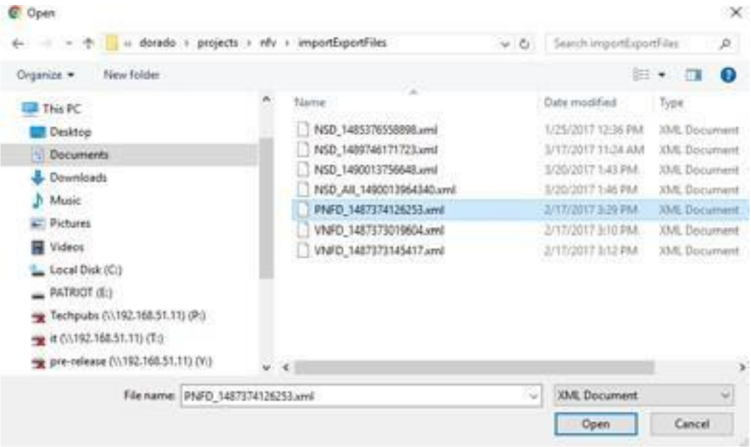

Name	Date modified	Type
NSD_1485376558898.xml	1/25/2017 12:36 PM	XML Document
NSD_1489746171723.xml	3/17/2017 11:24 AM	XML Document
NSD_1490013756648.xml	3/20/2017 1:43 PM	XML Document
NSD_All_1490013964340.xml	3/20/2017 1:46 PM	XML Document
PNFD_1487374126253.xml	2/17/2017 3:29 PM	XML Document
VNFD_1487373019604.xml	2/17/2017 3:10 PM	XML Document
VNFD_1487373145417.xml	2/17/2017 3:12 PM	XML Document

File name: PNFD_1487374126253.xml

XML Document

3 Navigate to the XML file.

4 Select the file to import.

For example: `PNFD_1487374126253.xml`

5 Click Open.

The Import Physical Network Function Descript or(s) Import from Disk button is activated.

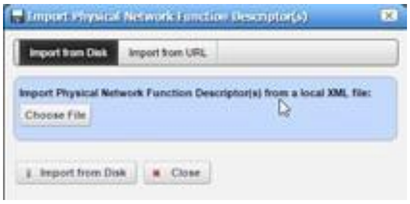

Import Physical Network Function Descriptor(s) from a local XML file:

6 Click Import from Disk.

The selected descriptor is added to the PNF list.

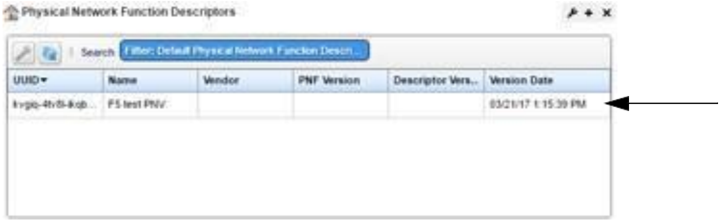
Physical Network Function Descriptors
Search
Filter: Default Physical Network Function Descri...

UUID▾	Name	Vendor	PNF Version	Descriptor Vers...	Version Date
kvgio-4tv8l-ikqb...	F5 test PNV				03/21/17 1:15:39 PM

Modifying a PNF Descriptor

Modify a physical network function (PNF) from the Physical Network Function portlet as follows.

1 Right-click *descriptor* > Edit.

The Editing Physical Network Descriptor window is displayed.

2 Click Edit.

3 Make any necessary changes.

4 Save your modifications.

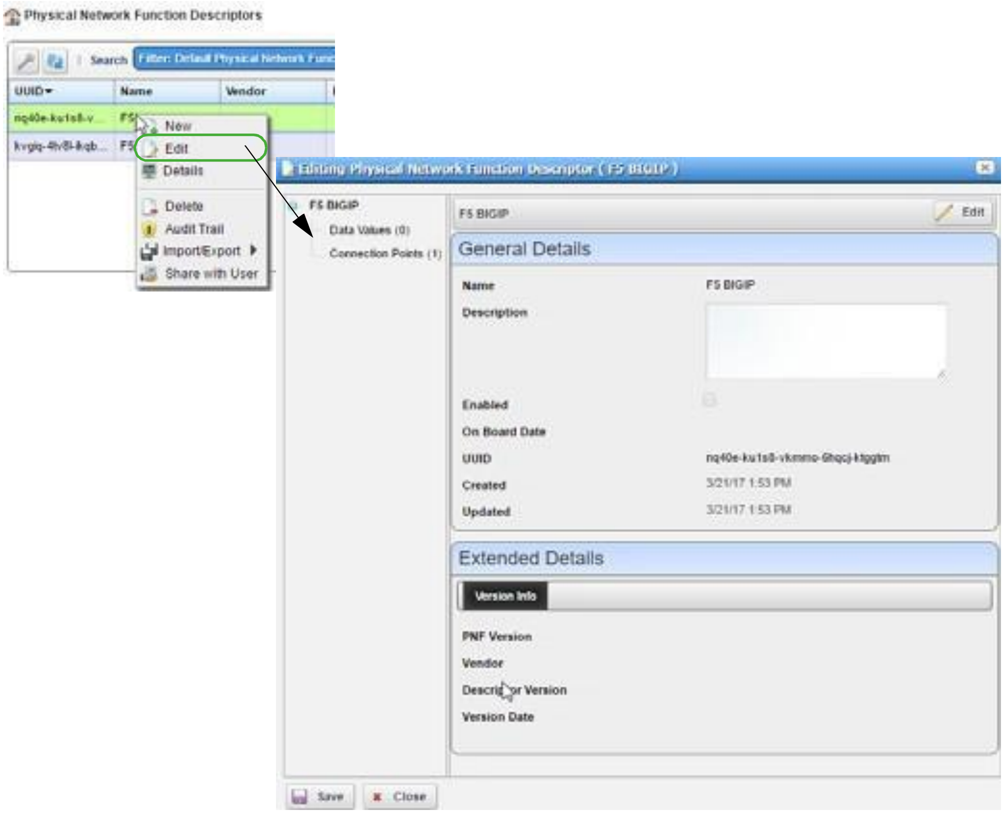

Name F5 BIGIP
Description
Enabled □
On Board Date
UUID nq40e-ku1s8-vkmmo-6hqcj-ktggtm
Created 3/21/17 1:53 PM
Updated 3/21/17 1:53 PM

PNF Version
Vendor
Descriptor Version
Version Date

Deleting a PNF Descriptor

Delete a physical network function (PNF) from the Physical Network Function portlet as follows.

1 Right-click *descriptor* > Delete.

A confirmation message is displayed.

2 Click Delete.

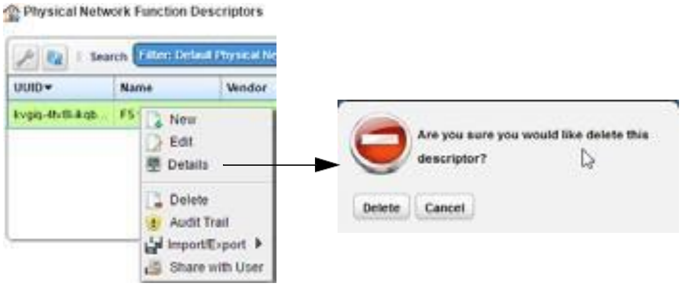

Are you sure you would like delete this descriptor?

3 Verify that the descriptor is no longer listed.

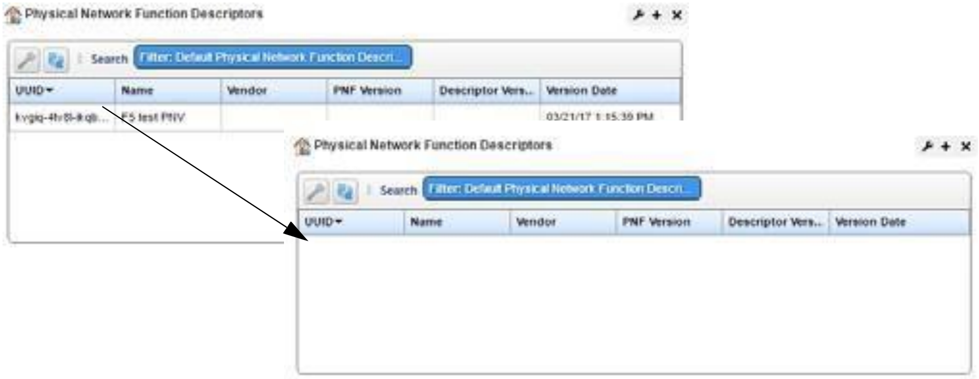
Physical Network Function Descriptors
Search
Filter: Default Physical Network Function Descri...
Physical Network Function Descriptors
Search
Filter: Default Physical Network Function Descri...

UUID	Name	Vendor	PNF Version	Descriptor Vers...	Version Date
kvgiq-4tv8l-ikqb...	F5 test PNV				03/21/17 1:15:38 PM

UUID	Name	Vendor	PNF Version	Descriptor Vers...	Version Date

Automating SMB vCPE Deployment

Full deployment automation of an SMB v CPE service including VPN and internet access. For this example, the provider’s order management (OM) system manages the customer catalog service and coordinates with the Dorado Software Network Functions Virtualization Orchestration/ Virtualized Network Function Manager (NFVO/VNFM) for full automated instantiation of the virtualized network functions.

- Current State
- Order Management
- Service Orchestration
- Final State
- NFV Catalogs

Current State

Thecurrent Network Functions Virtualization Infrastructure (NFVI) Pointof Presence (POP) is configured and has one multi-vendor Virtualized Network Function (VNF) deployed.

The system provides a wide arrayof ongoing management tools to access the health and configuration of the network function and environment, such as:

- Topology
- Resource Management
- Performance Monitoring
- Infrastructure Capacity

These portals indicate the current NFVI POP state where the virtual customer premise equipment (vCPE) is to be deployed.

Order Management

Your order management system and the Dorado Software Management and Orchestration (MANO) product use a REST interface to provide a seamless service deployment.

Catalogs defined in both systems provide the platform for top-to-bottom automated order management.

Key elements include:

- Product Catalogs
- Service Templates
- Network Service Descriptors
- Virtualized Network Function Descriptors
- Current Network Capability (VNFR)

Service Orchestration

Create and deploy a customer vCPE service. Trac k a service from initial order entry through final function configuration.

Key tasks include:

- Create a Customer Record
- Select a Service
- Complete Service Order
- Submit the Order
- Instantiate virtual components
- Configure customer attachment
- Configure virtual functions
- Monitor Progress

Final State

Upon completion of vCPE instantiation, the management tools reflect the deployed VNFs. The topology reflects the positioning of the new vCPE and VNF instances. Resource Management indicatestheadditional VNFsunder management. Infrastructure capacity reflects the resources utilized for thescaled capability.

NFV Catalogs

The NFV catalogs contain the building blocks (Figure 16-6) that the orchestrators use to manage services, functions, and infrastructure for virtual network ecosystems.

Managed elements include:

- VNF Descriptor
- NS Descriptor
- VNF Images
- VIM Images
- VIM Capacity
- NS Records

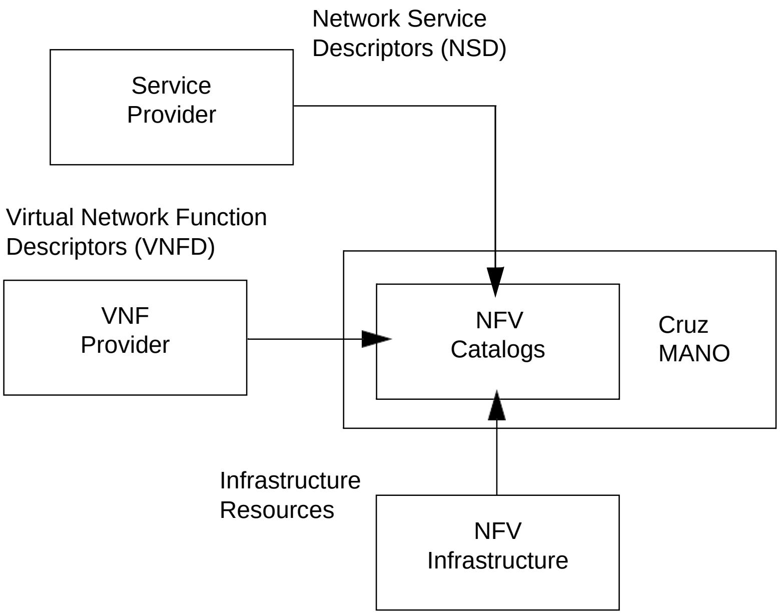

Figure 16-6. Orchestration Building Blocks

Managing Network Services(NS) and Virtualized Network Function(VNF) descriptorseffectively as the Network Functions Virtualization (NFV) standards, functions, and processes mature is essential to allow both service and network function providers to recognize benefits.

F5 LTM Onboarding
Review the full process of instantiation of a VNF from onboarding through function turnup.

VIM Management
The application manages each VIM as a resource. This allows all the Cruz core management processes to be applied to the VIM resources.

Additionally, NFV functions are supported, such as reservation.

Image Management
The Cruz NFV functions provide image management capabilities to support both the demand of production and lab environments.

REST API
The Cruz application provides REST access to both its NFV capabilities and its entire suite of service management features.

17

Virtualization Management

As virtualized network functionality becomes central to a service provider's business model, key Network Functions Virtualization (NFV) capabilities need to be brought under management.

With the Cruz NFV applications and the third-party OpenStack ® operating system, service providers can quickly bring virtual network functions (VNFs) under full resource management.

This section is intended foroperations technicians, systemadministrators, networkadministrators, or anyone using Cruz and OpenStack infrastructure to perform daily operations, such as stage, deploy, and undeploy virtual network functions (VNFs).

This section covers the following topics related to the Cruz VNF daily operations, such as stage, deploy, and undeploy:

Getting Started

Before you get started managing virtualized network functions (VNFs) from the Cruz application, you should know how to start the application, set up your environment, and test your setup.

- Signing In/Out
- Setting Up Cruz Application
- Testing Your Cruz Setup

Signing In/Out

To get started using the Cruz Network Functions Virtualization (NFV) features, you need to know how to sign into the Cruz application and sign out when finished.

Sign into your Cruz NFV installation from a Web browser as follows and then log off when finished.

1 Enter the Cruz URL.

```
http://appServerHost:portNumber
```

Replace *appServerHost* with your Cruz NFV server host name or IP address. The default *portNumber* is 8080.

The sign in request is displayed.

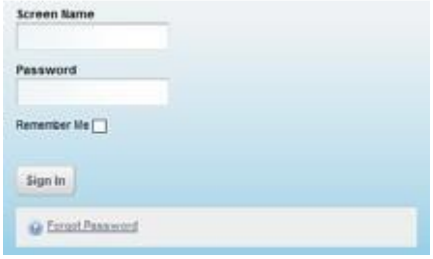

2 Enter your screen name (username) and password.

3 Click Sign In.

The Cruz NFV Home page is displayed.

If you entered an incorrect usernameor password, a message is displayed. Enter the information again.

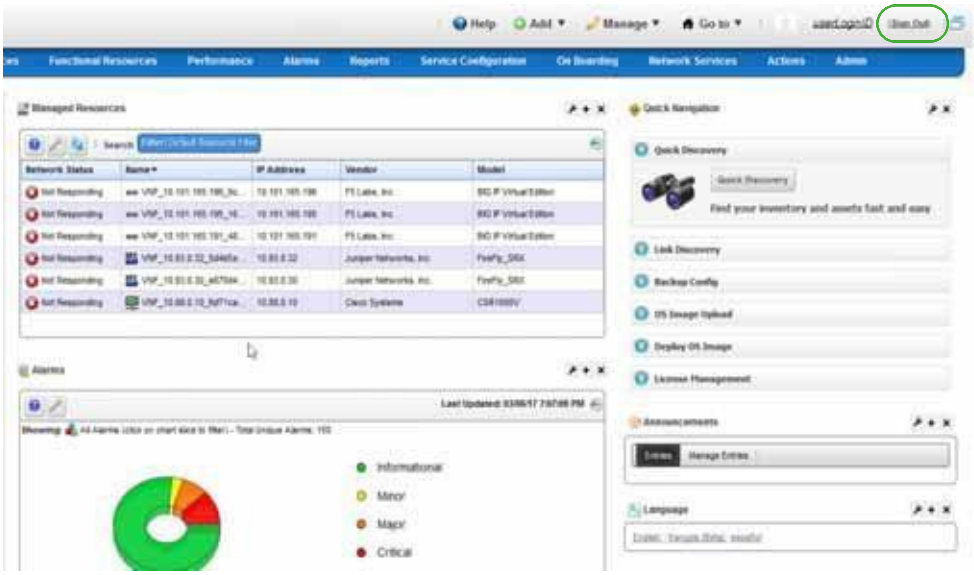

Network Status	Name	IP Address	Vendor	Model
Not Responding	VNF_10.101.165.198_9c...	10.101.165.198	F5 Labs, Inc.	BIG IP Virtual Edition
Not Responding	VNF_10.101.165.195_16...	10.101.165.195	F5 Labs, Inc.	BIG IP Virtual Edition
Not Responding	VNF_10.101.165.191_40...	10.101.165.191	F5 Labs, Inc.	BIG IP Virtual Edition
Not Responding	VNF_10.93.0.32_5d4e5a...	10.93.0.32	Juniper Networks, Inc.	FireFly_SRX
Not Responding	VNF_10.93.0.30_a67b94...	10.93.0.30	Juniper Networks, Inc.	FireFly_SRX
Not Responding	VNF_10.88.0.10_4d71ca...	10.88.0.10	Cisco Systems	CSR1000V

Last Updated: 03/06/17 7:07:08 PM

Showing: All Alarms (click on chart slice to filter) – Total Unique Alarms: 153

4 Click Sign Out when finished with the application.

You are returned to the sign in request page.

Setting Up CruzApplication

Before you perform the tasks described in this guide, make sure that the NFV -specific portlets you need are added to the existing Cruz portal.

Which portlets you need, de pends on your role (Table 17-1) in the management of your virtualized network environment and the deployment of network services and virtualized network functions (VNFs).

Table 17-1. User Roles and Tasks

Role	Tasks
Administrator	Bringing VIM under management (creating the controller record) Setting up the OpenStack environment (uploading vendor base images to the OpenStack controller, modifying resource flavors, creating VIM snapshots) Preparing vendor-supplied images from Cruz (creating and editing VIM images) Maintaining the VIMs under management, descriptors, virtual reservations, and virtual requirements
Network Designer/ Engineer	Defining descriptors Implementing virtual reservations Monitoring virtual requirements
Operator	Instantiating VNFs (stage, deploy) Monitoring virtual requirements Performing other daily operations, such as undeploying, scaling, monitoring system health, resolving issues found, *etc.*

Depending on your permissions and how your company wants to define the work environment for all users, you have the option to add the Network Virtualization applications (portlets) to existing Cruz pages or create pages and add portlets more appropriately for each users work environment. Once the pages/portlets setup is complete, verify this setup.

Here is a list of NFV portlets (Table 17-2) and how each portlet is used by the different user roles. For a detailed description of these portlets, see Network Virtualization Portlets on page 866. If any of these portlets are **not** available, a message is displayed when you try to add it to a page.

Table 17-2. NFV Portlets (Sheet 1 of 2)

Portlet	Operator	Administrator	Network Designer
License Accounts	Maintain a list of licensed accounts	Maintain a list of licensed accounts	Maintain a list of licensed accounts
License Descriptors	View and understand license descriptors	Maintain available license descriptors	Define licensing integration and characteristics
License Records	View and understand license records	View and understand license records	View and understand license records
NFV Monitoring Attributes	View and understand NFV monitoring attributes	Maintain monitoring attributes	Maintain monitoring attributes
Network Service Descriptors	View and understand NSDs	Maintain network service descriptors	Define monitoring attribute characteristics
Network Service Records	Stage, deploy, and undeploy (manage) network services	No activity performed by the administrator	Test an implemented network service
OSS Instance	View and understand OSS instances	Implement OSS instances	No activity performed by the network designer
Physical Network Function Descriptors	View and understand PNFDs	Maintain PNF descriptors	Implement physical network functionality
Physical Network Function Records	Stage, deploy, and undeploy (manage) PNFs	Test new or modified descriptors	Test implemented PNFs
SDN Controllers	View and understand SDN controllers	Maintain software-defined network (SDN) controllers.	Test SDN controllers
Software Images	View and understand software images	Create a software image descriptor for a snapshot image, deploy software images, and edit the VIM software descriptor parameters	Create a software image descriptor for a snapshot image, deploy software images, and edit the VIM software descriptor parameters
VIM Images	View and understand VIM images	Maintain consistency across the VIM instances	Test that software images deployed/undeployed were added/removed from the VIM images list
Virtual Network Function Descriptors	View and understand VNFDs	Maintain VNF descriptors	Implement virtualized network functions
Virtual Network Function Records	Stage, deploy, and undeploy (manage) VNFs	No activity performed by the administrator	Test an implemented network service

Table 17-2. NFV Portlets (Sheet 2 of 2)

Portlet	Operator	Administrator	Network Designer
Virtual Requirements	View virtual domain resource requirements and usage (such as memory, CPU, and disk). You also have the option to modify a domain's description.		
Virtual Reservations	Verify that a VDU's resources were reserved after they stage a VNF record	Maintain virtual reservations	Implement virtual reservations
Virtualized Infrastructure Managers	View available VIMs and their resources before deploying services or VNFs	Maintain the VIMs under resource management	No activity performed by the network designer

Adding a Page

Pages allow you to organize commonly used portlets into a single location for easy access or to group portlets by t asks you perform.For example, group tasks used to monitor and resolve network health on a single page. You have options to add NFV pages to the menu bar, add child pages to existing pages, or a combination of both.

Add any needed pages to the Cruz application as follows.

1 Select Manage > Page.

The Manage Page window is displayed.

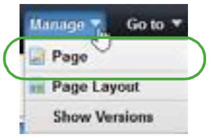

2 Add pages you want to access from the navigation panel.

a. Select Private Pages.

b. Specify the look and feel.

c. Click Add Page.

The Add Page window is displayed.

d. Enter a name for the page.

e. Select a type.

f. Click Add Page.

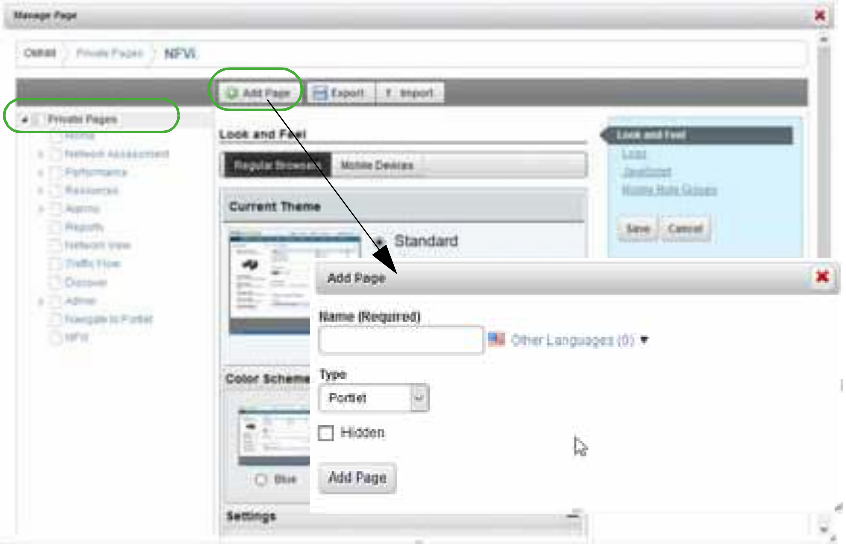

3 Repeat steps 1 and 2 for each page you add to the navigationpanel.

4 Add a child page to an existing page.

a. Select Manage > Page.

The Manage Page window is displayed.

b. Select the page under which to add the child page.

c. Click Add Child Page.

The Add Child Page window is displayed.

d. Enter a name for the child page.

e. Select the type.

f. Select whether to copy the parent or make the page hidden.

g. Click Add Page.

h. Repeat these steps for each child page you want to add.

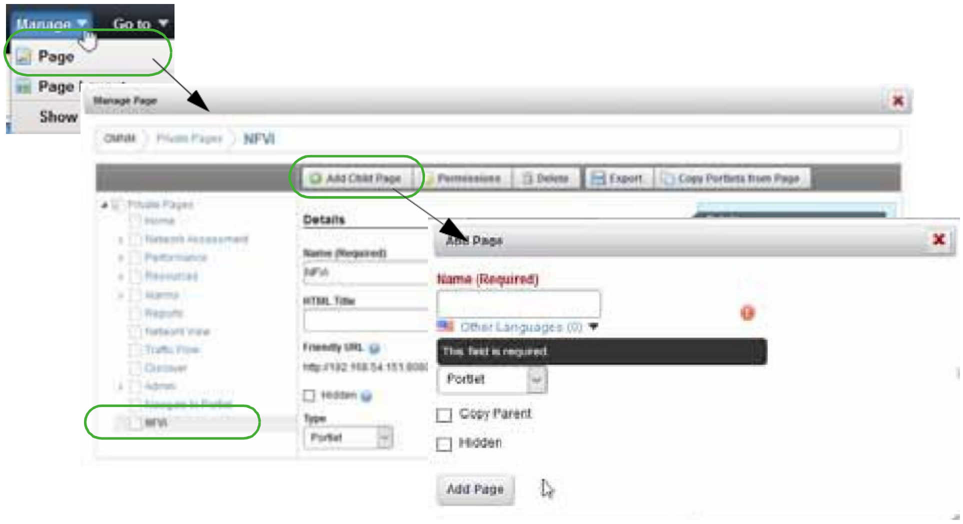

Now you are ready to add portlets (applications) to t he appropriate pages.

Adding Portlets to a Page

These instructions assume that the pages to which you plan to ad d portlets already exist. I f the pages do not exist, see Adding a Page before continuing.

Add portlets to a p age from the Cruz application as follows.

1 Click the page label to which you want to add a portlet.

The selected page is displayed.

2 Select Add > Applications.

The applications list is displayed.

3 Expand the Network Virtualization list.

See Setting Up Cruz Application on page 855 for a brief description of each portlet, or Network Virtualization Portlets on page 866 more details.

4 Click Add for each portlet you want to add to the selectedpage.

Note that the purple indicator means that you can add the portlet only once to a page. All other portlets you can add multiple instances to a page.

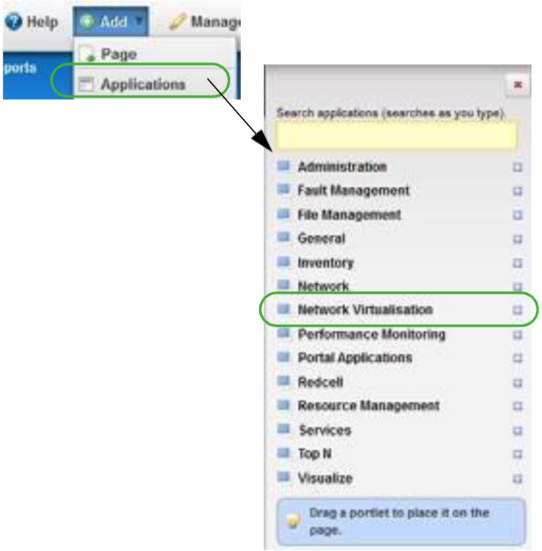

5 Refresh the page to show the portlet.

If you do not have permissions to view the portlet, a message is displayed. Delete the portlet or consult your system administrator if you feel you should have access to the portlet.

6 Repeat steps 4 and 5 for each portlet you want to add to the selected page.

7 Rearrange the portlets as needed using the drag-and-drop method.

8 Repeat steps 1 through 7 for each page to which you want to add portlets.

Testing Your Cruz Setup

It does not matter if you set up your Cruz Network Functional Virtualized (NFV) work environment (pages/portlets) or if it was defined by your system administrator , you should run **some** tests to make sure you have what you need and you can perform some basic tasks on a record, such as discover, stage, deploy, modify, undeploy, and delete. This test example uses an NFV record and assumes that you have access to the Cruz application and your OpenStack project.

Test your NFV setup from the Cruz application as follows.

1 Sign out and then sign back in if you are already logged in.

2 Verify that your pages are accessible from the menu bar.

3 Make any necessary changes.

4 Navigate to the Virtual Network Function Records portlet.

This is where you will complete steps 5 through 10.

5 Verify that you can **discover** VNF records.

a. Right-click > Discover/Resync.

The Discover VNF Records window is displayed.

b. Select a VIM entity.

The VNF Discovery parameters are displayed.

c. Specify namespaces.

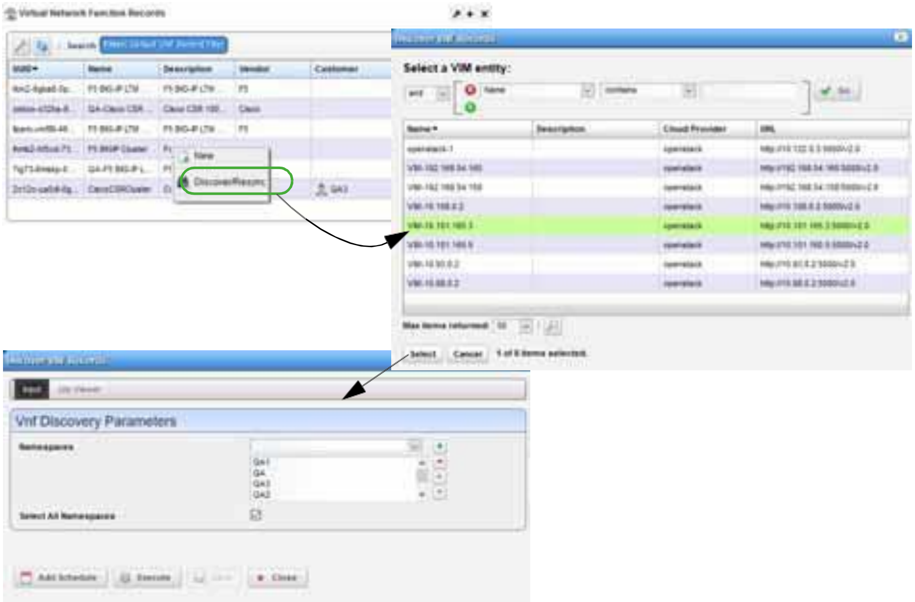

d. Click Execute.

The Job Viewer displays the execution progress, any errors, and success/failure, followed by the Results information.

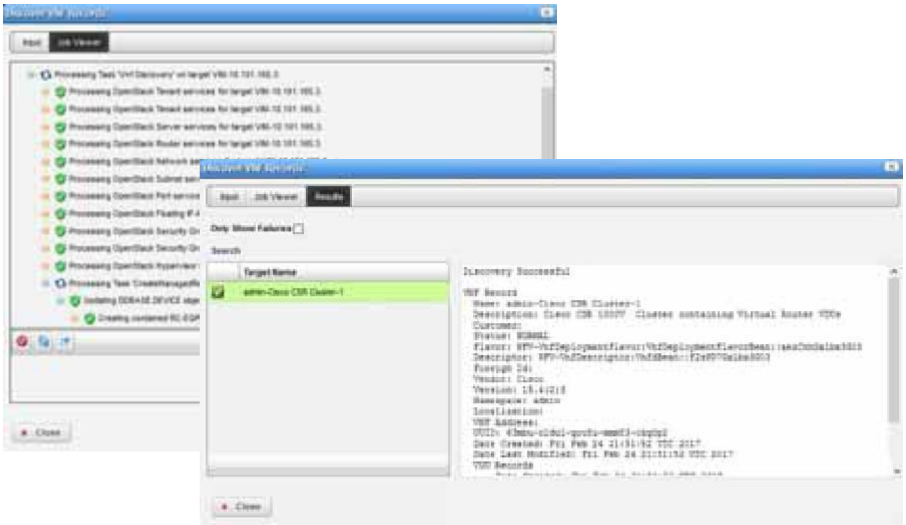

```
Processing Task 'Vnf Discovery' on target VIM-10.101.165.3
  Processing OpenStack Tenant services for target VIM-10.101.165.3
  Processing OpenStack Tenant services for target VIM-10.101.165.3
  Processing OpenStack Server services for target VIM-10.101.165.3
  Processing OpenStack Router services for target VIM-10.101.165.3
  Processing OpenStack Network se[illegible]
  Processing OpenStack Subnet se[illegible]
  Processing OpenStack Port service[illegible]
  Processing OpenStack Floating IP [illegible]
  Processing OpenStack Security Gr[illegible]
  Processing OpenStack Security Gr[illegible]
  Processing OpenStack Hypervisor[illegible]
  Processing Task 'CreateManagedR[illegible]
    Updating DDBASE DEVICE obje[illegible]
      Creating contained RC-EQP[illegible]
```

☑ admin-Cisco CSR Cluster-1

```
Discovery Successful

VNF Record
  Name: admin-Cisco CSR Cluster-1
  Description: Cisco CSR 1000V  Cluster containing Virtual Router VDUs
  Customer:
  Status: NORMAL
  Flavor: NFV-VnfDeploymentFlavor:VnfDeploymentFlavorBean::[illegible]
  Descriptor: NFV-VnfDescriptor:VnfdBean::[illegible]
  Foreign Id:
  Vendor: Cisco
  Version: 15.4(2)S
  Namespace: admin
  Localization:
  VNF Address:
  UUID: [illegible]
  Date Created: Fri Feb 24 21:31:52 UTC 2017
  Date Last Modified: Fri Feb 24 21:31:52 UTC 2017
  VDU Records
```

e. Close the Discover VNF Records window.

The discovery/resync process makes sure th at you have any changes from the target VIM (Virtualized Infrastructure Manager).

6 Verify that you can **stage** a record.

a. Select a record.

a. Right-click > New.

The Stage VNF window displays a list of descriptor entities.

b. Select the service you want to instantiate.

c. Click Select.

The Vnf Stage Parameters window is displayed.

d. Enter a name for the VNF and a description that includes specifics on the services.

e. Specify the parameters, such as the VIM on which to deploy.

The Name, VIM, Namespace, and Flavor fields are required.

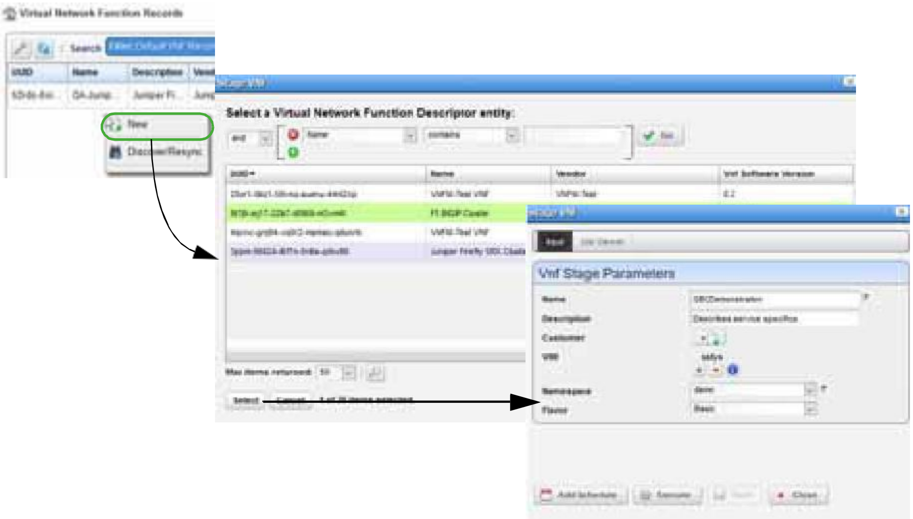

f. Click Execute.

The Job Viewer displays the execution progress, any errors, and success/failure. Upon completion, the stage results are displayed.

If staging is successful, the record is created and staged based on the descriptor information found and then resources are reserved. However, nothing is deployed to the OpenStack controller.

If staging fails, the reason for the failure is displayed.

g. Close the Stage VNF window.

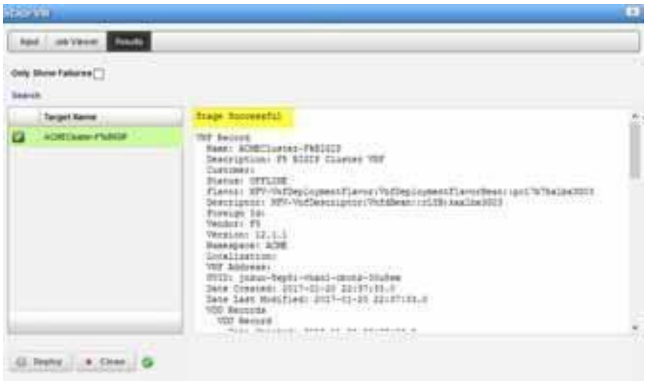

```
Stage Successful
VNF Record
  Name: ACMECluster-F5BIGIP
  Description: F5 BIGIP Cluster VNF
  Customer:
  Status: OFFLINE
  Flavor: NFV-VnfDeploymentFlavor:VnfDeploymentFlavorBean:[illegible]
  Descriptor: NFV-VnfDescriptor:VnfdBean:[illegible]
  Foreign Id:
  Vendor: F5
  Version: 12.1.1
  Namespace: ACME
  Localization:
  VNF Address:
  UUID: [illegible]
  Date Created: 2017-01-20 22:37:[illegible]
  Date last Modified: 2017-01-20 22:37:[illegible]
  VDU Records
    VDU Record
```

7 Verify that you can **deploy** the record successfully staged in step 6.

a. Click Deploy.

The default VNF deploy parameter values are displayed.

b. Change the default parameters as needed.

c. Click Execute.

The Job Viewer displays the steps taken and the results.

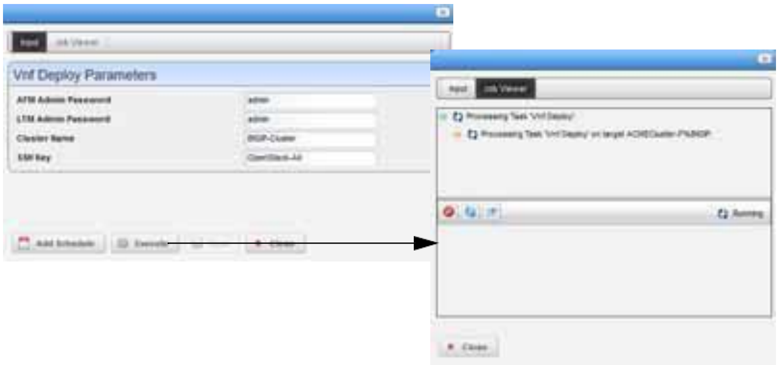

Vnf Deploy Parameters

ATM Admin Password	admin
LTM Admin Password	admin
Cluster Name	BIGP-Cluster
ESM Key	OpenStack-All

Processing Task 'Vnf Deploy'

Processing Task 'Vnf Deploy' on target ACMECluster-F5BIGP

d. Close the Job Viewer.

e. Verify that the record status changed to Normal.

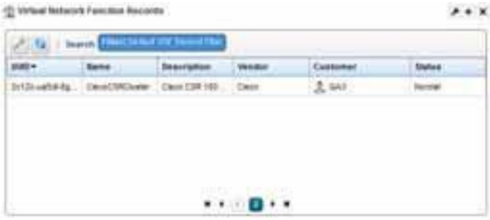

Virtual Network Function Records

VNFD ID	Name	Description	Vendor	Customer	Status
2c12c-ua5d4-6g...	CiscoCSRCluster	Cisco CSR 100...	Cisco	SAJ	Normal

f. Go to your OpenStack dashboard for your project.

g. Click Instances.

A list of instances and their information is displayed.

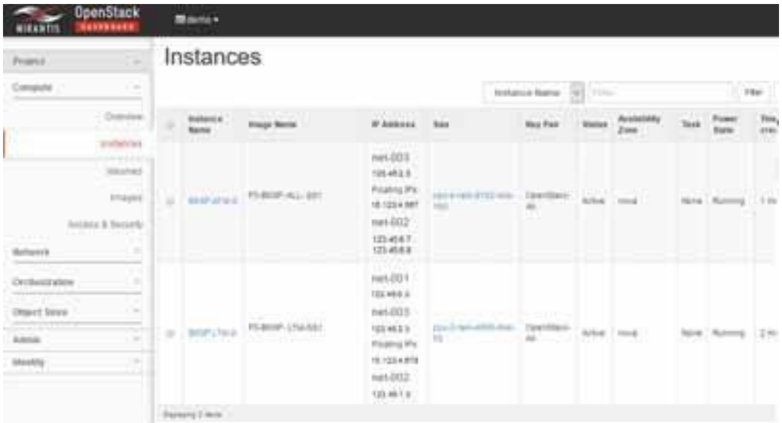

Instances

	Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created
□	[illegible]	F5-BIGIP-ALL-[illegible]	net-003 [illegible] Floating IPs: [illegible] net-002 [illegible]	[illegible]	OpenStack-All	Active	nova	None	Running	[illegible]
□	[illegible]	F5-BIGIP-LTM-[illegible]	net-001 [illegible] net-003 [illegible] Floating IPs: [illegible] net-002 [illegible]	[illegible]	OpenStack-All	Active	nova	None	Running	[illegible]

h. Verify that the resource record is listed.

This indicates connectivity to your target VIM.

i. Return to the Cruz application.

8 Verify that you can **modify** the record staged in step 6.

a. Select a record.

b. Right-click > Edit.

The Editing VNF Record window is displayed.

c. Click Edit.

d. Modify the name or description.

e. Click Apply.

f. Click Save.

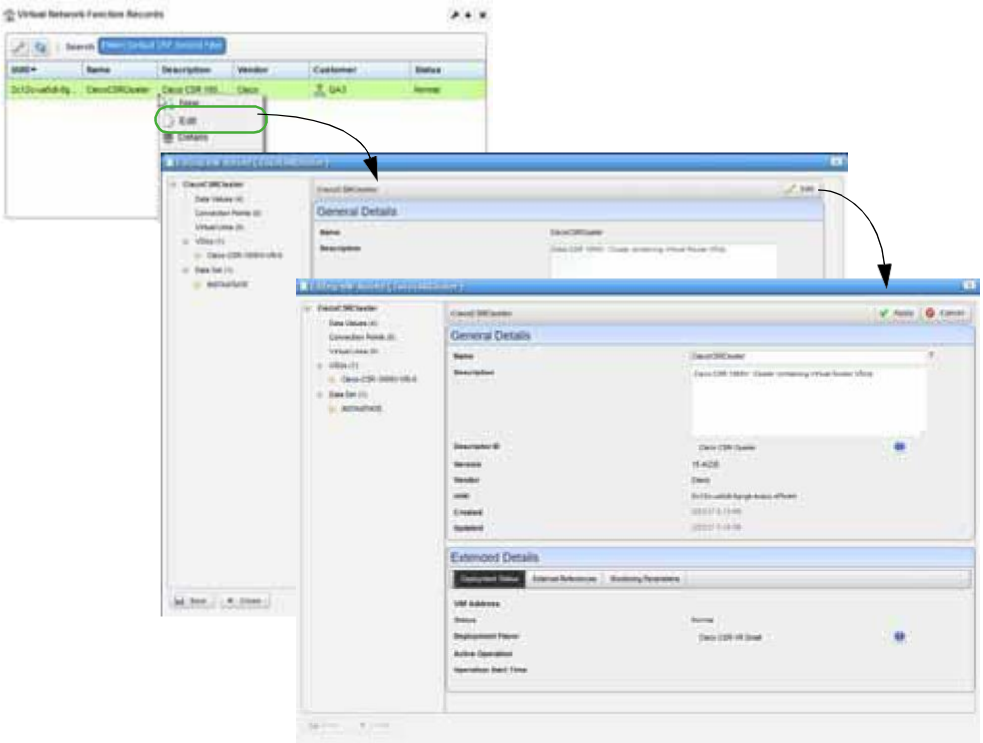

9 Verify that you can **undeploy** the record deployed in step 7.

a. Select a record.

b. Right-click > Undeploy.

A confirmation message is displayed.

c. Click Undeploy.

The Job Viewer displays the steps taken and the results.

d. Click Close.

e. Verify the record’s status is now Offline.

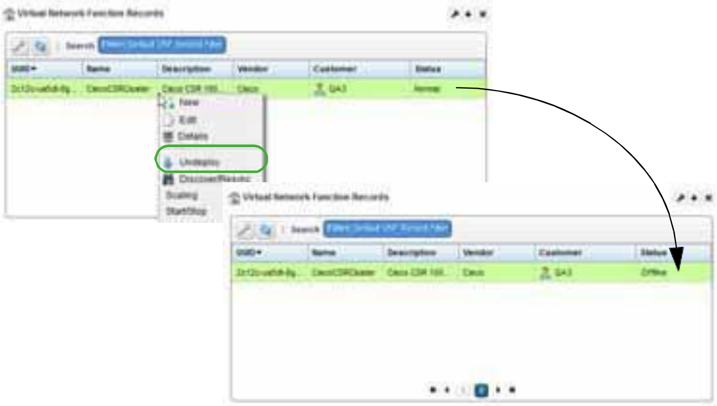

10 Verify that you can **delete** a record.

a. Select a record.

b. Right-click > Edit.

A confirmation message is displayed.

c. Click Delete.

The Job Viewer displays the steps taken and the results.

d. Close the Job Viewer.

e. Verify that the record is no longer listed in the Virtual Network Function Records portlet.

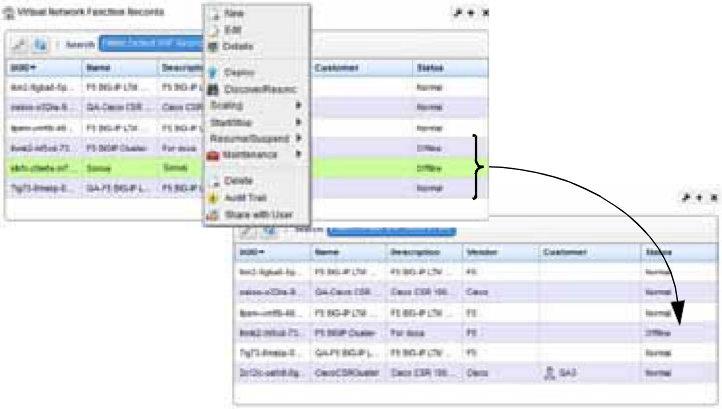

Network Virtualization Portlets

This section describes the portlets, windows, menu options, and other aspects for the Network Virtualization (NV) user interface.

- Image Portlets
- Network Service Portlets
- NFV Monitoring Attributes Portlet
- OSS Instance Portlet
- Physical Network Function Portlets
- SDN Controllers Portlet
- Vendor Licensing Portlets
- Virtualized Network Function Portlets
- Virtual Requirements Portlet
- Virtual Reservations Portlet
- Virtualized Infrastructure Managers Portlet

To understand the common portlet behavior, options, and so on, see General Portlet Information on page 127.

Image Portlets

Use the following image portlets to view, create, and manage image descriptors that are used to manage, deploy, or migrate images:

- Software Images Portlet
- VIM Images Portlet

Software Images Portlet

Use the Software Images portlet (Figure 17-1) to viewand manage software image descriptors. The software image descriptors are used to manage, deploy, or migrateimages.

From the Software Images portlet, a system administrator creates and maintains the image snapshots used to create descriptors, a network designer could also create and maintain image snapshots, and operators (if they have the proper permissions) view descriptors to gain knowledge.

The Software Images portlet also has a maximized view. Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

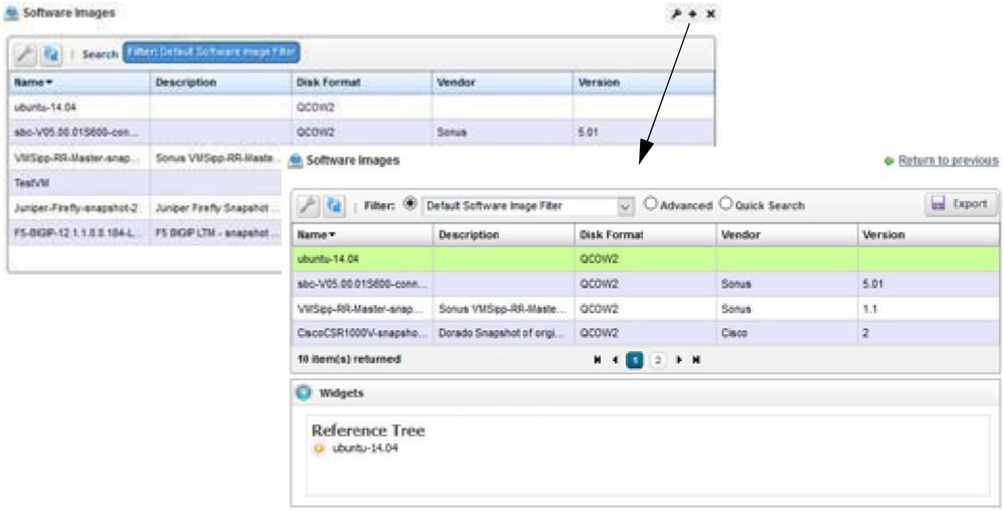

Software Images

Search Filter: Default Software Image Filter

Name	Description	Disk Format	Vendor	Version
ubuntu-14.04		QCOW2		
sbc-V05.00.01S600-con...		QCOW2	Sonus	5.01
VMSipp-RR-Master-snap...	Sonus VMSipp-RR-Maste...			
TestVM				
Juniper-Firefly-snapshot-2	Juniper Firefly Snapshot ...			
F5-BIGIP-12.1.1.0.0.184-L...	F5 BIGIP LTM - snapshot ...			

Software Images

Return to previous

Filter: ● Default Software Image Filter ○ Advanced ○ Quick Search

Name	Description	Disk Format	Vendor	Version
ubuntu-14.04		QCOW2		
sbc-V05.00.01S600-conn...		QCOW2	Sonus	5.01
VMSipp-RR-Master-snap...	Sonus VMSipp-RR-Maste...	QCOW2	Sonus	1.1
CiscoCSR1000V-snapsho...	Dorado Snapshot of origi...	QCOW2	Cisco	2

10 item(s) returned

Widgets

Reference Tree

ubuntu-14.04

Figure 17-1. Software Images Portlet

Pop-Up Menu

The Software Images pop-up menu provides access to the following options. Right-click a row to access these options (Table 17-3).

Table 17-3. Software Images Pop-Up Menu Options

Menu Option	Description
New	Opens the Creating new Software Image window, where you specify the details used to create an image. You can also create an image using the Manage option.
Edit	Opens the Editing Software Image window, where you modify the general image details or simply view the image details. Click Vim Images > Edit icon to view the VIM image details or select whether the image is a master copy. By default, Master Copy is not selected.
Details	Displays the general details provided when you created the image, a reference tree, alarms, and event history details for the selected image. Click Return to previous to return to the Software Images portlet.
Manage	Opens the Manage new Software Image from VIM Parameters window , where you create a software image so it is available from the catalog to deploy to other VIMs. This method of creating a software image is useful when you are not sure of all the parameters needed when creating an image using the New option.
Download	Opens the Save VIM Image to Disk Parameters window, where you specify where to store the image backup and then execute the download.
Deploy	Pushes the software image out to VIMs under management so that the same image is available across all your VIMs or datacenters. A deployed image can then be turned up (activated and provisioned) for a specific virtualization deployment unit (VDU) within a virtualized network function (VNF).
Undeploy	Removes the selected image’s availability from your VIMS or datacenters. Some reasons for undeploying an image that is no longer needed are the service is no longer offered or perhaps the VDU was updated to a newer image and the only image is no longer needed. **Caution:** This also removes the OpenStack project image instance, which can cause s ues is with anything using that image.

Table 17-3. Software Images Pop-Up Menu Options

Menu Option	Description
Delete	Removes the selected images from the system. **Caution:** Deleting an image can cause anything using that image to fail.
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected image. Select an audit record and the job record displays its status and any other job-related information.
Import/Export	Provides the following actions when available for the selected image: • Import retrieves a file containing XML image descriptions. Some imports can come from a URL. • Export Selection exports the selected image description to an XML file. • Export All exports all image descriptions to an XML file. Click Download Export File to specify where to save the file. The Import/Export option is useful as a backup or to share descriptors with other projects.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection.

Columns

Other than the general navigationand configuration options, the Software Images portlet includes the following columns/fields (Table 17-4).

You can view the value for most of the hidden columns by right clicking a row and selecting Details. The other option is to add the column byclicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-4. Software Images Columns (Sheet 1 of 2)

Column	Description
Name	The user-specified software image identifier.
Description	The optional details about the software image, such as its purpose. For example, mentions customizations made for your use and the original vendor-provided image from which you created the software image.
Disk Format	The file format in which the image is to be created, such as QCOW2. Valid formats are AKI, AMI, ARI, ISO, QCOW2, RAW, UNRECOGNIZED, VDI, VHD, VMDK.
Vendor	The vendor’s name that supplied the initial image used to create the software image descriptor.
Version	The vendor-supplied image version.
Boot Wait Time	The number of seconds to wait for an image to finish booting. By default, this field is not shown.
Checksum	The data used to detect errors introduced during transmission or storage. By default, this field is not shown.
Cloud Provider	The target locations where you plan to deploy the image. By default, this field is not shown.

Table 17-4. Software Images Columns (Sheet 2 of 2)

Column	Description
Container Format	The target format to which you plan to deploy the image. For example, a target Bare Metal (BARE) OpenStack controller. Valid values are AKI, AMI, ARI, BARE, DOCKER, OVF, and UNRECOGNIZED. By default, this field is not shown.
Created	The system-generated timestamp for the image instance creation. By default, this field is not shown.
Discovery Profile	The resource profile selected for the image. By default, this field is not shown.
Min Disk	The minimum disk size required in gigabytes (such as 2GB) for the image to boot and operate. By default, this field is not shown.
Min Ram	The minimum RAM required in megabytes for the image to boot and operate, such as 2048. By default, this field is not shown.
Size	The binary image size in gigabytes. By default, this field is not shown.
Sys Object Id	The system's object ID number provided during creation. This ID is a universally unique identifier that tells what type of resource the device is. For example: `1.3.6.1.4.1.3375.2.1.3.4.43` By default, this field is not shown.
URL	The local environment's file server or disk location where the snapshot image file (.qcow2) is stored, for example: `ftp://caserverIP/vnf/images/juniper/firefly/Juniper-Firefly-shapshot-2.qcow2` By default, this field is not shown.
UUID	A unique 128-bit, system generated value assigned to each network service image descriptor created. By default, this field is not shown.
Updated	The system-generated timestamp for the latest image changes. By default, this field is not shown.

Creating new Software Image Window

Use the Creating new Software Image window (Figure 17-2) to create software images, bring that images under management so that you can deploy the images to the various VIMs that you have under management. This promotes consistency across the VIMs and simplifies the process of deploying a new image to multiple VIMs (Openstack installations).

Youcan create a software image from an existing image snapshot listed in the Software Image portlet or a new snapshot image from the original image located in the specified URL.

Accessthis windowbynavigating to the Software Images window, right clicking the portlet ora selected image, and then selecting New.

After you provide the required information, you can save the image and then edit the image to provide the additional information later.

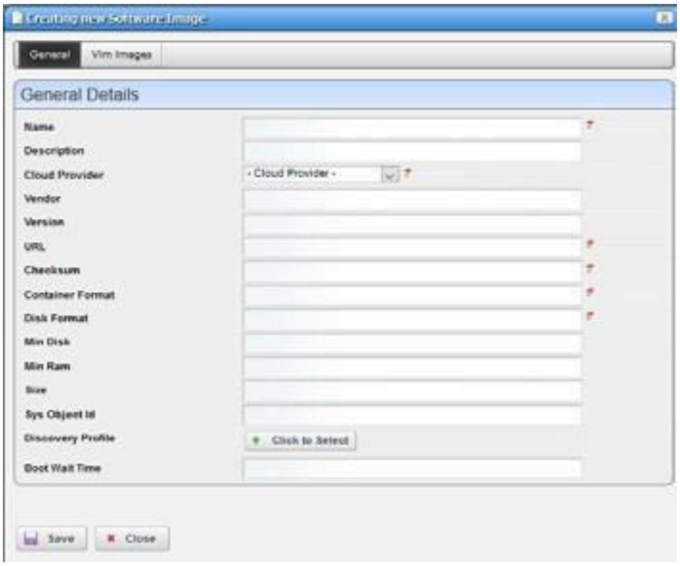

Figure 17-2. Creating new Software Image Window

The Creating new Software Image windowincludes the followingeditable fieldsand options (Table 17-5).

Table 17-5. Creating New Software Image Fields/Options (Sheet 1 of 2)

Field/Options	Description
Name	Enter a name that allows you to identify the software image.
Description	Enter more details about the software image. For example, mention customizations made for your use and the original vendor-provided image from which you created the software image.
Cloud Provider	Select the target container where you plan to deploy the image. When you select a provider, the Container Format and Disk Format lists are populated based.
Vendor	Enter a name for the vendor that supplied the initial image used to create this software image descriptor.
Version	Enter the version for the vendor-supplied image.
URL	Enter the local environment’s file server or disk location where the snapshot image file (.qcow2) is stored, for example: `ftp://caserverIP/vnf/images/juniper/firefly/Juniper-Firefly-shapshot-2.qcow2`
Checksum	Enter the value assigned to the image snapshot instance located in your OpenStack environment. For example, an F5-BIGIP LTM-SS1 image could look like this: `9275adf3c578c60754b3d56dcbbed5c8` This value is used to detect errors introduced during transmission or storage and it is provided by the image manufacturer.
Container Format	Select a format for the target container to which you deploy the image. For example, if your target container is a Bare Metal OpenStack controller, select BARE. Valid values are AKI, AMI, ARI, BARE, DOCKER, OVF, and UNRECOGNIZED.

Table 17-5. Creating New Software Image Fields/Options (Sheet 2 of 2)

Field/Options	Description
Disk Format	Select the file format of the image provided by the manufacturer, such as QCOW2. Valid formats are AKI, AMI, ARI, ISO, QCOW2, RAW, UNRECOGNIZED, VDI, VHD, VMDK.
Min Disk	Enter a minimum disk size required in gigabytes, such as 2. The default is 0 (zero) if you do not enter a value. This identifies how much disk space is required for the image to boot and operate.
Min Ram	Enter the minimum RAM required in megabytes, such as 2048. The default is 0 (zero) if you do not enter a value. This identifies how much RAM is required for the image to boot and operate.
Size	Enter the binary image size in gigabytes, such as 3.5. The default is 0 (zero) if you do not enter a value.
Sys Object Id	The system's object ID number provided during creation. This ID is a universally unique identifier that tells what type of resource the device is. For example: `1.3.6.1.4.1.3375.2.1.3.4.43`
Discovery Profile	Select a resource profile.
Boot Wait Time	Enter how many seconds to wait for an image to finish booting. The default is 0 (zero) if you do not enter a value.
Vim Images	View the list of VIM images after you deploy the image by clicking the Vim Images tab.
Save	Save the image you created and exits the window.
Close	Exist the current window without creating the image.

Editing Software Image Window

Use the Editing Software Image window (Figure 17-3) to modify the existing software image details.

Access this window by navigating to the System Images portlet, right clickingan image, and then selecting Edit.

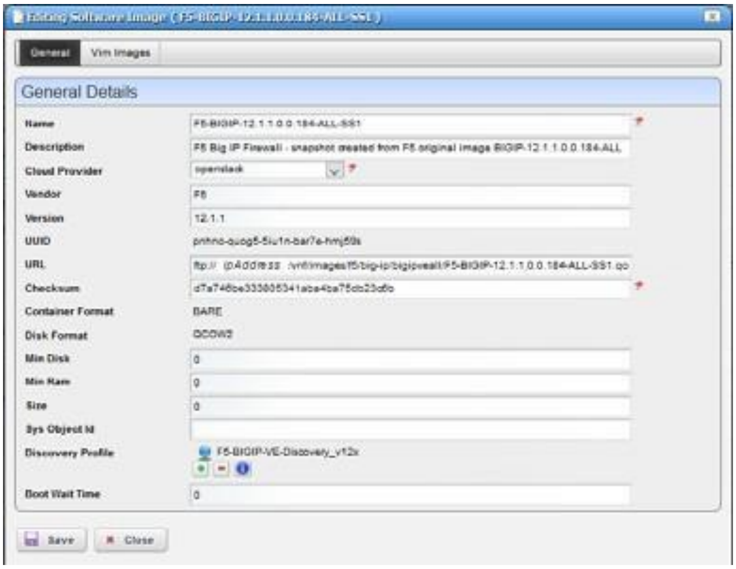

Name	F5-BIGIP-12.1.1.0.0.184-ALL-SS1
Description	F5 Big IP Firewall - snapshot created from F5 original image BIGIP-12.1.1.0.0.184-ALL
Cloud Provider	openstack
Vendor	F5
Version	12.1.1
UUID	pnhno-quog5-5iu1n-bar7e-hmj59s
URL	ftp:// ipAddress /vnf/images/f5/big-ip/bigipwall/F5-BIGIP-12.1.1.0.0.184-ALL-SS1.qo
Checksum	d7a746be33380534 1aba4ba75db23d6b
Container Format	BARE
Disk Format	QCOW2
Min Disk	0
Min Ram	0
Size	0
Sys Object Id	
Discovery Profile	F5-BIGIP-VE-Discovery_v12x
Boot Wait Time	0

Figure 17-3. Editing Software Image Window

The Editing Software Image window allows you to change the following fields and options as required (Table 17-6).

Table 17-6. Editing Software Image Fields/Options

Field/Options	Description
Name	Modify the name that allows you to identify the software image.
Description	Provide more details about the software image. For example, mention customizations made for your use and the original vendor-provided image from which you created the software image.
Cloud Provider	Select a different target container where you plan to deploy the image. When you select a provider, the Container Format and Disk Format lists are populated based.
Vendor	Modify the name for the vendor that supplied the initial image used to create this software image descriptor.
Version	Modify the version for the vendor-supplied image.
URL	Provide a different local environment’s file server or disk location where the snapshot image file (.qcow2) is stored, for example: `ftp://`*`caserverIP`*`/vnf/images/juniper/firefly/Juniper-Firefly-shapshot-2.qcow2`
Checksum	Enter the value assigned to the image snapshot instance located in your OpenStack environment. For example, an F5-BIGIP LTM-SS1 image could look like this: `9275adf3c578c60754b3d56dcbbed5c8` This value is used to detect errors introduced during transmission or storage and it is provided by the image manufacturer.
Min Disk	Modify the minimum disk size required in gigabytes, such as 2. The default is 0 (zero) if you do not enter a value. This identifies how much disk space is required for the image to boot and operate.
Min Ram	Modify the minimum RAM required in megabytes, such as 2048. The default is 0 (zero) if you do not enter a value. This identifies how much RAM is required for the image to boot and operate.
Size	Modify the binary image size in gigabytes, such as 3.5. The default is 0 (zero) if you do not enter a value.
Sys Object Id	The system’s object ID number provided during creation. This ID is a universally unique identifier that tells what type of resource the device is. For example: `1.3.6.1.4.1.3375.2.1.3.4.43`
Discovery Profile	Select a different resource profile.
Boot Wait Time	Modify how many seconds to wait for an image to finish booting. The default is 0 (zero) if you do not enter a value.
Vim Images	Select whether a VIM image is the master copy by clicking the Vim Images tab, clicking the Edit this entry icon for the appropriate image, selecting/deselecting the Master Copy option, and then saving the change.
Save	Saves any changes made and closes the window.
Close	Exits the current window without saving the changes.

Manage new Software Image from VIM Parameters Window

Use the Manage new Software Image from VIM parameters window (Figure 17-4) to bring an image under management that is already deployed to a VIM. This method of creating a software image is useful when you are not sure of all the parameters needed when creating an image using the New option.

Access this window by navigating to the Software Images portlet, right clicking an image, and selecting Manage.

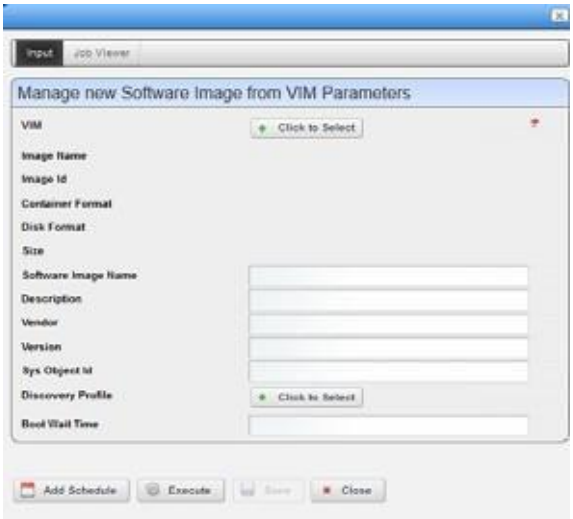

Figure 17-4. Manage new Software Image VIM Parameters Window

The Managingnew Software ImagefromVIMParameterswindowincludesthe followingfieldsand options (Table 17-7).

Table 17-7. Manage New Software Image Fields/Options (Sheet 1 of 2)

Field/Options	Description
VIM	Select the VIM where the source software image is deployed. The Image Name list is populated based on the selected VIM.
Image Name	Select the appropriate image from the list. Based on the image selected, the following fields are populated: • **Image Id** is the system-generated identifier. • **Container Format** is the target format specified during imagecreation. • **Disk Format** is the format in which the image was saved. • **Size** is a value in gigabytes. • **Software Image Name** is an editable field to specify a name identifier.
Description	Enter more details about the software image. For example, mention customizations made for your use and the original vendor-provided image from which you created the software image.
Vendor	Enter a name for the vendor that supplied the initial image used to create this software image descriptor.
Version	Enter the version for the vendor-supplied image.
Sys Object Id	The system's object ID number provided during creation. This ID is a universally unique identifier that tells what type of resource the device is. For example: `1.3.6.1.4.1.3375.2.1.3.4.43`

Table 17-7. Manage New Software Image Fields/Options (Sheet 2 of 2)

Field/Options	Description
Discovery Profile	Select a resource profile.
Boot Wait Time	Enter how many seconds to wait for an image to finish booting. The default is 0 (zero) if you do not enter a value.
Add Schedule	Set the following schedule parameters used to automatically execute the manage new software image from VIM task. • Starting On sets the date and time to execute the task. • Recurrence sets whether the task executes Only Once, Only at Startup, every X minutes (Increment), or Every *X* minutes, hours, days, weekdays, weekend days, weeks, months, years. • Stopping On sets when the task stops executing. Valid values are Never, By Occurrence, or By Date and Time.
Execute	Runs the manage new software image from VIM task.
Save	Saves the parameters and exits the window.
Close	Exits the current window without saving the VIM parameters.

Save VIM Image to Disk Parameters Window

Use the Save VIM Image to Disk Parameters window (Figure 17-5) to define the settings used to download a VIM image to disk.

Access this window by navigating to the Software Images portlet, right clicking an image, and then selecting Download.

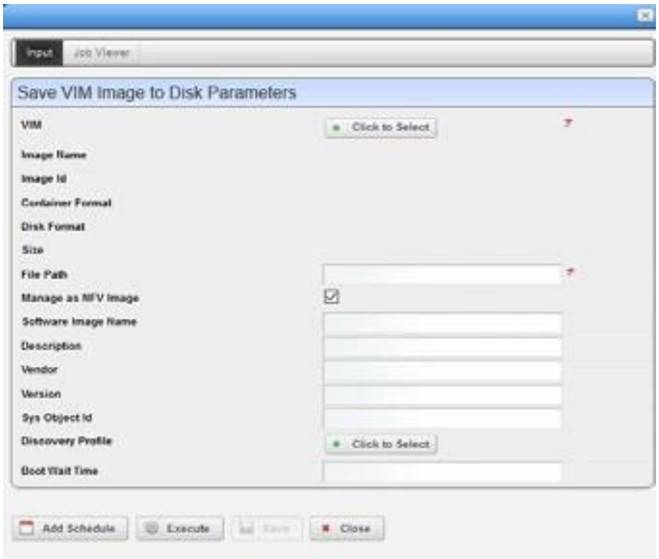

Figure 17-5. Save VIM Image to Disk Parameters Window

The Save VIM Image to Disk Parameters window includes the following fields and options (Table 17-8).

Table 17-8. Save VIM Image to Disk Parameters Fields/Options

Field/Options	Description
VIM	Select the VIM where the source software image is deployed. The Image Name list is populated based on the selected VIM.
Image Name	Select the appropriate image from the list. Based on the image selected, the following fields are populated: • **Image Id** is a system-generated identifier. • **Container Format** is the target format specified during imagecreation. • **Disk Format** is the format in which the image was saved. • **Size** is a value in gigabytes. • **Software Image Name** is an editable field to specify a name identifier.
File Path	Enter the full path to where you want to save the image.
Manage as NFV Image	Select whether to manage this image as a network functions virtualized image. This option is selected by default.
Description	Enter more details about the VIM image. For example, mention customizations made for your use and the original vendor-provided image from which you created the software image.
Vendor	Enter a name for the vendor that supplied the initial image used to create this software image descriptor.
Version	Enter the version for the vendor-supplied image.
Sys Object Id	The system's object ID number provided during creation. This ID is a universally unique identifier that tells what type of resource the device is. For example: `1.3.6.1.4.1.3375.2.1.3.4.43`
Discovery Profile	Select a resource profile.
Boot Wait Time	Enter how many seconds to wait for an image to finish booting. The default is 0 (zero) if you do not enter a value.
Add Schedule	Set the following schedule parameters used to automatically execute the manage new software image from VIM task. • Starting On sets the date and time to execute the task. • Recurrence sets whether the task executes Only Once, Only at Startup, every X minutes (Increment), or Every *X* minutes, hours, days, weekdays, weekend days, weeks, months, years. • Stopping On sets when the task stops executing. Valid values areNever, By Occurrence, or By Date and Time.
Execute	Runs the manage new software image from VIM task.
Save	Saves any changes and exits the window.
Close	Exits the current window.

VIM Images Portlet

Use the VIM Images portlet (Figure 17-6) to view a list of VIMs to which software images are deployed and details about the images or share software images with other users on your current system.

From the VIM Images portlet, a system administrator maintains consistency across the VIM instances or network designer views the list of VIMs to test that the software image they deployed/ undeployed was added to the list or removed from the list, respectively.

The VIM Images portlet also has a maximized view. Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

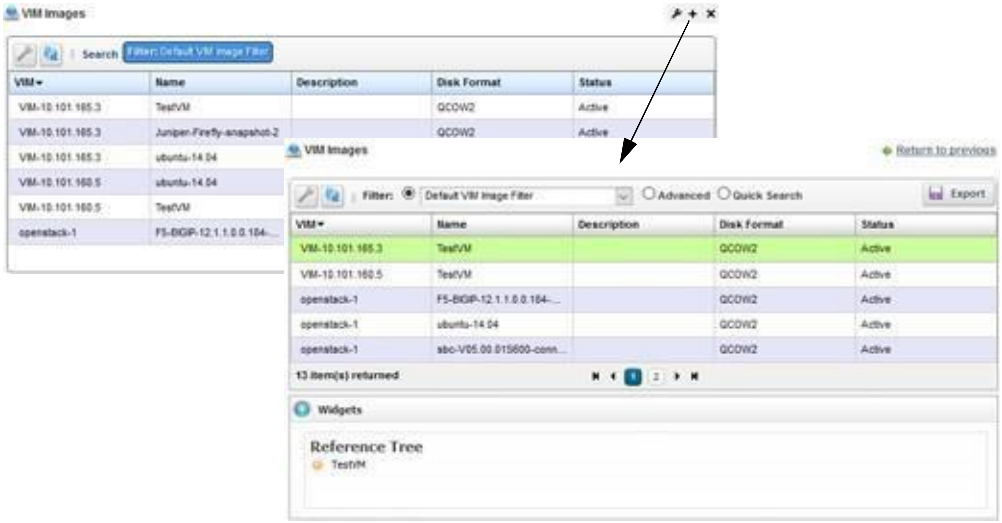
VIM Images
Search
Filter: Default VIM Image Filter

VIM	Name	Description	Disk Format	Status
VIM-10.101.165.3	TestVM		QCOW2	Active
VIM-10.101.165.3	Juniper-Firefly-snapshot-2		QCOW2	Active
VIM-10.101.165.3	ubuntu-14.04			
VIM-10.101.160.5	ubuntu-14.04			
VIM-10.101.160.5	TestVM			
openstack-1	F5-BIGIP-12.1.1.0.0.184-...			

VIM Images
Return to previous
Filter: ● Default VIM Image Filter
○ Advanced ○ Quick Search
Export

VIM	Name	Description	Disk Format	Status
VIM-10.101.165.3	TestVM		QCOW2	Active
VIM-10.101.160.5	TestVM		QCOW2	Active
openstack-1	F5-BIGIP-12.1.1.0.0.184-...		QCOW2	Active
openstack-1	ubuntu-14.04		QCOW2	Active
openstack-1	abc-V05.00.015600-conn...		QCOW2	Active

13 item(s) returned

Figure 17-6. VIM Images Portlet

Pop-Up Menu

The VIM Images pop-up menu providesaccess to the followingoptions. Right-clickarowto access these options (Table 17-9).

Table 17-9. VIM Images Pop-Up Menu Options

Menu Option	Description
Details	Displays details provided when you created the image and reference tree for the selected image. Click Return to previous to return to the Software Images portlet.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection.

Columns

Other than the general navigation and configuration options, the VIM Images portlet includes the following columns/fields (Table 17-10).

Youcan viewthe value for most of the hidden columns byright clicking a rowand selecting Details. The other option is to add the column by clicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate columns, and the applying the change.

Table 17-10. VIM Images Columns

Column	Description
VIM	The name of the virtualized Infrastructure Manager that houses the image.
Name	The user-specified image identifier.
Description	Details that explain more about the image, such as its purpose. For example, mentions customizations made for your use and the original vendor-provided image from which you created the software image.
Disk Format	The file format in which the image is to be created, such as QCOW2. Valid formats are AKI, AMI, ARI, ISO, QCOW2, RAW, UNRECOGNIZED, VDI, VHD, VMDK.
Status	The image status, such as Active.
Checksum	The data used to detect errors introduced during transmission or storage. By default, this field is not shown.
Container Format	The target format to which you deploy the image. For example, if the target is a Bare Metal OpenStack controller, the container format is BARE. Valid values are AKI, AMI, ARI, BARE, DOCKER, OVF, and UNRECOGNIZED. By default, this field is not shown.
Created	The system-generated timestamp for the image instance creation. By default, this field is not shown.
Image Id	The number assigned to the image when it was created. By default, this field is not shown.
Image URL	The local environment’s file server or disk location where the snapshot image file (.qcow2) is stored, for example: `ftp://caserverIP/vnf/images/juniper/firefly/Juniper-Firefly-shapshot-2.qcow2` By default, this field is not shown.
Master Copy	An indicator that shows whether this is a master copy or not. By default, this field is not shown.
Min Disk	The minimum disk size required in gigabytes (such as 2GB) for the image to boot and operate. By default, this field is not shown.
Min Ram	The minimum RAM required in megabytes for the image to boot and operate, such as 2048. By default, this field is not shown.
Size	The binary image size in gigabytes. By default, this field is not shown.
UUID	A unique 128-bit, system generated value assigned to each image descriptor created. By default, this field is not shown.
Updated	The system-generated timestamp for the latest image changes. By default, this field is not shown.
VIM URL	The address used to access the VIM from your browser, for example: `http://ipAddress:5000/v2.0` By default, this field is not shown.
Version	The vendor-supplied image version. By default, this field is not shown.

Network Service Portlets

This section describes the following Network Service portlets:

- Network Service Descriptors Portlet
- Network Service Records Portlet
- Details Portlets

Network Service Descriptors Portlet

Use the Network Service Descriptors portlets (Figure 17-7) to maintain a list of network service descriptors. These descriptors are templates used to create network service records that provision virtualized network functions (VNFs) and physical network functions (PNFs) used to produce connectivity.

From the Network Service Descriptors portlet, a system administrator creates and maintains the descriptors, anetworkdesigner implements networkservices, and operators(iftheyhavethe proper permissions) view descriptors to gain knowledge.

The Network Service Descriptors portlet also provides a maximized view that includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

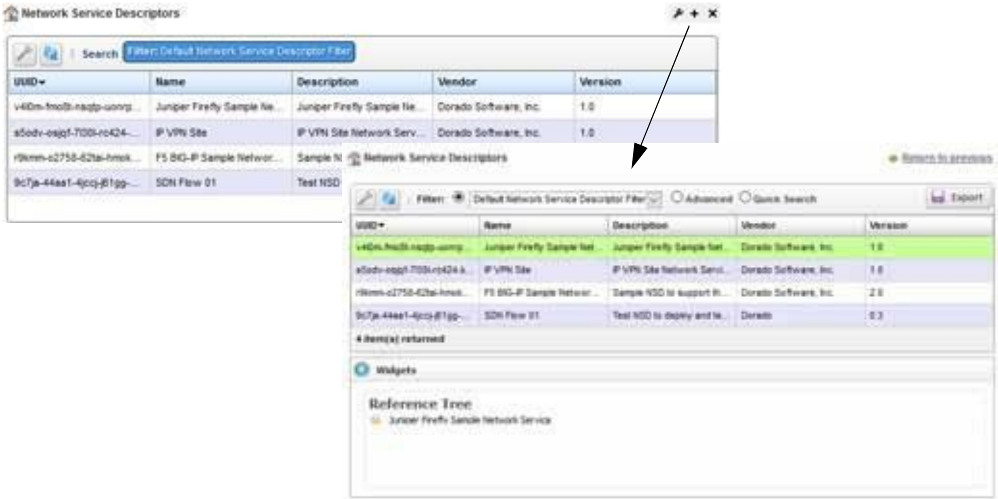
Network Service Descriptors
Search
Filter: Default Network Service Descriptor Filter
Network Service Descriptors
Return to previous
Filter:
● Default Network Service Descriptor Filter
○ Advanced
○ Quick Search
Export
Widgets
Reference Tree
Juniper Firefly Sample Network Service

UUID▾	Name	Description	Vendor	Version
v40m-fmo8t-nagtp-uonrp...	Juniper Firefly Sample Ne...	Juniper Firefly Sample Ne...	Dorado Software, Inc.	1.0
a5odv-osjqf-7t00i-ro424-...	IP VPN Site	IP VPN Site Network Serv...	Dorado Software, Inc.	1.0
r9kmm-c2758-62tai-hmok...	F5 BIG-IP Sample Networ...	Sample N		
9c7ja-44aa1-4jccj-j81gg-...	SDN Flow 01	Test NSD		

UUID▾	Name	Description	Vendor	Version
v40m-fmo8t-nagtp-uonrp...	Juniper Firefly Sample Net...	Juniper Firefly Sample Net...	Dorado Software, Inc.	1.0
a5odv-osjqf-7t00i-ro424-k...	IP VPN Site	IP VPN Site Network Servi...	Dorado Software, Inc.	1.0
r9kmm-c2758-62tai-hmok...	F5 BIG-IP Sample Networ...	Sample NSD to support th...	Dorado Software, Inc.	2.0
9c7ja-44aa1-4jccj-j81gg-...	SDN Flow 01	Test NSD to deploy and te...	Dorado	0.3

4 item(s) returned

Figure 17-7. Network Service Descriptors Portlets

Pop-Up Menu

The Network Service Descriptors portlet provides access to the following pop-up menu options. Right-click a row to access these options (Table 17-11).

Table 17-11. Network Service Descriptor Pop-Up Menu Options (Sheet 1 of 2)

Option	Description
New	Opensthe Editing Network Service Descriptor window,whereyou define a new descriptor.

Table 17-11. Network Service Descriptor Pop-Up Menu Options (Sheet 2 of 2)

Option	Description
Edit	Opens the Editing Network Service Descriptor window,where you can change the description for the selected descriptor.
Details	Displays the general details, such as identification information, versioninformation, monitor parameters, a list of VNF forwarding descriptors, and a reference tree. You can also view network functions details, connection details, service flavor details, alarms and event details, and history details by selecting the appropriate tab.
Topology	Opens the Topology portlet, where you define a multi-layered, customizable topology view of your network to help track network devices state. See Topology Portlet on page 235 for more details on configuring the Topology portlet.
Stage	Opens the Stage new Network Service window, where you set the staging parameters and then execute thestaging task.
Delete	Deletes the selected network service.
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected descriptor. Select an audit record and the job record displays its status and any other job-related information.
Import/Export	Provides the following actions when available for the selected image: • Import retrieves a file containing XML network service descriptions. Some imports can come from a URL. • Export Selection exports the selected network service description to an XML file. • Export All exports all network service descriptions to an XML file. Click Download Export File to specify where to save the file. The Import/Export option is useful as a backup or to share descriptors with other projects.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection.

Columns

Other than the general navigation and configuration options, the Network Service Descriptors portlet includes the following columns (Table 17-12).

Youcan viewthe value for most of the hidden columns byright clicking a rowand selecting Details. The other option is to add the column by clicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-12. Network Service Descriptors Columns (Sheet 1 of 2)

Column	Description
UUID	A unique 128-bit, system generated value assigned to each network service descriptor created.
Name	The user-provided name that identifies the network service descriptor.

Table 17-12. Network Service Descriptors Columns (Sheet 2 of 2)

Column	Description
Description	More details about the network service descriptor, such as its functionality and purpose.
Vendor	The name that identifies from which vendor image this descriptor was created.
Version	The vendor image version.
Created	The system-generated timestamp for the descriptor creation. By default, this field is not shown.
Enabled	An indicator that shows whether the descriptor is available (checkmark) or not (X). By default, this field is not shown.
Monitoring Parameters	Displays any monitoring parameters provided when the descriptor was created. By default, this field is not shown.
Notifications	This feature is not currently implemented. By default, this field is not shown.
On Board Date	The system-generated timestamp when the service was instantiated. By default, this field is not shown.
Updated	The system-generated timestamp for the latest descriptor changes. By default, this field is not shown.
Version Date	The timestamp provided during descriptor creation. By default, this field is not shown.

Editing Network Service Descriptor Window

Use the Editing Network Service Descriptor window (Figure 17-8) to define a new network service (NS) descriptor, modify an existing NS descriptor, or view an NS descriptor configuration.

NOTE:

Network Service descriptor creation and maintenance requires administrative permissions.

The navigation tree allows you to select and view a network service descriptor's detail for:

- VNF Dependencies lists any defined VNF dependencies.
- Data Values lists any defined data values.
- Lifecycle Events provides access to any defined list of events for the Heal, Shutdown, Configure, Scale in, Instantiate, Terminate, Scale out, and Diagnostics lifecycle events.
- Connection Points lists any defined connection points.
- Forwarding Graphics lists any defined forwarding graphs.
- Deployment Flavors lists the default or defined deployment flavors.

  Therearedeploymentflavorsavailablebydefault. However, yoursystemadministrator can create and maintain the deployment flavors available.

- Virtual Links lists any defined virtual links.
- PNF and VNF Descriptors lists any defined PND/VNF descriptors.

Select an item from the navigation tree and the available definitions are displayed in the general or extended details. If you have administrative permissions, you can add, modify, or deletes definitions.

Accessthiswindowbynavigating to the Network Service Descriptors portlet, right clicking a descriptor or the portlet, and then selecting Edit or New.

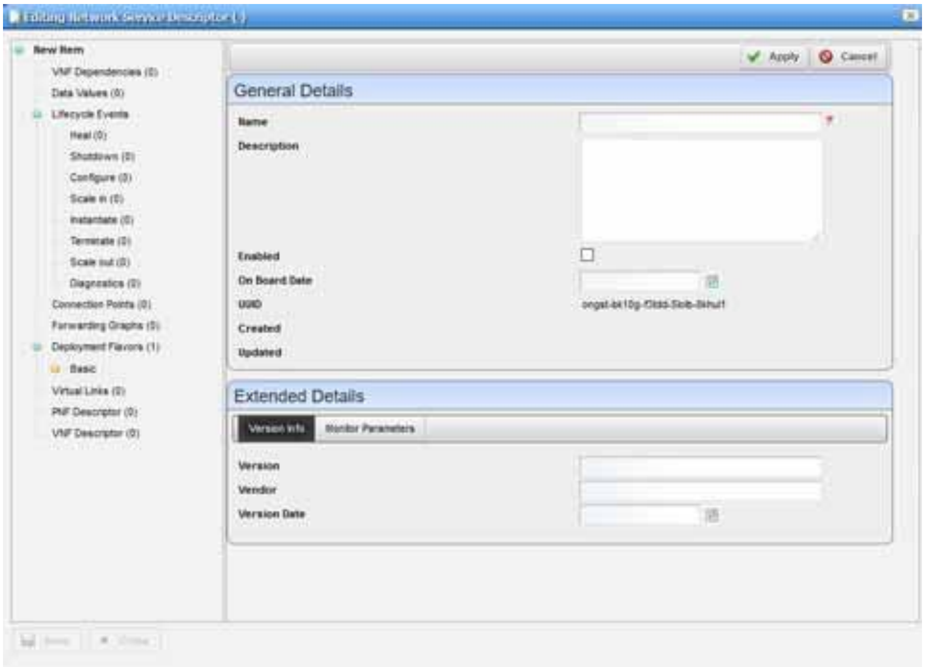

Enabled □

UUID ongst-bk10g-f3dd-5oib-8khu1

Figure 17-8. Editing Network Service Descriptor Window

The Editing Network Service Descriptor window has the following editable fields and options (Table 17-13) by default. The fields not listed are system-generated fields, such UUID, Created, and Updated.

Table 17-13. Editing Network Service Descriptor Fields/Options

Field/Option	Description
Name	Enter a name that identifies the network service descriptor.
Description	Provide more details about the network service descriptor purpose.
Enabled	Select whether the descriptor is available. Select this option to make the descriptor available. The default is no (not selected).
On Board Date	Select a date and time to instantiate the network service. Once instantiated, it displays the date of instantiation.
Version Info	Enter the network service descriptor version, vendor/provider, and date.
Monitor Parameters	Enter any monitoring parameters. This feature is not currently implemented.

Stage new Network Service Window

Use the Stage new Network Service window (Figure 17-9) to specify networkservice staging parameters.

Accessthis windowbynavigating to the Network Service Descriptors portlet, right clicking a descriptor, and then selecting Stage.

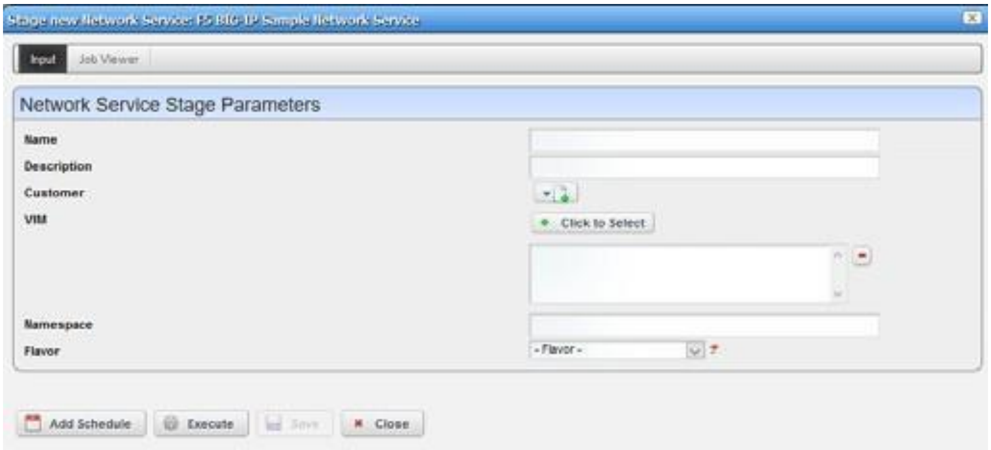

Figure 17-9. Stage new Network Service Parameters Window

The Stage new Network Service window has the following fields and options (Table 17-14).

Table 17-14. Stage new Network Service Parameters Fields/Options

Field/Option	Description
Name	Enter a name that identifies the service you want to stage.
Description	Enter a more detailed description of the service you want to stage.
Customer	Select the person, group, or organization that requested the service. If the customer is not listed, create it.
VIM	Select the Virtualized Infrastructure Manager to which you to instantiate the network service.
Namespace	Select the name for the project that houses the VIM. This list is populated when you select a deployment flavor.
Flavor	Select a target deployment model (flavor). Valid values vary depending on the service selected. Selecting a flavor populates the Namespace list.
Add Schedule	Set the following schedule parameters used to automatically execute the network service staging task. • Starting On sets the date and time to execute the task. • Recurrence sets whether the task executes Only Once, Only at Startup, every X minutes (Increment), or Every *X* minutes, hours, days, weekdays, weekend days, weeks, months, years. • Stopping On sets when the task stops executing. Valid values are Never, By Occurrence, or By Date and Time.
Execute	Runs the network service stage task.
Save	Preserves your staging configuration.
Close	Exits the Stage new Network Service window.

Network Service Records Portlet

Use the Network Service Records portlets (Figure 17-10) to maintain a list of records that provision virtualized network functions (VNFs) connectivity.

Fromthe Network Services Recordsportlet, operators can stage, deploy, andundeploynetwork services. A network designer would use this portlet to test a new or modified network service descriptor.

The Network Service Records portlet also providesa maximized viewthat includes additional filtering and the ability to export the table to a PDF, Excel, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

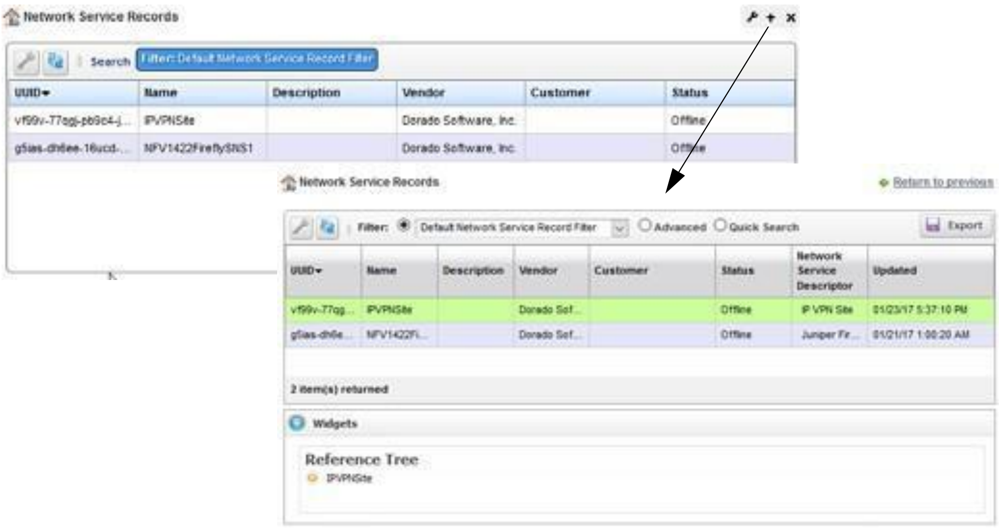

UUID▾	Name	Description	Vendor	Customer	Status
vf99v-77qgj-pb9c4-j...	IPVPNSite		Dorado Software, Inc.		Offline
g5ias-dh6ee-16ucd-...	NFV1422FireflySNS1		Dorado Software, Inc.		Offline

UUID▾	Name	Description	Vendor	Customer	Status	Network Service Descriptor	Updated
vf99v-77qg...	IPVPNSite		Dorado Sof...		Offline	IP VPN Site	01/23/17 5:37:10 PM
g5ias-dh6e...	NFV1422Fi...		Dorado Sof...		Offline	Juniper Fir...	01/21/17 1:00:20 AM

2 item(s) returned

Figure 17-10. Network Service Records Portlets

Pop-Up Menu

The Network Service Records portlet provides access to the following pop-up menu options. Right-click a row to access these options (Table 17-15).

Table 17-15. Network Service Record Pop-Up Menu Options (Sheet 1 of 2)

Option	Description
New	Opens the Stage NSR window, where you define a network service record. Select the service descriptor from which you want to create the record.
Edit	Opens the Editing Network Service Record window, where you view the selected service record definition or modify the: • Record name, description, or monitoring parameters • Connection point name or description • Virtual links name, description, or QOS options
Details	Displays the general details, such as identification information, VNF health, and a reference tree showing related virtual link records and VNF records. You can also view network functions details, connection details , alarms and event details, and history details by selecting the appropriate tab.

Table 17-15. Network Service Record Pop-Up Menu Options (Sheet 2 of 2)

Option	Description
Topology	Opens the Topology portlet, where you define a multi-layered, customizable topology view of your network to help track network devices state. See Topology Portlet on page 235 for more details on configuring the Topology portlet.
Deploy	Implements a network service in the specified VIM by establishing network connectivity, processing the VNF instantiate event, and re-syncing VIM resources.
Undeploy	Removes the network service from the specified VIM by processing the VNF terminate event and re-syncing VIM resources.
Deploy Virtual Links	Executes the Service Link Deploy task on the target NFV, processing the Instantiate event for each defined network service.
Discover/Resync	Executes the Network Service Discovery task to identify services in the network and performs a resync action to ensure services are up to date.
Scaling	Provides access to the following scaling options: • **Scale In** opens the Network Service Sale In Parameters window, where you specify a deployment flavor name and increment number, and then execute the Scale-in event to reduce the service scope. • **Scale Out** opens the Network Service Scale Out Parameters window, where you specify a deployment flavor name and increment number, and then execute the Scale-out event to increase the service scope.
Maintenance	Executes the following network service tasks on the target VNF: • **Diagnostics** runs the Diagnostics event for the selected service and the VNF. The user defines these events. • **Heal** runs the Heal event for the network service and the VNF. • **Upgrade** runs the Upgrade event for the selected service and the VNF to ensure that the latest descriptor is used. • **Disaster Recovery** runs the Disaster Recovery event for the selected service and the VNF. The user defines these events.
Delete	Removes the selected network service record from the system.
Shutdown	Gracefully shuts down the selected network service.
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected record. Select an audit record and the job record displays its status and any other job-related information.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection.

Columns

Other than the general navigation and configuration options, the Network Service Record portlet includes the following columns (Table 17-16).

You can view the value for most of the hidden columns by right clicking a row and selecting Details. The other option is to add the column byclicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-16. Network Service Records Columns

Column	Description
UUID	A unique 128-bit, system generated value assigned to each network service record created.
Name	A name that identifies the network service record.
Description	More details about the network service record, such as its functionality and purpose.
Vendor	The name for the vendor that created the record, such as Dorado Software.
Customer	The customer selected during record creation.
Status	The network service record’s status, such as Offline, Normal, *etc*.
Active Operation	The current operation being performed on the network service. Because network service operations are long running, this lets you know when the network service is ready to receive another operation. By default, this field is not shown.
Alarm Severity	The alarm severity for each record, such as Critical, Major, Minor, Informational, Normal, Warning. By default, this field is not shown.
Alarm Suppression Description	A text description of alarm suppression. By default, this field is not shown.
Alarm Suppression Mode	An indicator that shows whether alarm suppression is on or off. By default, this field is not shown.
Created	The system-generated timestamp for the record creation. By default, this field is not shown.
Deployment Flavor	The deployment selected for each record, such as Basic, Standard Basic, *etc*. By default, this field is not shown.
Monitoring Parameters	The attributes (parameters) described in the descriptor and that you are monitoring. By default, this field is not shown. This feature is not currently implemented.
Network Service Descriptor	The descriptor used to create the network service record. By default, shows only on the maximized view.
OSS ID	The identifier for the higher-level system for which the service was instantiated. In addition to being able to see the service associated with a particular OSS, the OSS ID is used to filter northbound notifications to the right place, making sure that the wrong OSS is not notified about an event that does not pertain to them. By default, this field is not shown.
Operation Start Time	The time the active operation started. By default, this field is not shown.
Updated	The system-generated timestamp for the latest record modifications. Initially this date is the same as the Created date. By default, shows only on the maximized view.
Version	The network service version. By default, this field is not shown.

Stage NSR Window

Use the Stage NSR window (Figure 17-11) to define and stage a network service record.

Access this window by navigating to the Network Service Records portlet, right clicking the portlet, and the selecting New.

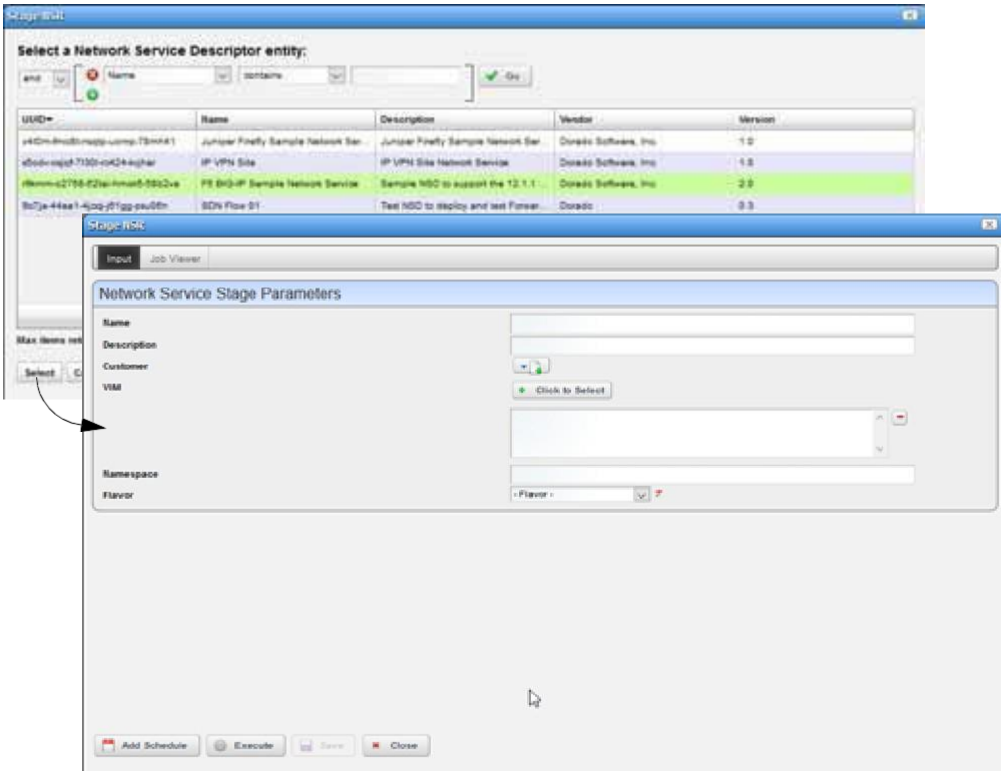

Select a Network Service Descriptor entity:

UUID	Name	Description	Vendor	Version
[illegible]	Juniper Firefly Sample Network Ser...	Juniper Firefly Sample Network Ser...	Dorado Software, Inc	1.0
[illegible]	IP VPN Site	IP VPN Site Network Service	Dorado Software, Inc	1.0
[illegible]	F5 BIG-IP Sample Network Service	Sample NSD to support the 12.1.1 ...	Dorado Software, Inc	2.0
[illegible]	SDN Flow 01	Test NSD to deploy and test Forwar...	Dorado	0.3

Figure 17-11. Stage NSR Window

The Stage NSR window has the following fields and options (Table 17-17).

Table 17-17. Stage NSR Fields/Options (Sheet 1 of 2)

Field/Option	Description
Conditional Search Parameters	Define conditional search criteria to find the appropriate descriptor.
Available Descriptors List	Select a descriptor from which to create a network service record and then click Select. The Network Service Stage Parameters input fields are displayed.
Name	Enter a name for the network service record you are creating.
Description	Enter a detailed description to understand the service function.
Customer	Select the customer requesting the service. If the customer is not on the list, add it.
VIM	Select the managed VIM to which you want to instantiate the network service.
Namespace	Select the name for the project that houses the VIM. This list is populated when you select a deployment flavor.
Flavor	Select a target deployment model (flavor). Valid values vary depending on the service selected. Selecting a flavor populates the Namespace list.

Table 17-17. Stage NSR Fields/Options (Sheet 2 of 2)

Field/Option	Description
Add Schedule	Set the following schedule parameters used to automatically execute the network service record staging task. • Starting On sets the date and time to execute the task. • Recurrence sets whether the task executes Only Once, Only at Startup, every X minutes (Increment), or Every *X* minutes, hours, days, weekdays, weekend days, weeks, months, years. • Stopping On sets when the task stops executing. Valid values areNever, By Occurrence, or By Date and Time.
Execute	Runs the network service stage task.
Save	Preserves your staging configuration.
Close	Exits the Stage new Network Service window.

Editing Network Service Record Window

Use the Editing Network Service Record window (Figure 17-12) to view and maintain the service record details, such as record name, description, and monitoring parameters. Youcan also modify some other details, such as connection points, virtual links, VNF records, and so on.

Accessthis windowbynavigating to the Network Service Recordsportlet, right clicking arecord, and then selecting Edit.

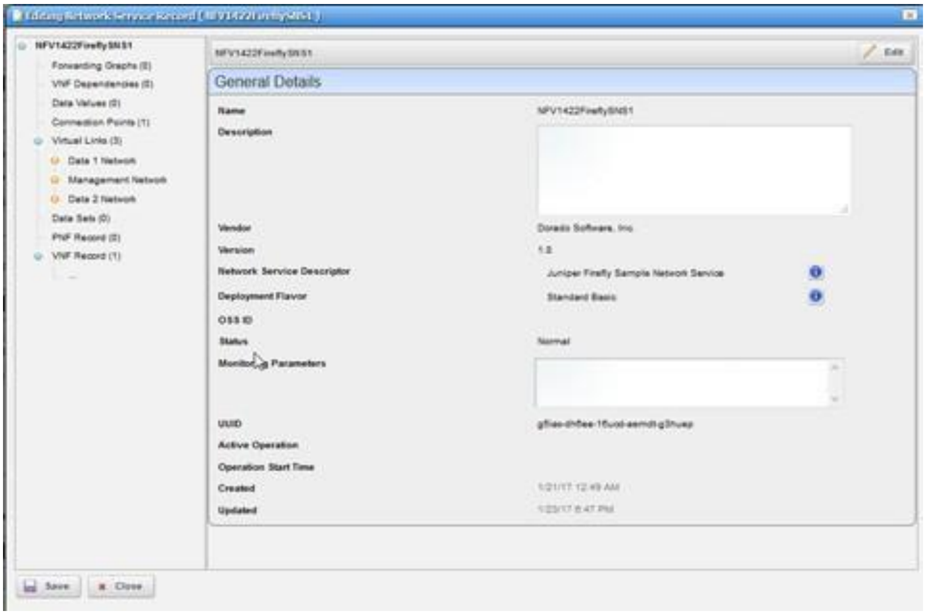

Name	NFV1422FireflySNS1
Description	
Vendor	Dorado Software, Inc
Version	1.0
Network Service Descriptor	Juniper Firefly Sample Network Service
Deployment Flavor	Standard Basic
OSS ID	
Status	Normal
Monitoring Parameters	
UUID	gflas-dhfles-1fuod-aemdt.g3huep
Active Operation	
Operation Start Time	
Created	1/21/17 12:49 AM
Updated	1/23/17 8:47 PM

Figure 17-12. Editing Network Service Record Window

The Editing Network Service Record window allows you to modify some fields for the following record elements (Table 17-18).

Table 17-18. Navigation Tree Elements

Element	Description
Service Record	Modify the record name, description, and monitoring parameters by clicking the Edit button.
Forwarding Graphs	Modify a forwarding graph name and description by selecting Forwarding Graphs, clicking the Edit this entry icon, making your changes, and then clicking Apply.
VNF Dependencies	Modify a VNF dependency name and description by selecting VNF Dependencies, clicking the Edit this entry icon, making your changes, and then clicking Apply.
Data Values	Modify a data value name and description by selecting Data Values, clicking the Edit this entry icon, making your changes, and then clicking Apply.
Connection Points	Modify a connection point name and description by selecting Connection Points, clicking the Edit this entry icon, making your changes, and then clicking Apply.
Virtual Links	Modify a virtual link name, description, and QOS options by selecting Virtual Links, clicking the Edit this entry icon, making your changes, and then clicking Apply.
Data Sets	Modify a data set name and description by selecting Data Sets, clicking the Edit this entry icon, making your changes, and then clicking Apply.
PFN Record	Cannot modify this record element.
VNF Record	Cannot modify this record element.

Details Portlets

The Details portlets display the general details for the selected network service descriptor or record. Except where noted, the general details are the same for the network service descriptor and network service record.

You can also view network functions details, connection details, service flavor details (descriptor only), alarms and event details , and history details related to the selected Network Service descriptor or record.

General Details

Use the Network Service General details (Figure 17-13) to view identification and version information, and a reference tree of items related to the selected network service descriptor or record. You can also view:

- Anymonitoring parameter definitions and related VNF forwardinggraphdescriptorsforthe selected network service descriptor
- Service health for the selected network service record

Access this information by navigating to the appropriate Network Service portlet, right clicking a row, and then selecting Details.

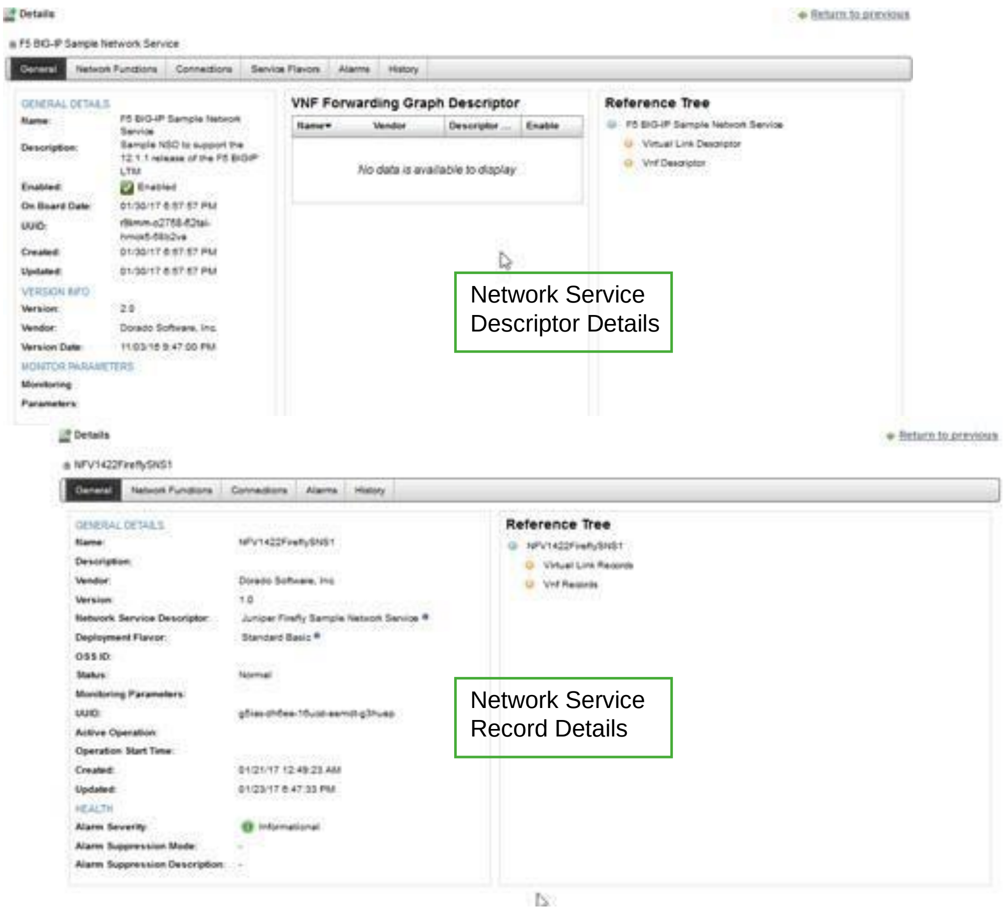

F5 BIG-IP Sample Network Service

GENERAL DETAILS

Name: F5 BIG-IP Sample Network Service

Description: Sample NSD to support the 12.1.1 release of the F5 BIGIP LTM

Enabled: ☑ Enabled

On Board Date: 01/30/17 6:57:57 PM

UUID: rf6mm-o2768-82tal-hmox5-66b2va

Created: 01/30/17 6:57:57 PM

Updated: 01/30/17 6:57:57 PM

VERSION INFO

Version: 2.0

Vendor: Dorado Software, Inc

Version Date: 11/03/16 9:47:00 PM

MONITOR PARAMETERS

Monitoring Parameters:

VNF Forwarding Graph Descriptor

Name	Vendor	Descriptor ...	Enable

No data is available to display

Reference Tree

- F5 BIG-IP Sample Network Service
  - Virtual Link Descriptor
  - Vnf Descriptor

NFV1422FireflySNS1

GENERAL DETAILS

Name: NFV1422FireflySNS1

Description:

Vendor: Dorado Software, Inc

Version: 1.0

Network Service Descriptor: Juniper Firefly Sample Network Service

Deployment Flavor: Standard Basic

OSS ID:

Status: Normal

Monitoring Parameters:

UUID: g5ias-dhfee-1fuot-eemdt-g3huep

Active Operation:

Operation Start Time:

Created: 01/21/17 12:49:23 AM

Updated: 01/23/17 6:47:33 PM

HEALTH

Alarm Severity: Informational

Alarm Suppression Mode: -

Alarm Suppression Description: -

Reference Tree

- NFV1422FireflySNS1
  - Virtual Link Records
  - Vnf Records

Figure 17-13. Network Service General Details Portlet

Network Functions Details

Use the Network Functions details (Figure 17-14) to view the related VNF record, PNF record, and dependency details.

You can perform the same action s as those available from the Virtual Network Function Descriptors Portlet on page 917 (New, edit, Details, Discover Vnfd, Branch, Stage, Delete, Audit Trail, Import/Export, Share with User).

Access this information by navigating to the appropriate Network Service portlet, right clicking a row, selecting Details, and then clicking the Network Functions tab.

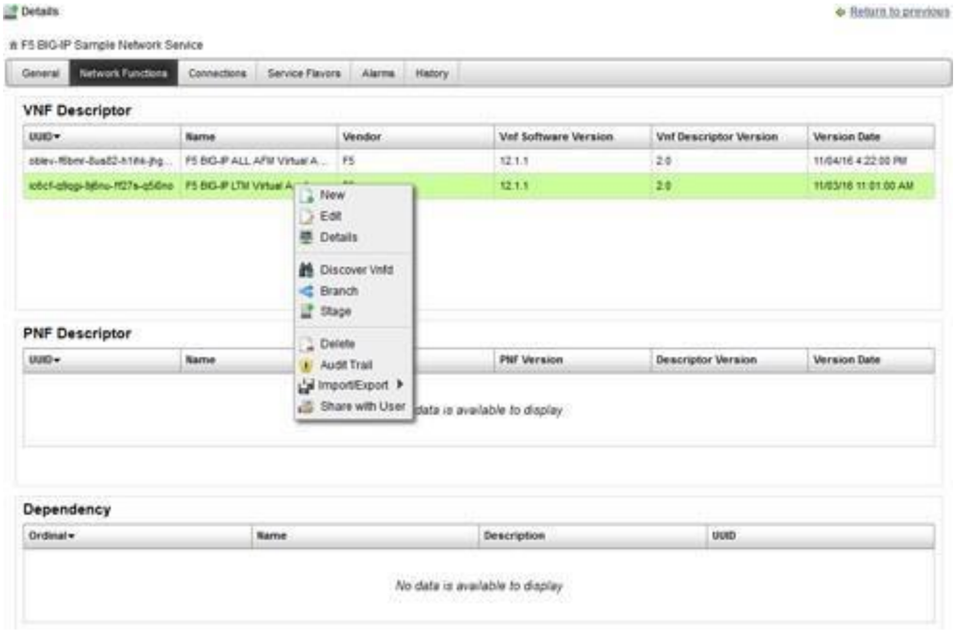

VNF Descriptor

UUID	Name	Vendor	Vnf Software Version	Vnf Descriptor Version	Version Date
ob1ev-f6bmr-8us82-h1th6-jhg...	F5 BIG-IP ALL AFM Virtual A...	F5	12.1.1	2.0	11/04/16 4:22:00 PM
io6cf-a9qgi-9j6nu-ff27s-q56no	F5 BIG-IP LTM Virtual A...	F5	12.1.1	2.0	11/03/16 11:01:00 AM

PNF Descriptor

UUID	Name		PNF Version	Descriptor Version	Version Date

data is available to display

Dependency

Ordinal	Name	Description	UUID

No data is available to display

Figure 17-14. Network Functions Details

Connection Details

Use the Network Service Connection details (Figure 17-15) to view moregranular detailsabout the connection point and virtual link descriptor/record. This information is useful when trying to troubleshoot connectivity issues.

You have the option to drill down to more details or share these details with another user.

If the Connections details are for a selected network service record, you can edit or delete a VL record.

Access this information by navigating to the appropriate Network Service portlet, right clicking a row, selecting Details, and then clicking the Connections tab.

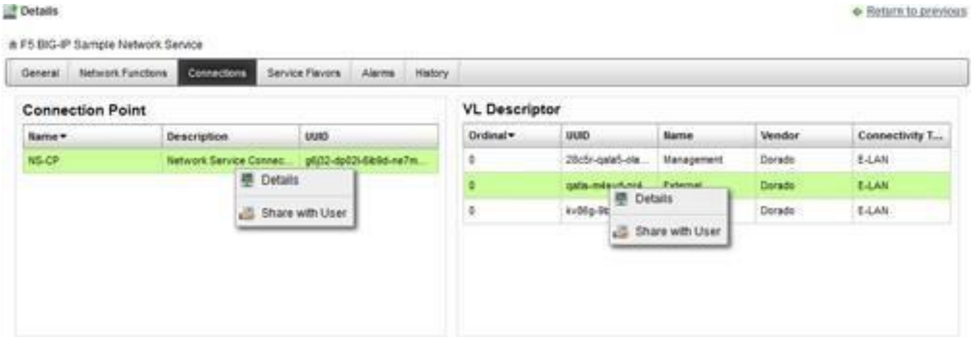

Connection Point

Name	Description	UUID
NS-CP	Network Service Connec...	g6j32-dp02i-6b9d-ne7m...

VL Descriptor

Ordinal	UUID	Name	Vendor	Connectivity T...
0	28c5r-qala5-ola...	Management	Dorado	E-LAN
0	qala-m4eu...	External	Dorado	E-LAN
0	kv06g-9t...		Dorado	E-LAN

Figure 17-15. Network Service Connection Details

Service Flavor Details

Use the Network Service Flavors details (Figure 17-16) to viewa list of deployment flavors for the selected network service descriptor. These deployment flavors are useful when there is a need to scale in/out the number of instances to decrease/increase network service capacity, respectively.

You have the option to share a service flavor with another user on your system.

Access this information by navigating to the Network Service Descriptors portlet, right clicking a row, selecting Details, and then clicking the Service Flavors tab.

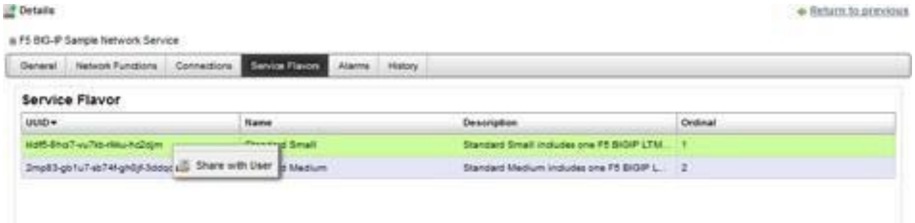

Service Flavor

UUID	Name	Description	Ordinal
lldf5-8ho7-vu7kb-rkku-hc2djm	Standard Small	Standard Small includes one F5 BIGIP LTM...	1
2mp83-gb1u7-sb74f-gh0jf-3ddqc	Standard Medium	Standard Medium includes one F5 BIGIP L...	2

Figure 17-16. Network Service Flavors Details

AlarmsandEventDetails

Use the Network Service Alarms details (Figure 17-17) to monitor and manage alarms and view event history for the selected service.

You can perform the same actions on an alarm or event history from this location as you can from the Alarms and Event History portlets, respectively. See Alarms on page 276 for details about the Alarms and Event History portlets.

Access this information by navigating to the appropriate Network Service portlet, right clicking a row, selecting Details, and then clicking the Alarms tab.

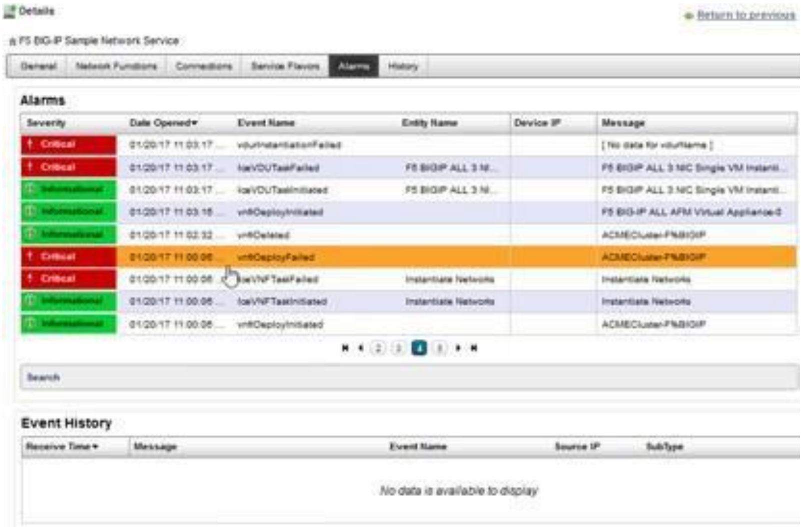

Alarms

Severity	Date Opened	Event Name	Entity Name	Device IP	Message
Critical	01/20/17 11:03:17 ...	vduInstantiationFailed			[No data for vduName]
Critical	01/20/17 11:03:17 ...	IcaVDUTaskFailed	F5 BIGIP ALL 3 NI...		F5 BIGIP ALL 3 NIC Single VM Instanti...
Informational	01/20/17 11:03:17 ...	IcaVDUTaskInitiated	F5 BIGIP ALL 3 NI...		F5 BIGIP ALL 3 NIC Single VM Instanti...
Informational	01/20/17 11:03:16 ...	vnfDeployInitiated			F5 BIG-IP ALL AFM Virtual Appliance-0
Informational	01/20/17 11:02:32 ...	vnfDeleted			ACMECluster-F%BIGIP
Critical	01/20/17 11:00:06 ...	vnfDeployFailed			ACMECluster-F%BIGIP
Critical	01/20/17 11:00:06 ...	IcaVNFTaskFailed	Instantiate Networks		Instantiate Networks
Informational	01/20/17 11:00:06 ...	IcaVNFTaskInitiated	Instantiate Networks		Instantiate Networks
Informational	01/20/17 11:00:06 ...	vnfDeployInitiated			ACMECluster-F%BIGIP

Event History

Receive Time	Message	Event Name	Source IP	SubType

No data is available to display

Figure 17-17. Network Service Alarms Details

History Details

Use the Network Service History details (Figure 17-18) to view a list of actions (audit trail) performed on the selected service.

Access this information by navigating to the appropriate Network Service portlet, right clicking a row, selecting Details, and then clicking the History tab.

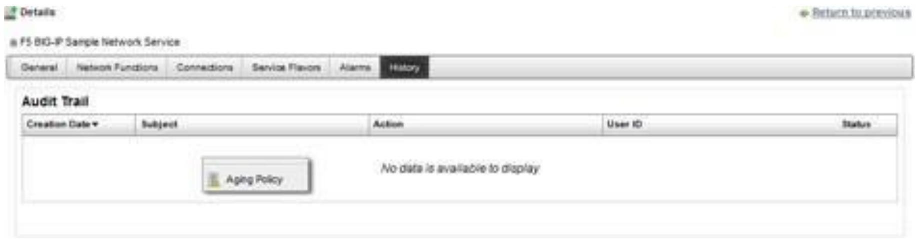

Creation Date ▾	Subject	Action	User ID	Status

No data is available to display

Figure 17-18. Network Service History Details

NFV MonitoringAttributes Portlet

Use the NFV Monitoring Attributes portlet (Figure 17-19) to view, create and maintain virtualized network function monitoring (VNFM) descriptors.

Fromthe NFVMonitoring Attributesportlet, asystemadministratorcreatesand maintains VNFM descriptors, a network designer defines monitoring attribute characteristics, and operators (if they have the proper permissions) view descriptors to gain knowledge.

The NFV Monitoring Attributes portlet also provides a maximized view that includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

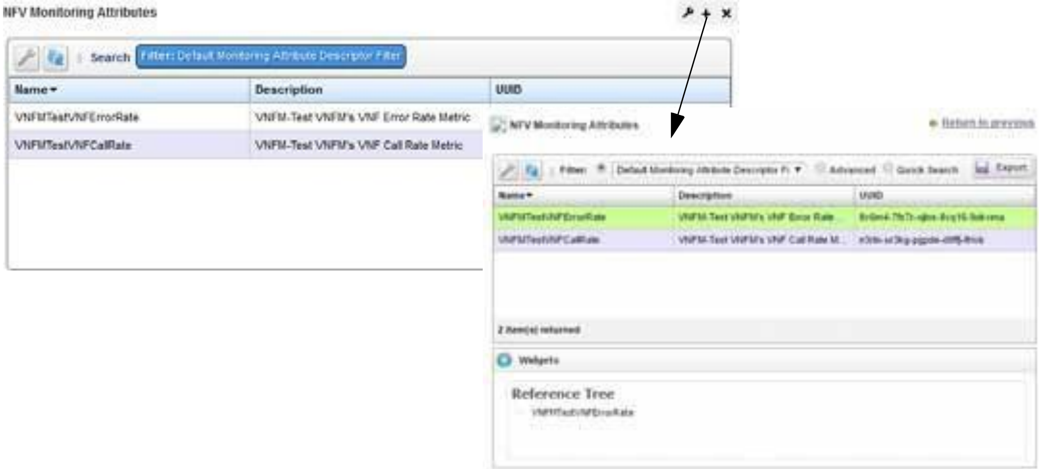
NFV Monitoring Attributes
Search
Filter: Default Monitoring Attribute Descriptor Filter
NFV Monitoring Attributes
Return to previous
Filter:
Default Monitoring Attribute Descriptor Fi ▾
Advanced
Quick Search
Export
Widgets
Reference Tree
VNFMTestVNFErrorRate

Name ▾	Description	UUID
VNFMTestVNFErrorRate	VNFM-Test VNFM's VNF Error Rate Metric	
VNFMTestVNFCallRate	VNFM-Test VNFM's VNF Call Rate Metric	

Name ▾	Description	UUID
VNFMTestVNFErrorRate	VNFM-Test VNFM's VNF Error Rate	[illegible]
VNFMTestVNFCallRate	VNFM-Test VNFM's VNF Call Rate M...	[illegible]

2 item(s) returned

Figure 17-19. NFV Monitoring Attributes Portlet

Pop-Up Menu

The NFV Monitoring Attributes portlet provides access to the following pop-up menu options. Right-click a row to access these options (Table 17-19).

Table 17-19. NFV Monitoring Attributes Pop-Up Menu Options (Sheet 1 of 2)

Option	Description
New	Opens the Editing Monitoring Attribute Descriptor window , where you create a VNFM descriptor.

Table 17-19. NFV Monitoring Attributes Pop-Up Menu Options (Sheet 2 of 2)

Option	Description
Edit	Opens the Editing Monitoring Attribute Descriptor window , where you modify a VNFM descriptor.
Details	Displays the general VNFM details, such as identification information, attribute properties, version information, and a reference tree.
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected VNFM. Select an audit record and the job record displays its status and any other job-related information.
Delete	Removes the selected VNFM descriptor from the system. A confirmation message is displayed giving you a chance to change your mind.
Import/Export	Provides the following actions when available for the selected image: • Import retrieves a file containing XML VNFM descriptions. Some imports can come from a URL. • Export Selection exports the selected VNFM description to an XML file. • Export All exports all VNFM descriptions to an XML file. Click Download Export File to specify where to save the file. The Import/Export option is useful as a backup or to share descriptors with other projects.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection.

Columns

Other than the general navigation and configuration options, the NFV Monitoring Attributes portlet includes the following columns (Table 17-20).

You can view the value for most of the hidden columns by right clicking a row and selecting Details. The other option is to add the column by clicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-20. NFV Monitoring Attributes Columns (Sheet 1 of 2)

Column	Description
Name	The text that identifies each VNFM descriptor.
Description	A detailed description of the VNFM's purpose.
UUID	A unique 128-bit, system generated value assigned to each VNFM descriptor created.
Attribute Id	Unique identifier for an attribute. By default, this field is not shown.
Composite	An indicator that shows whether the monitored attribute is a composite value or not. By default, this field is not shown.
Created	The system-generated timestamp in which the VNFM descriptor was created. By default, this field is not shown.
Foreign Id	The attribute identifier from a VNFM that relates an NFV descriptor with the VNFM attribute definition. By default, this field is not shown.
Maximum Value	An attribute's minimum valid value. By default, this field is not shown.
Metric Type	The type of metrics collected. Valid values are Boolean, Counter, Gauge, or Says Up Time. By default, this field is not shown.

Table 17-20. NFV Monitoring Attributes Columns (Sheet 2 of 2)

Column	Description
Minimum Value	An attribute’s maximum value. By default, this field is not shown.
Minimum Version	The descriptor version to which the monitoring attribute is applied. By default, this field is not shown.
Owner	The name of the vendor from which the VNF was obtained so that the system knows which attribute descriptors apply to which VNF descriptors. By default, this field is not shown.
Scope	The VNFM descriptor scope is for a Network Service, Vdu, or Vnf. By default, this field is not shown.
Type	The attribute data type. Valid values are Boolean, Double, Float, Integer, IP, Long, or String. By default, this field is not shown.
Units	Units of measure, such as calls per seconds, errors per seconds, and so on. By default, this field is not shown.
Updated	The system-generated timestamp for the latest record modifications. Initially this date is the same as the Created date. By default, this field is not shown.

Editing Monitoring Attribute Descriptor Window

Use the Editing Monitoring Attribute Descriptor window (Figure 17-20) to create and modify monitoring attribute descriptors.

Accessthiswindowbynavigating to the NFV Monitoring Attributes portlet, right clicking arow, and then selecting Edit or New.

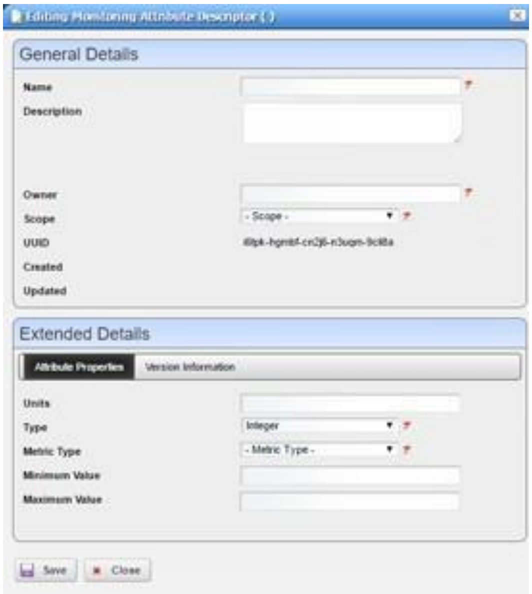

Scope: - Scope -

UUID: i8pk-hgmbf-cn2j6-n3uqm-9ck8a

Type: Integer

Metric Type: - Metric Type -

Figure 17-20. Editing Monitoring Attribute Descriptor Window

The Editing Monitoring Attribute Descriptor window has the following input fields and options (Table 17-21). The fields not listed are system-generated fields, such UUID, Created, and Updated.

Table 17-21. Editing Monitoring Attribute Descriptor Fields/Options

Field/Option	Description
Name	Enter a name for the monitoring descriptor you are creating. This field is required.
Description	Enter an optional description that provides more details about the descriptors purpose.
Owner	Enter the name of the vendor from which the VNF was obtained so that the system knows which attribute descriptors apply to which VNF descriptors. This field is required.
Scope	Select whether this descriptor is a Network Service, Vdu, or Vnf. This field is required.
Attribute Properties	Enter the following attribute properties: • Units of measure • Type of data (Boolean, Double, Float, Integer, IP, Long, or String) • Metric type (Boolean, Counter, Gauge, or Says Up Time) • Minimum Value (available only if you select Double, Float, Integer, or Long type) • Maximum Value (available only if you select Double, Float, Integer, or Long type) **Note:** Type and Metric Type are required.
Version Information	Enter the following version information: • Foreign Id (required) • Minimum Version • Maximum Version

OSS Instance Portlet

Use the OSS Instance portlet (Figure 17-21) to view, create and maintain OSS instances.

From the OSS Instance portlet, a system administrator implements OSS instances and operators (if they have the proper permissions) view OSS instances to gain knowledge.

The OSSInstance portlet also providesa maximized viewthat includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

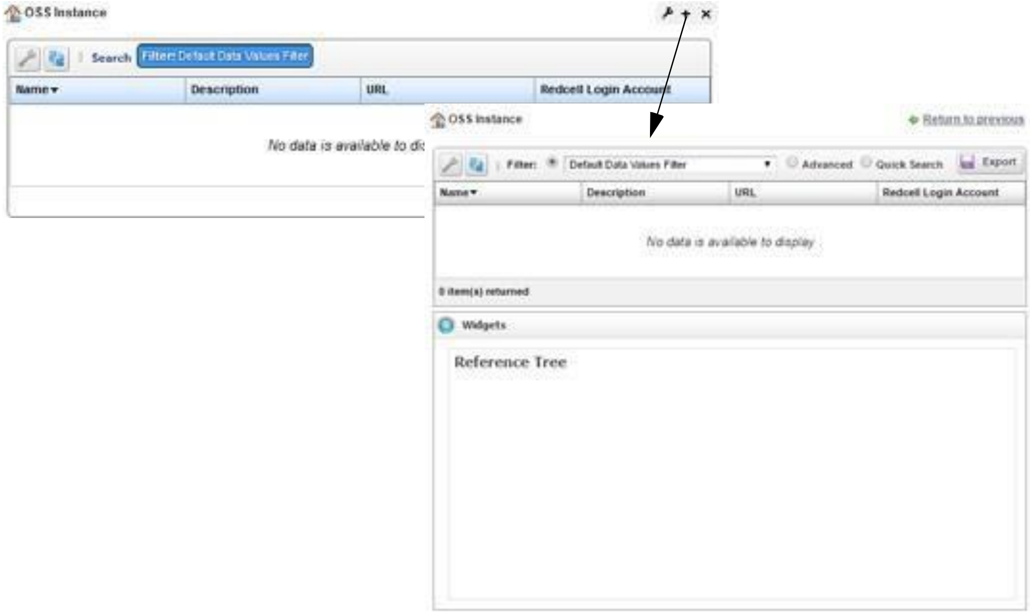

Figure 17-21. OSS Instance Portlet

Pop-Up Menu

The OSS Instance portlet provides access to the following pop-up menu options. Right-click a row to access these options (Table 17-22).

Table 17-22. OSS Instance Pop-Up Menu Options

Option	Description
New	Opens the Editing Data Values window , where you create an OSS instance.
Edit	Opens the Editing Data Values window , where you modify an OSS instance.
Details	Displays the general OSS instance details, such as identification information, version information, a reference tree, and so on.
Delete	Removes the selected OSS instance from the system. A confirmation message is displayed giving you a chance to change your mind.
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected PNF. Select an audit record and the job record displays its status and any other job-related information.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection.

Columns

Other than the general navigationand configuration options, the OSSInstance portlet includesthe following columns (Table 17-23).

Youcan viewthe value for most of the hidden columns byright clicking a rowand selecting Details. The other option is to add the column by clicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-23. OSS Instance Columns

Column	Description
Name	The text that identifies each the OSS instance.
Description	A detailed description of the OSS instance’s purpose.
URL	The file server or disk location where the OSS instance is stored.
Redcell Login Account	The remote login account for users from the OSS (operations support system).
Created	The system-generated timestamp in which the OSS instance was created. By default, this field is not shown.
OSS Id	A unique identifier for the OSS. In addition to being able to see services associated with a OSS, the OSS ID is used to filter northbound notifications to the right place, making sure that the wrong OSS is not notified about an event that does not pertain to them. By default, this field is not shown.
Password	The password associated with the User ID assigned to this OSS instance. By default, this field is not shown.
UUID	A unique 128-bit, system generated value assigned to each OSS instance created. By default, this field is not shown.
Updated	The system-generated timestamp for the latest record modifications. Initially this date is the same as the Created date. By default, this field is not shown.
User	The user ID assigned to this OSS instance during creation. By default, this field is not shown.

Editing Data Values Window

Use the Editing Data Values window (Figure 17-22) to create and modify monitoring OSS Instances.

Accessthis windowbynavigating to the OSSInstance portlet, right-clicking, and then selecting New.

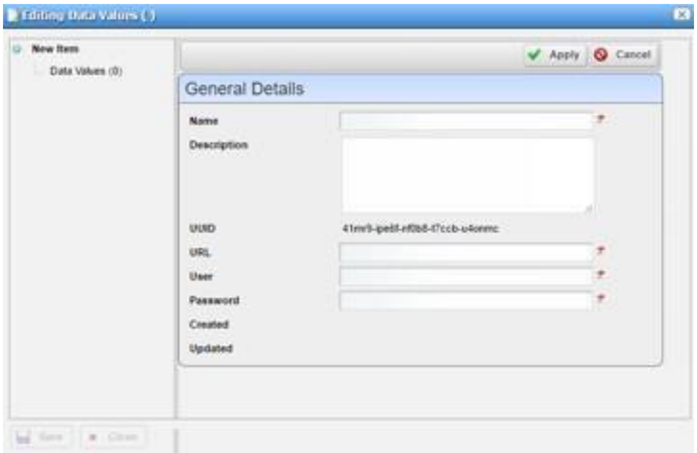

UUID 41mr9-ipe8l-nf0b8-l7ccb-u4onmc

Figure 17-22. Editing Data Values Window

The Editing Data Values window has the following input fields and options (Table 17-24). The fields not listed are system-generated fields, such UUID, Created, and Updated.

Table 17-24. Editing Data Values Fields/Options

Field/Option	Description
Name	Enter a name for the OSS instance that you are creating. This field is required.
Description	Enter an optional description that provides more details about the OSS instance’s purpose.
URL	Enter file server or disk location where the OSS instance is stored. This field is required.
User	Enter the user ID. This field is required.
Password	Enter the password assigned to the provided User ID. This field is required.

Physical Network FunctionPortlets

This section describes the following Physical Network Function (PNF) portlets:

- Physical Network Function Descriptors Portlet
- Physical Network Function Records Portlet
- Detail Portlets

Physical Network Function Descriptors Portlet

Usethe Physical Network Function Descriptors portlet (Figure 17-23) to view, create, and maintain PNF descriptors. The PNF descriptors are used to define network connections between a PNF and other network instances, such as network service, virtualized network functions, virtual links, and so on.

From the Physical Network Functions Descriptors portlet, a system administrator creates and maintains the PNF descriptors, a network designer PNFs, and operators (if they have the proper permissions) view descriptors to gain knowledge.

The Physical Network Function Descriptors portlet also provides a maximized view that includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

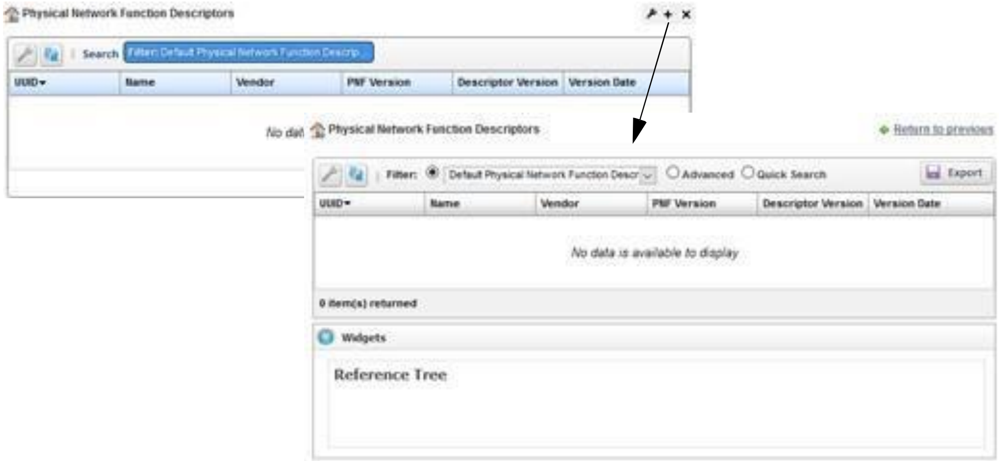

Figure 17-23. Physical Network Function Descriptors Portlet

Pop-Up Menu

The Physical Network Function Descriptors portlet providesaccess to the following pop-up menu options. Right-click a row to access these options(Table 17-25).

Table 17-25. Physical Network Function Descriptors Pop-Up Menu Options (Sheet 1 of 2)

Option	Description
New	Opens the Editing Physical Network Function Descriptor window, where you create a physical network function descriptor.
Edit	Opens the Editing Physical Network Function Descriptor window, where you modify a physical network function descriptor.
Details	Displays the general details, such as identification information, version information, connection points, and a reference tree. You can also view alarms and event details and history details by selecting the appropriate tab.
Delete	Removes the selected PNF descriptor from the system. A confirmation message is displayed giving you a chance to change your mind.
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected PNF. Select an audit record and the job record displays its status and any other job-related information.
Import/Export	Provides the following actions when available for the selected image: • Import retrieves a file containing XML PNF descriptions. Some imports can come from a URL. • Export Selection exports the selected PNF description to an XML file. • Export All exports all PNF descriptions to an XML file. Click Download Export File to specify where to save the file. The Import/Export option is useful as a backup or to share descriptors with other projects.

Table 17-25. Physical Network Function Descriptors Pop-Up Menu Options (Sheet 2 of 2)

Option	Description
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection.

Columns

Other than the general navigation and configuration options, the Physical Network Function Descriptors portlet includes the following columns (Table 17-26).

You can view the value for most of the hidden columns by right clicking a row and selecting Details. The other option is to add the column byclicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-26. Physical Network Function Descriptors Columns

Column	Description
UUID	A unique 128-bit, system generated value assigned to each PNF descriptor created.
Name	The text that identifies each PNF descriptor.
Vendor	The vendor generating this PNF descriptor.
Pnf Version	The version for the PNF that this PNF descriptor is describing. Provided by the network service engineer creating the descriptor.
Descriptor Version	The version of the PNF descriptor.
Version Date	The timestamp for the PNF descriptor.
Created	The system-generated timestamp in which the PNF descriptor was created. By default, this field is not shown.
Description	A detailed description of the PNF's purpose. By default, this field is not shown.
Enabled	An indicator that shows whether the PNF descriptor is available (checkmark) or not (X). By default, this field is not shown.
On Board Date	The user-specified timestamp in which to instantiate the PNF descriptor. By default, this field is not shown.
Update	The system-generated timestamp that the PNF descriptor was last modified. By default, this field is not shown.

Editing Physical Network Function Descriptor Window

Use the Editing Physical Network Function Descriptor window (Figure) to createand modify PNF descriptors. This is where you define network connections between a PNF and other network instances, such as network service, virtualize d network functions, virtual links, and so on.

Accessthiswindowbynavigating to the Physical Network Function Descriptorsportlet, right-clicking a row, and then selecting Edit.

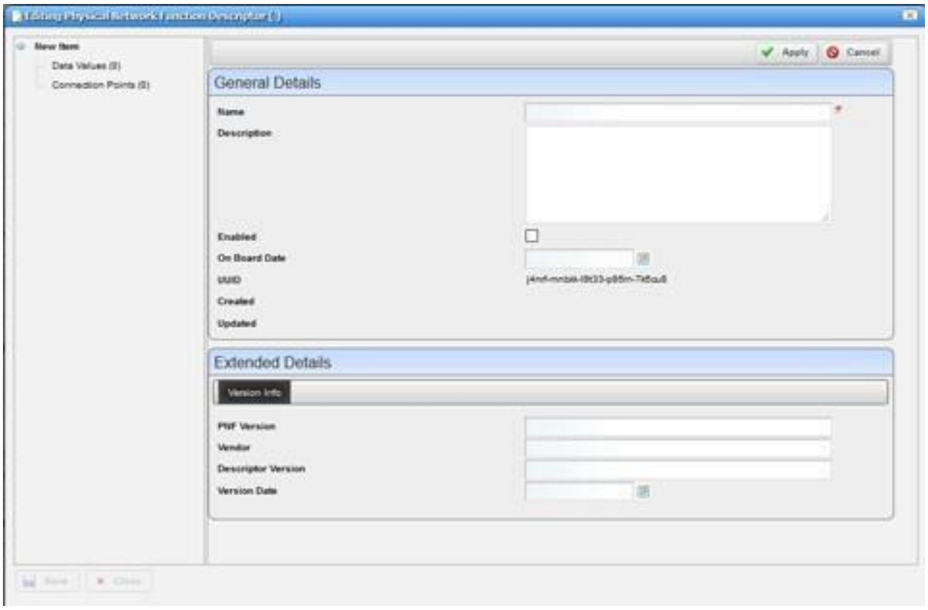

Figure 17-24. Editing Physical Network Function Descriptor Window

The Editing Physical Network Function Descriptor window has the following input fields and options (Table 17-27). The fields not listed are system-generated fields, such UUID, Created, and Updated.

Table 17-27. Editing Physical Network Function Descriptor Fields/Options (Sheet 1 of 2)

Field/Option	Description
Name	Enter a name for the PNF descriptor you are creating. If you are modifying the PNF descriptor, you can change the name. This field is required.
Description	Enter an optional description that provides more details about the PNF.
Enable	Indicates whether the device is available for use. If you select this option during PNF descriptor creation, you must define connection points . Otherwise, an error is displayed when save the descriptor.
On Board Date	Enter the date and time that you want to instantiate the PNF descriptor.
Vendor Info	Enter the PNF version, vendor, descriptor version, and the version date.
Data Values	Define one or more data value records. Enter a name, description, data type, and value. The name and data type are required fields. The system generates the UUID, Created, and Updated information. Data values are typically captured at deploy time. They are either entered directly by the user or extracted from the resulting deployed VDU. The types of values is completely dependent upon the network service. For example: • A VNF may need to know what the gateway IP address is for the service so that it can properly configure the service. • A value is extracted from the OpenStack system is the heat stack ID, that ID is then used at a later date when undeploying the service.

Table 17-27. Editing Physical Network Function Descriptor Fields/Options (Sheet 2 of 2)

Field/Option	Description
Connection Points	Define one or more connection points between the PNF and other network instances, such as network service, virtualized network functions, virtual links, *etc*. The name and connection type are required. Valid connection types are EthernetPort, IPVPN, OTHER, PNIC, PPORT, VNIC (virtual network interface controller), VPORT, VPORTLAN, VPORTWAN.

Physical Network Function Records Portlet

Use the Physical Network Function Records portlet (Figure 17-25) to view, create, and maintain PNF records.

From the Physical Network Function Records portlet, operators create PNF records that specify networkconnectivitybetween aphysicaldeviceand virtual networkinstances. Anadministrator or network designer would use this portlet to test a new or modified descriptor or test an implemented PNF.

The Physical Network Function Records portlet also provides a maximized view that includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

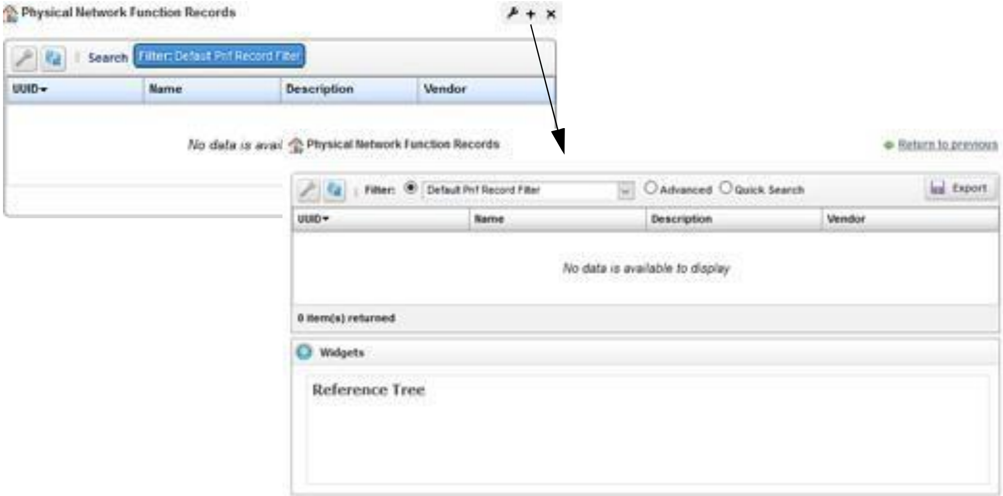

Filter: ● Default Pnf Record Filter

No data is available to display

0 item(s) returned

Figure 17-25. Physical Network Function Records Portlet

Pop-Up Menu

The Physical Network Function Records portlet provides access to the following pop-up menu options. Right-click a row to access these options (Table17-28).

Table 17-28. Physical Network Function Records Pop-Up Menu Options (Sheet 1 of 2)

Option	Description
New	Opens the Editing Pnf Record window, where you create a PNF record.
Edit	Opens the Editing Pnf Record window, where you modify a PNF record name or description.

Table 17-28. Physical Network Function Records Pop-Up Menu Options (Sheet 2 of 2)

Option	Description
Details	Displays the general details, such as identification information, version information, connection points, and a reference tree. You can also view alarms and event details and history details by selecting the appropriate tab.
Delete	Removes the selected PNF descriptor from the system.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection

Columns

Other than the general navigation and configuration options, the Physical Network Function Records portlet includes the following columns (Table 17-29).

Youcan viewthe value for most of the hidden columns byright clicking a rowand selecting Details. The other option is to add the column by clicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-29. Physical Network Function Records Columns

Column	Description
UUID	A unique 128-bit, system generated value assigned to each PNF record created.
Name	The text that identifies the PNF record.
Description	Details that provide more information about the PNF record, such as its function and purpose.
Vendor	The vendor that generated this PNF descriptor.
Alarm Severity	The alarm severity for each record, such as Critical, Major, Minor, Informational, Normal, Warning. By default, this field is not shown.
Alarm Suppression Description	A text description of alarm suppression. By default, this field is not shown.
Alarm Suppression Mode	An indicator that shows whether alarm suppression is on or off. By default, this field is not shown.
Created	The timestamp that the PNF record was created. By default, this field is not shown.
DescriptionId	The identifier for the PNF descriptor from which the PNF record was created. By default, this field is not shown.
Equipment	A reference to the physical resource that the network service is using. By default, this field is not shown.
IP Address	The resource (equipment) IP address. By default, this field is not shown.
Network Service	The reference to the network service record in which the PNF is participating. By default, this field is not shown.
OAM Reference	This is not currently implemented. By default, this field is not shown.
Updated	The timestamp for the latest PNF record changes. By default, this field is not shown.
Version	The PNF descriptor version from which the PNF record was created. By default, this field is not shown.

Editing Pnf Record Window

Use the Editing Pnf Record window (Figure 17-26) to createand modify PNF records. This is where you define network connections between a physical device and other virtual network instances, such as network service, virtualized network functions, virtual links, and so on.

Accessthiswindowbynavigating to the Physical Network Function Recordsportlet, right clicking a row, and then selecting New or Edit.

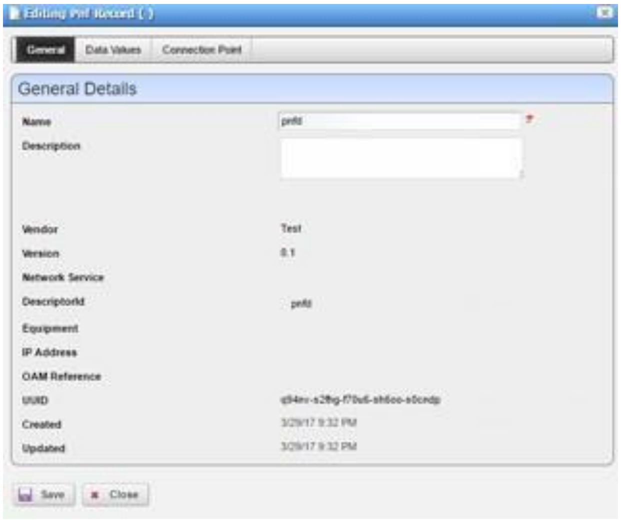

Name: pnfd

Description:

Vendor: Test

Version: 0.1

Network Service:

DescriptorId: pnfd

Equipment:

IP Address:

OAM Reference:

UUID: q94nv-s2fhg-f70u6-sh6oo-s0cndp

Created: 3/29/17 9:32 PM

Updated: 3/29/17 9:32 PM

Figure 17-26. Editing Pnf Record Window

The Editing Physical Network Function Descript or window has the following input fields and options (Table 17-27). The fields not listed are system-generated fields, such UUID, Created, and Updated.

Table 17-30. Editing Physical Network Function Descriptor Fields/Options (Sheet 1 of 2)

Field/Option	Description
Name	Enter a name for the PNF record you are creating. If you are modifying the PNF descriptor, you can change the name. This field is required.
Description	Enter an optional description that provides more details about the PNF.
Vendor	Enter the vendor that generated this PNF descriptor.
Version	Enter the PNF descriptor version from which the PNF record was created.
Network Service	Enter a reference to the network service record in which the PNF is participating.
DescriptorId	Enter an identifier for the PNF descriptor from which the PNF record was created.
Equipment	Enter a reference to the physical resource that the network service is using.
IP Address	Enter the resource (equipment) IP address.
OAM Reference	This is not currently implemented.

Table 17-30. Editing Physical Network Function Descriptor Fields/Options (Sheet 2 of 2)

Field/Option	Description
Data Values	Define one or more data value records. Enter a name, description, data type, and value. The name and data type are required fields. The system generates the UUID, Created, and Updated information. Data values are typically captured at deploy time. They are either entered directly by the user or extracted from the resulting deployed VDU. The types of values is completely dependent upon the network service. For example: • A VNF may need to know what the gateway IP address is for the service so that it can properly configure the service. • A value is extracted from the OpenStack system is the heat stack ID, that ID is then used at a later date when undeploying the service.
Connection Points	Define one or more connection points between the PNF and other network instances, such as network service, virtualized network functions, virtual links, *etc*. The name and connection type are required. Valid connection types are EthernetPort, IPVPN, OTHER, PNIC, PPORT, VNIC (virtual network interface controller), VPORT, VPORTLAN, VPORT WAN.

Detail Portlets

The Details portlets display the general details for the selected network descriptor or record. Exceptwhere noted, the general details are the samefor the network servicedescriptor and network service record.

You can also view alarms and event details and history details related to the selected Physical Network Function descriptor or record.

General Details

Use the Physical Network Function (PNF) General details (Figure 17-27) to viewidentification and version information, connection points, and a reference tree of items related to the selected physical network function descriptor or record.

Access this information by navigating to the appropriate PNF portlet, right clicking arow, and then selecting Details.

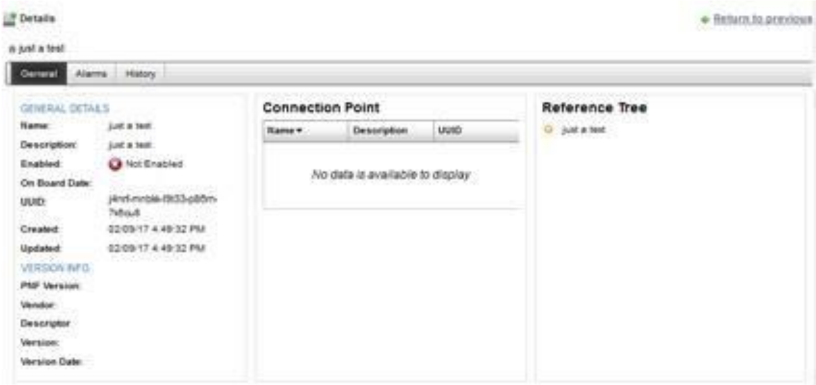

a just a test

GENERAL DETAILS

Name: just a test

Description: just a test

Enabled: Not Enabled

On Board Date:

UUID: j4nd-mnble-l9t33-p88m-7x8ou8

Created: 02/09/17 4:49:32 PM

Updated: 02/09/17 4:49:32 PM

VERSION INFO

PNF Version:

Vendor:

Descriptor

Version:

Version Date:

Connection Point

Name ▾	Description	UUID

No data is available to display

Reference Tree

just a test

Figure 17-27. Physical Network Function General Details

Alarms and Event Details

Use the Physical Network Function (PNF) Alarms details (Figure 17-28) to monitor and manage alarms and view event history for the selected PNF descriptor or record.

You can perform the same actions on an alarm or event history from this location as you can from the Alarms and Event History portlets, respectively. See Alarms on page 276 for details about the Alarms and Event History portlets.

Accessthisinformation bynavigating to theappropriate PNF portlet, right clicking arow, selecting Details, and then clicking the Alarms tab.

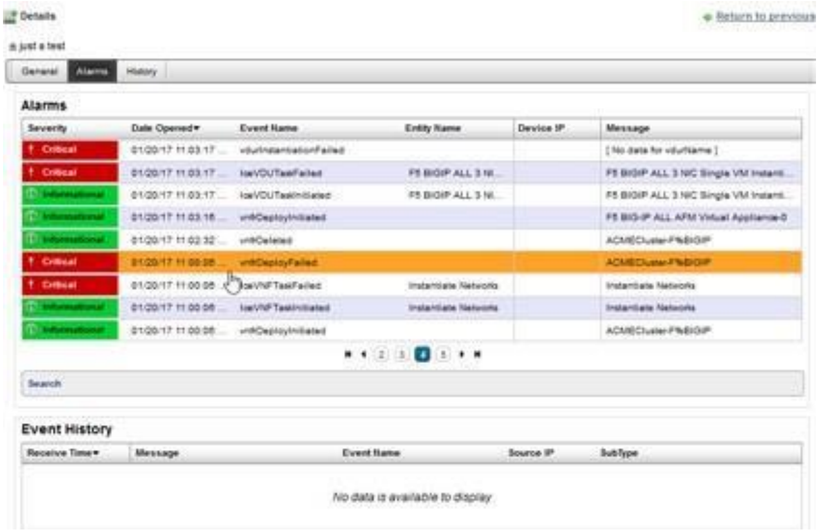

Alarms

Severity	Date Opened▾	Event Name	Entity Name	Device IP	Message
Critical	01/20/17 11:03:17 ...	vdurInstantiationFailed			[No data for vdurName]
Critical	01/20/17 11:03:17 ...	iaeVDUTaskFailed	F5 BIGIP ALL 3 NI...		F5 BIGIP ALL 3 NIC Single VM Instant...
Informational	01/20/17 11:03:17 ...	iaeVDUTaskInitiated	F5 BIGIP ALL 3 NI...		F5 BIGIP ALL 3 NIC Single VM Instant...
Informational	01/20/17 11:03:16 ...	vnfDeployInitiated			F5 BIG-IP ALL AFM Virtual Appliance-0
Informational	01/20/17 11:02:32 ...	vnfDeleted			ACMECluster-F%BIGIP
Critical	01/20/17 11:00:08 ...	vnfDeployFailed			ACMECluster-F%BIGIP
Critical	01/20/17 11:00:08 ...	iaeVNFTaskFailed	Instantiate Networks		Instantiate Networks
Informational	01/20/17 11:00:08 ...	iaeVNFTaskInitiated	Instantiate Networks		Instantiate Networks
Informational	01/20/17 11:00:08 ...	vnfDeployInitiated			ACMECluster-F%BIGIP

Event History

Receive Time▾	Message	Event Name	Source IP	SubType

No data is available to display

Figure 17-28. Physical Network Function Alarms Details Portlet

History Details

Use the Physical Network Function (PNF) History details (Figure 17-29) to view a list of actions (audit trail) performed on the selected PNF descriptor or record.

For the selected action, you can view the job, aging policy, delete the action, view as a portable document format (PDF), or share this action with other users on your system.

Access this information by navigating to the appropriate Physical Network Function portlet, right-clicking a row, selecting Details, and then clicking the History tab.

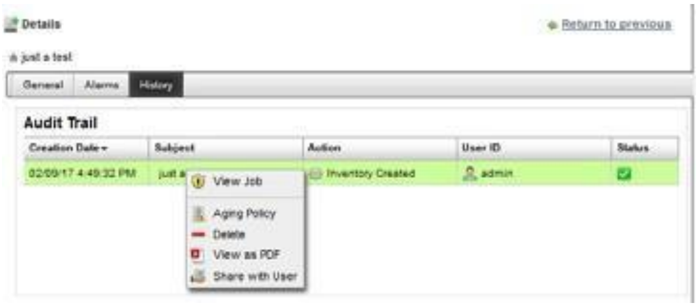

Audit Trail

Creation Date▾	Subject	Action	User ID	Status
02/09/17 4:49:32 PM	just a	Inventory Created	admin	✓

Figure 17-29. Physical Network Function History Portlet

SDN Controllers Portlet

Use the SDN Controllers portlet (Figure 17-30) to view, create and maintain software-defined network (SDN) controllers.

From the SDN Controllers portlet, asystem administrator createsand maintains SDN controllers and operators (if they have the proper permissions) view SDN controllers to gain knowledge.

The SDN Controllers portlet also provides amaximized viewthat includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

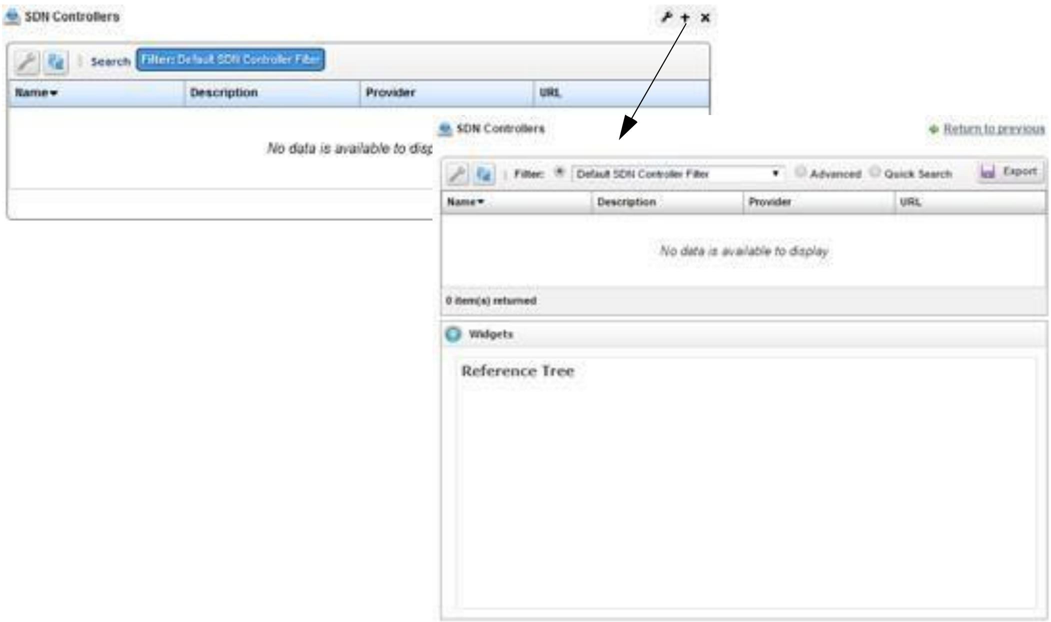

Figure 17-30. SDN Controllers Portlet

Pop-Up Menu

The SDN Controllers portlet provides access to the following pop-up menu options. Right-click a row to access these options (Table 17-31).

Table 17-31. SDN Controllers Pop-Up Menu Options (Sheet 1 of 2)

Option	Description
New	Opens the Editing SDN Controller window , where you create an SDN controller.
Edit	Opens the Editing SDN Controller window, where you modify an SDN controller.
Details	Displays the general details, such as identification information, version information, a reference tree, and so on.
Delete	Removes the selected SDN controller from the system. A confirmation message is displayed giving you a chance to change your mind.
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected PNF. Select an audit record and the job record displays its status and any other job-related information.

Table 17-31. SDN Controllers Pop-Up Menu Options (Sheet 2 of 2)

Option	Description
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection.

Columns

Other than the general navigation and configuration options, the SDN Controllers portlet includes the following columns (Table 17-32).

You can view the value for most of the hidden columns by right clicking a row and selecting Details. The other option is to add the column byclicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-32. SDN Controllers Columns

Column	Description
Name	The text that identifies each the SDN controller.
Description	A detailed description of the SDN controller’s purpose.
Provider	The vendor providing the SDN controller, such as OpenDaylight.
URL	The file server or disk location where the SDN controller is stored.
Created	The system-generated timestamp in which the SDN controller was created. By default, this field is not shown.
Password	The password associated with the provided user ID. By default, this field is not shown.
Priority	A relative number indicating which SDN controller should be chosen over another, where multiple SDN controllers are available. By default, this field is not shown.
Service ID	The internal service identifier for the SDN controller captured in the service model for reference purposes. By default, this field is not shown.
UUID	A unique 128-bit, system generated value assigned to each SDN controller created. By default, this field is not shown.
Updated	The system-generated timestamp for the latest record modifications. Initially this date is the same as the Created date. By default, this field is not shown.
User	The user ID assigned to this SDN controller during creation. By default, this field is not shown.

Editing SDN Controller Window

Use the Editing SDN Controller window (Figure 1 7-31) to create and modify SDN controllers.

Accessthiswindowbynavigatingto the SDN Controllersportlet, right-clicking, and thenselecting New.

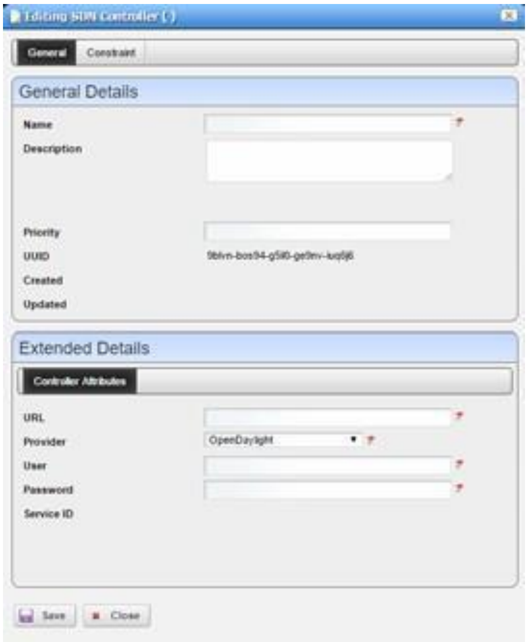

UUID 9blvn-bos94-g5l0-ge9mv-luq6j6

Provider OpenDaylight

Figure 17-31. Editing SDN Controller Window

The Editing SDN Controller window has the following input fields and options (Table 17-33). The fields not listed are system-generated fields, such UUID, Created, and Updated.

Table 17-33. Editing SDN Controller Fields/Options

Field/Option	Description
Name	Enter a name for the SDN controller that you are creating. This field is required.
Description	Enter an optional description that provides more details about the SDN controller’s purpose.
Priority	Enter a relative number indicating which SDN controller should be chosen over another, where multiple SDN controllers are available.
URL	Enter file server or disk location where the OSS instance is stored. This field is required.
Provider	Select the vendor providing the SDN controller. This field is required.
User	Enter the user ID. This field is required.
Password	Enter the password assigned the provided user ID. This field is required.

Vendor Licensing Portlets

Use the following license portlets to view, create, and manage vendor license information:

- License Accounts Portlet
- License Descriptors Portlet
- License Records Portlet

License Accounts Portlet

Use the License Accounts portlet (Figure 17-32) to view, create and maintain vendor license accounts.

Fromthe License Accountsportlet, asystemadministrator, networkdesigner, andoperators(ifthey have the proper permissions) creates and maintains licenseaccounts.

The License Accounts portlet also providesa maximized view that includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

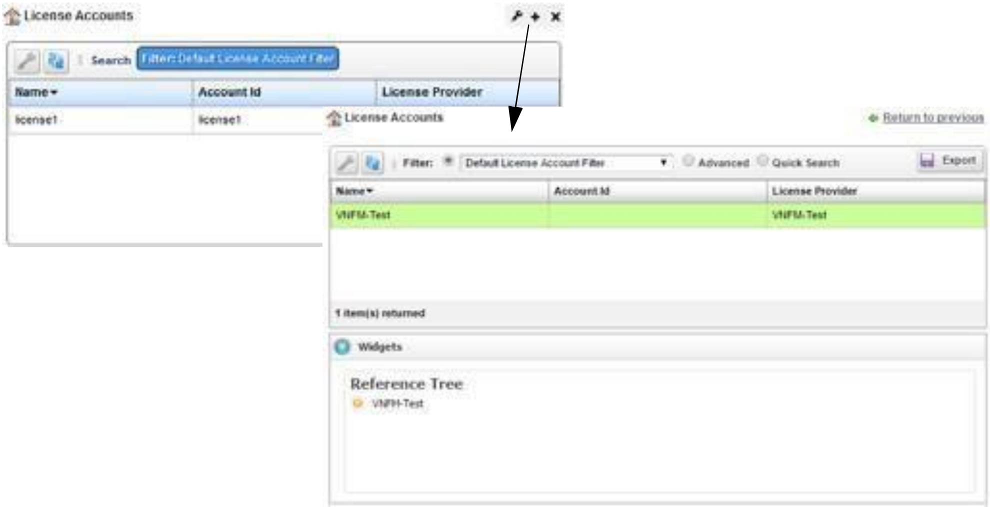

Name ▾	Account Id	License Provider
license1	license1	

Name ▾	Account Id	License Provider
VNFM-Test		VNFM-Test

1 item(s) returned

Reference Tree

VNFM-Test

Figure 17-32. License Accounts Portlet

Pop-Up Menu

The License Accounts portlet provides access to the following pop-up menu options. Right-click a row to access these options (Table 17-34).

Table 17-34. License Accounts Pop-Up Menu Options

Option	Description
New	Opens the Editing License Account window, where you create a license account.
Edit	Opens the Editing License Account window, where you modify a license account.
Details	Displays the general license details, such as identification information, account ID, and a reference tree.
Delete	Removes the selected license account from the system. A confirmation message is displayed giving you a chance to change your mind.
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected license account. Select an audit record and the job record displays its status and any other job-related information.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection.

Columns

Other than the general navigation and configuration options, the License Accounts portlet includes the following columns (Table 17-35).

Youcan viewthe value for most of the hidden columns byright clicking a rowand selecting Details. The other option is to add the column by clicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-35. License Accounts Columns

Column	Description
Name	The text that identifies each license account.
Account Id	An optional user-defined account identifier.
License Provider	The vendor that issued the license.
Created	The system-generated timestamp in which the license account was created. By default, this field is not shown.
Description	A detailed description of the license account’s purpose. By default, this field is not shown.
UUID	A unique 128-bit, system generated value assigned to each license account created. By default, this field is not shown.
Updated	The system-generated timestamp for the latest record modifications. Initially this date is the same as the Created date. By default, this field is not shown.

Editing License Account Window

Use the Editing License Account window (Figure 17-33) to create and modify license accounts.

Accessthiswindowbynavigatingtothe License Accountsportlet, right-clicking, andthenselecting Edit or New.

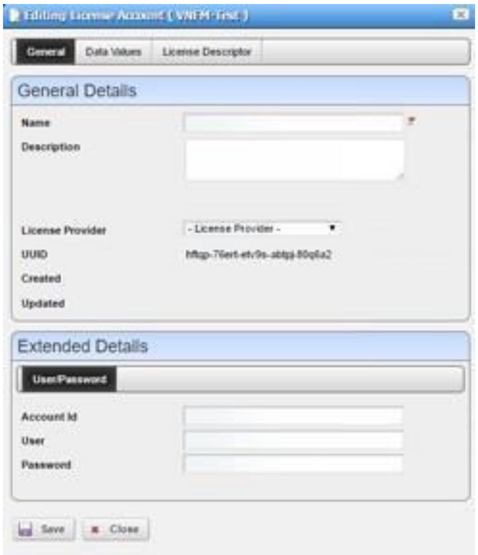

License Provider: - License Provider -

UUID: hftqp-76ert-etv9s-abtpj-80q6a2

Figure 17-33. Editing License Account Window

The Editing License Account window hasthe followinginput fieldsand options (Table 17-36). The fields not listed are system-generated fields, such UUID, Created, and Updated.

Table 17-36. Editing License Account Fields/Options (Sheet 1 of 2)

Field/Option	Description
Name	Enter a name for the license account that you are creating. This field is required.

Table 17-36. Editing License Account Fields/Options (Sheet 2 of 2)

Field/Option	Description
Description	Enter an optional description that provides more details about the license descriptor’s purpose.
License Provider	Enter the name for license issuer.
Account Id	Enter an optional user-defined account identifier.
User	Enter an optional user ID attached to this account.
Password	Enter the password associated to the user ID for this account.
Data Values	Enter any necessary data values.
License Descriptor	Select a valid license descriptor for this account.

License Descriptors Portlet

Use the License Descriptors portlet (Figure 17-34) to view, create and maintain vendor license descriptors.

From the License Descriptors portlet, a system administrator creates and maintains license descriptors, a network designer defineslicensingand characteristics, and operators (if they have the proper permissions) view license descriptors to gain knowledge.

The License Descriptors portlet also provides a maximized view that includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

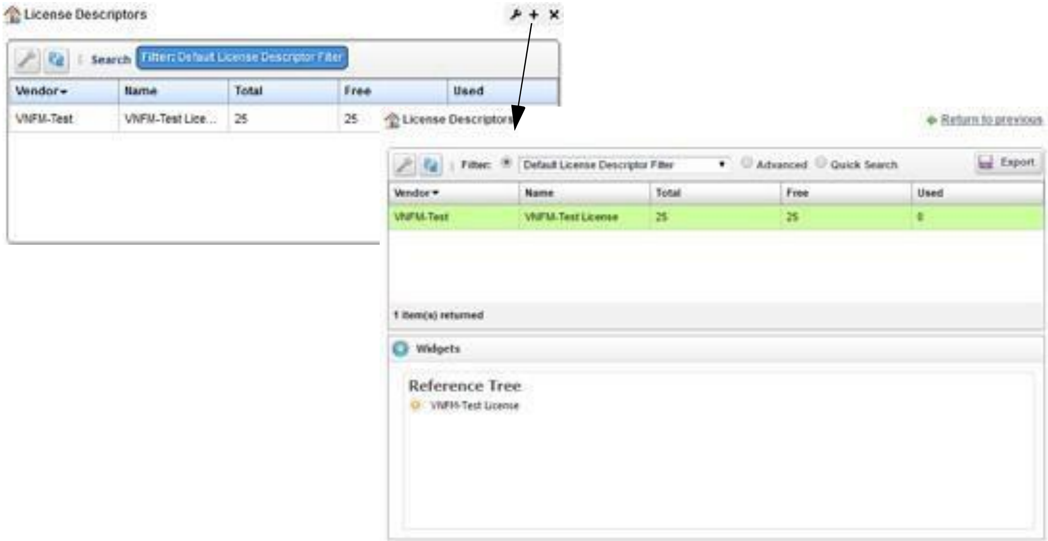

Vendor	Name	Total	Free	Used
VNFM-Test	VNFM-Test Lice...	25	25	

Vendor	Name	Total	Free	Used
VNFM-Test	VNFM-Test License	25	25	0

1 item(s) returned

Figure 17-34. License Descriptors Portlet

Pop-Up Menu

The License Descriptors portlet provides access to the followingpop-up menu options. Right-click a row to access these options (Table 17-37).

Table 17-37. License Descriptors Pop-Up Menu Options

Option	Description
New	Opens the Editing License Descriptor window, where you create a license descriptor.
Edit	Opens the Editing License Descriptor window, where you modify a license descriptor.
Details	Displays the generallicense details, suchasidentification information, properties, resources, and a reference tree.
Delete	Removes the selected license descriptor from the system. A confirmation message is displayed giving you a chance to change your mind.
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected license descriptor. Select an audit record and the job record displays its status and any other job-related information.
Import/Export	Provides the following actions when available for the selected descriptor: • Import retrieves a file containing XML license descriptions. Some imports can come from a URL. • Export Selection exports the selected license description to an XML file. • Export All exports all license descriptions to an XML file. Click Download Export File to specify where to save the file. The Import/Export option is useful as a backup or to share descriptors with other projects.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Right clicking Note: Sharing only handles one item so it uses the first one in the selection.**

Columns

Other than the general navigation and configuration options, the License Descriptors portlet includes the following columns (Table 17-38).

Youcan viewthe value for most of the hidden columns byright clicking a rowand selecting Details. The other option is to add the column by clicki ng the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-38. License Descriptors Columns (Sheet 1 of 2)

Column	Description
Vendor	The license account name.
Name	The text that identifies each the license descriptor.
Total	The total number of licenses available.
Free	The total number of licenses **not** registered (still available).
Used	The total number of licenses registered (in use).
Account Required	An indicator that shows whether the account is required (checkmark) or not (X). By default, this field is not shown.

Table 17-38. License Descriptors Columns (Sheet 2 of 2)

Column	Description
Created	The system-generated timestamp in which the license descriptor was created. By default, this field is not shown.
Description	A detailed description of the license descriptor's purpose. By default, this field is not shown.
Enabled	An indicator that shows whether the license descriptor is available (checkmark) or not (X). By default, this field is not shown.
Expiration Policy	The action to take when the license expires, such as Degrade, Fail, or Unknown. By default, this field is not shown.
Foreign Id	An external system identifier for reference purposes. By default, this field is not shown.
Key Source	The location where the license keys are managed. By default, this field is not shown.
License Provider	The name for license issuer.
Low Threshold	The value used to notify you when the license pool is low. By default, this field is not shown.
Release Policy	The policies used to release a license, such as Delete, Retain, and Reuse. By default, this field is not shown.
Type	The type of license. By default, this field is not shown.
UUID	A unique 128-bit, system generated value assigned to each license descriptor created. By default, this field is not shown.
Updated	The system-generated timestamp for the latest record modifications. Initially this date is the same as the Created date. By default, this field is not shown.
Version	The license descriptor version. By default, this field is not shown.

Editing License Descriptor Window

Use the Editing License Descriptor window (Figure 17-35) to create and modify license descriptors.

Access this window by navigating to the License Descriptors portlet, right-clicking, and then selecting Edit or New.

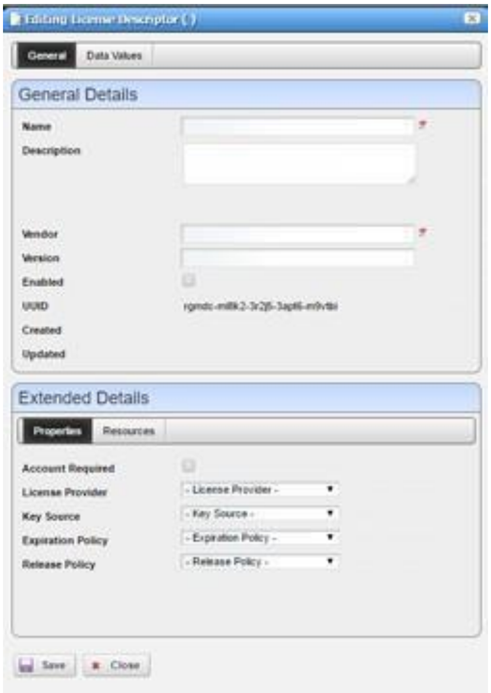

Figure 17-35. Editing License Descriptor Window

The Editing License Descriptor window has the following input fields and options (Table 17-39). The fields not listed are system-generated fields, such UUID, Created, and Updated.

Table 17-39. Editing License Descriptor Fields/Options

Field/Option	Description
Name	Enter a name for the license descriptor that you are creating. This field is required.
Description	Enter an optional description that provides more details about the license descriptor’s purpose.
Vendor	Enter the license account name. This field is required.
Version	Enter the license descriptor version.
Enabled	Select if you want to enable this descriptor. By default, the descriptor is not enabled.
Account Required	Select to capture the license account information.
License Provider	Enter the name for license issuer.
Key Source	Select the key source for this license descriptor, such as External, Internal, or MANO.
Expiration Policy	Select which expiration policy to use, such as Degrade, Fail, or Unknown.
Release Policy	Select which release policy to use, such as Delete, Retain, or Reuse.
Resources	Enter optional license values for the following resources: • Free • Used • Total • Low Threshold

License Records Portlet

Use the License Records portlet (Figure 17-36) to view the list of license records loaded into the system for distribution.

From the License Records portlet, a system administrator, network designer, and operators (if they have the proper permissions) view license records to gain knowledge of available licenses.

The License Records portlet also provides a maximized view that includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same columns by default the same. See Expanded Portlet on page 130 for a description of its standard options provided.

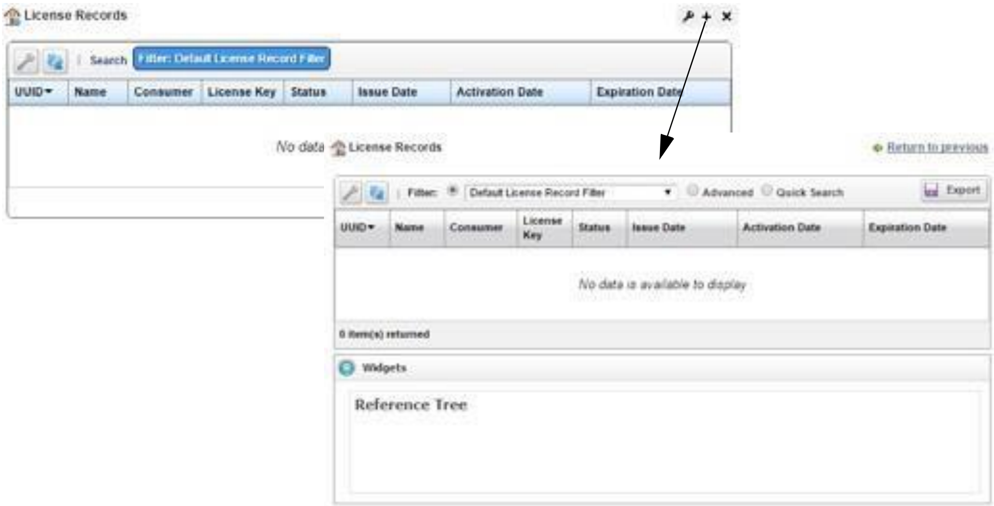

Figure 17-36. License Record Portlet

Other than the general navigation and configuration options, the License Descriptors portlet includes the following columns (Table 17-40).

You can view the value for most of the hidden columns by right clicking a row and selecting Details. The other option is to add the column byclicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-40. License Records Columns (Sheet 1 of 2)

Column	Description
UUID	A unique 128-bit, system generated value assigned to each license record created.
Name	The text that identifies each the license record.
Consumer	The network service or VNF record that is currently using the registered license.
License Key	The actual license key value.
Status	An indicator that shows whether the license is free or used.
Issue Date	The date the license was issued.
Activation Date	The date the license was activated.
Expiration Date	The date the license expires.
Created	The system-generated timestamp in which the license record was created. By default, this field is not shown.

Table 17-40. License Records Columns (Sheet 2 of 2)

Column	Description
Customer	The customer using the license - the customer associated with the consuming NFV record. By default, this field is not shown.
Description	A detailed description of the license record's purpose. By default, this field is not shown.
Descriptor	The license descriptor associated with the license inventory. By default, this field is not shown.
Release Date	The date the license was released. By default, this field is not shown.
Updated	The system-generated timestamp for the latest record modifications. Initially this date is the same as the Created date. By default, this field is not shown.
Vendor	The VNF vendor. By default, this field is not shown.

Virtualized Network Function Portlets

This section describes the following Virtualized Network Function (VNF) portlets:

- Virtual Network Function Descriptors Portlet
- Virtual Network Function Records Portlet
- Details Portlets

Virtual Network Function Descriptors Portlet

Use the Virtual Network Function Descriptors portlet (Figure 17-37) to maintain a list of virtualized network function (VNF) descriptors. These descriptorsare templates used to create VNF records.

From the Virtual Network Function Descriptors portlet, a system administrator creates and maintains the descriptors, a network designer implements VNFs, and operators (if they have the proper permissions) view descriptors to gain knowledge.

The Virtual Network Function Descriptors portlet also provides a maximized view that includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

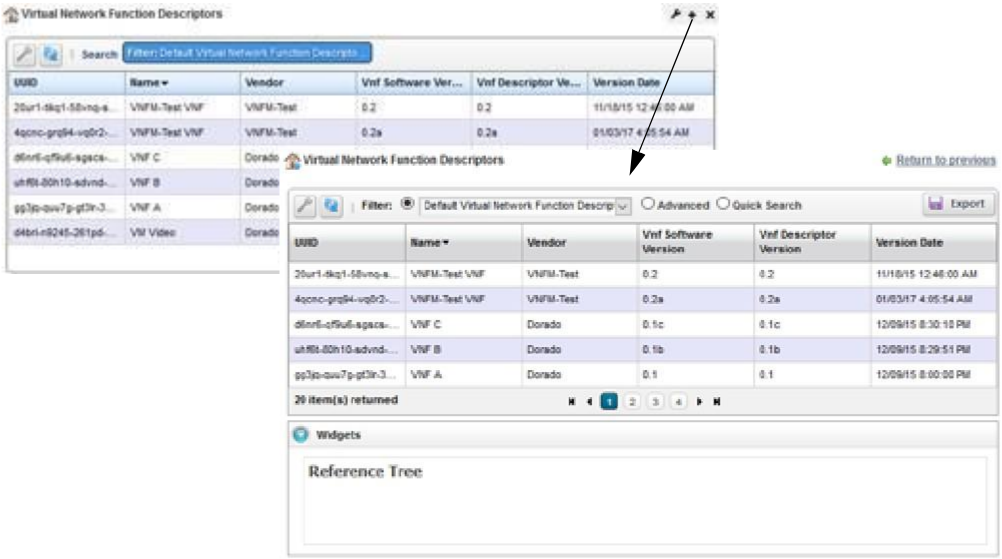

Virtual Network Function Descriptors
Search
Filter: Default Virtual Network Function Descripto...

UUID	Name ▾	Vendor	Vnf Software Ver...	Vnf Descriptor Ve...	Version Date
20ur1-dkq1-58vnq-a...	VNFM-Test VNF	VNFM-Test	0.2	0.2	11/18/15 12:46:00 AM
4qcnc-grq94-vq0r2-...	VNFM-Test VNF	VNFM-Test	0.2a	0.2a	01/03/17 4:05:54 AM
d6nr6-qf9u6-agsca-...	VNF C	Dorado			
uhf8t-80h10-advnd-...	VNF B	Dorado			
gg3jp-quu7p-gt3ir-3...	VNF A	Dorado			
d4brl-n9245-261pd-...	VM Video	Dorado			

Virtual Network Function Descriptors
Return to previous
Filter: ● Default Virtual Network Function Descrip
○ Advanced ○ Quick Search
Export

UUID	Name ▾	Vendor	Vnf Software Version	Vnf Descriptor Version	Version Date
20ur1-dkq1-58vnq-a...	VNFM-Test VNF	VNFM-Test	0.2	0.2	11/18/15 12:46:00 AM
4qcnc-grq94-vq0r2-...	VNFM-Test VNF	VNFM-Test	0.2a	0.2a	01/03/17 4:05:54 AM
d6nr6-qf9u6-agsca-...	VNF C	Dorado	0.1c	0.1c	12/09/15 8:30:10 PM
uhf8t-80h10-advnd-...	VNF B	Dorado	0.1b	0.1b	12/09/15 8:29:51 PM
gg3jp-quu7p-gt3ir-3...	VNF A	Dorado	0.1	0.1	12/09/15 8:00:00 PM

20 item(s) returned

Figure 17-37. Virtual Network Function Descriptors Portlet

Pop-Up Menu

The Virtual Network Function Descriptors portlet providesaccess to the following pop-up menu options. Right-click a row to access these options (Table17-41).

Table 17-41. Virtual Network Function Descriptors Pop-Up Menu Options (Sheet 1 of 2)

Option	Description
New	Opens the Editing Virtual Network Function Descriptor window , where you define a VNF descriptor and any related data values, lifecycle events, deployment flavors, connection points, VDU dependencies, virtual links, and VDUs.
Edit	Opens the Editing Virtual Network Function Descriptor window , where you modify a VNF descriptor and any related data values, lifecycle events, deployment flavors, connection points, VDU dependencies, virtual links, and VDUs.
Details	Displays the general details, such as identification information, version information, monitor parameters, and a reference tree. You can also view VDU information , deployment flavor details , connections details, alarms and event details , and history details by selecting the appropriate tab.
Discover Vnfd	Opens the Discover VNF Descriptor window, where you select a target VIM on which to create the VNF descriptor, create the VNF descriptor from OpenStack artifacts parameters, and then execute the Create Vnf Descriptor from OpenStack Artifacts task.
Branch	Opens the Branch VNF Descriptor window, where you execute the Vnf Descriptor Branch to new version task or schedule the task to run later. Executing this task ensures that you have the latest descriptor.
Stage	Opens the Stage new VNF Record window, where you set staging parameters and then run the VNF stage tasks or schedule them to run later. If the staging tasks run successfully, you can then deploy the VNF.
Delete	Removes the selected VNF descriptor from the system.

Table 17-41. Virtual Network Function Descriptors Pop-Up Menu Options (Sheet 2 of 2)

Option	Description
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected VNF. Select an audit record and the job record displays its status and any other job-related information.
Import/Export	Provides the following actions when available for the selected image: • Import retrieves a file containing XML VNF descriptions. Some imports can come from a URL. • Export Selection exports the selected VNF description to an XML file. • Export All exports all VNF descriptions to an XML file. Click Download Export File to specify where to save the file. The Import/Export option is useful as a backup or to share descriptors with other projects.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Right clicking Note: Sharing only handles one item so it uses the first one in the selection.**

Columns

Other than the general navigation and configuration options, the Virtual Network Function Descriptors portlet includes the following columns (Table 17-42).

You can view the value for most of the hidden columns by right clicking a row and selecting Details. The other option is to add the column by clicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-42. Virtual Network Function Descriptors Columns (Sheet 1 of 2)

Column	Description
UUID	A unique 128-bit, system generated value assigned to each VNF descriptor created.
Name	The user-entered name that identifies the VNF descriptor.
Vendor	The vendor VNF used to generate this VNF descriptor, such as F5, Cisco, *etc.*
Vnf Software Version	The version of VNF software used to generate this VNF.
Vnf Descriptor Version	The VNF descriptor version.
Version Date	The VNF descriptor timestamp.
Auto-configure	An indicator that shows whether configuration tasks are queued (checkmark) once the instantiation completes or not (X). By default, this field not shown.
Created	The system-generated timestamp for the description creation. By default, this field is not shown.
Description	The optional detailed description provided during descriptor creation. By default, this field is not shown.
Enabled	An indicator that shows whether the VNF descriptor is available (checkmark) or not (X). By default, this field is not shown.
Monitor Attributes	Monitor attributes provided during descriptor creation. By default, this field is not shown.
Monitoring Parameters	Monitoring parameters provided during descriptor creation. By default, this field is not shown.

Table 17-42. Virtual Network Function Descriptors Columns (Sheet 2 of 2)

Column	Description
Notifications	The latest activity for the descriptor, such as Created. By default, this field is not shown.
On Board Date	The timestamp for the descriptor onboarding. By default, this field is not shown.
Parent Version	The name of the descriptor from which the new descriptor was branched. By default, this field is not shown.
Stand-alone	An indicator that shows whether the VNF descriptor deploys without (checkmark) or with (X) a network service. By default, this field is not shown.
Updated	The system-generated timestamp for the latest descriptor changes. By default, this field is not shown.

Editing Virtual Network Function Descriptor Window

Use the Editing Virtual Network Function (VNF) Descriptor window (Figure 17-38) to view, define, or modify a VNF descriptor and any of the following related items listed in the navigation tree:

- Data Values lists any defined data values.
- Lifecycle Events provides access to any defined list of events for the Heal, Shutdown, Configure, Scale in, Instantiate, Terminate, Scale out, and Diagnostics lifecycle events.
- Deployment Flavors lists the default or defined deployment flavors.

  Therearedeploymentflavorsavailablebydefault. However, yoursystemadministrator can create and maintain the deployment flavors available.
- Connection Points lists any defined connection points.
- VDU (virtualization deployment unit) Dependencies lists anydefined VDU dependencies.
- Virtual Links lists any defined virtual links.
- VDUs lists any defined VDU descriptors.

Select an item from the navigation tree and the available definitions are displayed in the general or extendeddetails. Ifyouhaveadministrative permissions, youcanadd, modify, ordeletedefinitions.

Access thiswindowbynavigating to the Virtual Network Function Descriptors portlet, right-clicking a descriptor or the portlet, and the selecting New or Edit.

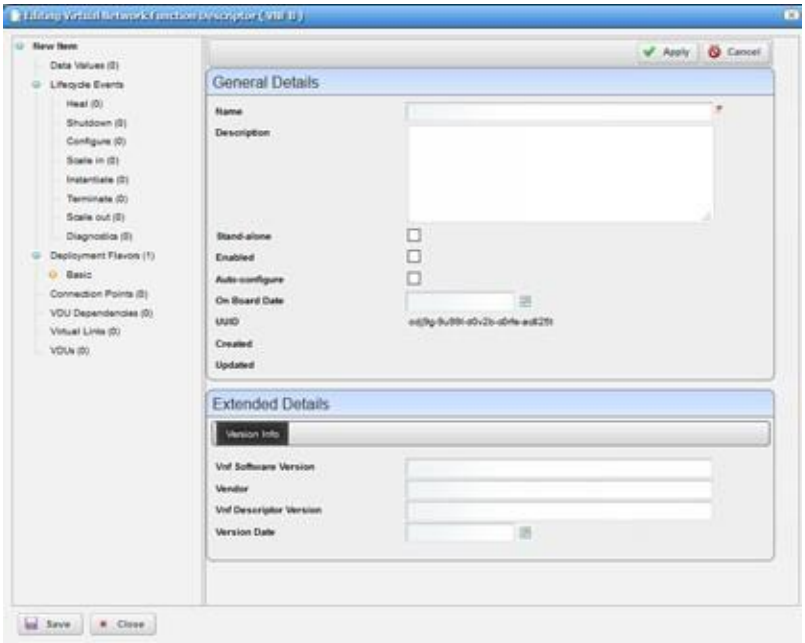

Figure 17-38. Editing Virtual Network Function Descriptor Window

The Editing Virtual Network Function Descriptor s window has the following editable fields and options (Table 17-43). The fields not listed are system-generated fields, such UUID, Created, and Updated.

Table 17-43. Editing Virtual Network Function Descriptor Fields/Options

Field/Option	Description
Name	Enter a VNF descriptor name that identifies its purpose.
Description	Enter more details that describe the VNF’s purpose, version, *etc.*
Stand-alone	Select to deploy the VNF descriptor deploys without a network service. The default is to deploy with a network service (not selected).
Enabled	Select to make the VNF descriptor is available. By default, the VNF descriptor is not available (not selected.)
Auto-configure	Select to automatically queue configuration tasks once the VNF descriptor instantiation completes. By default, configuration tasks are not automatically queued (not selected).
On Board Date	Select a date and time instantiate the VNF descriptor. Once instantiated, this timestamp changes to the instantiation date/time.
Extended Details	Enter the appropriate information for the VNF descriptor or the selected element, during creation: • For the VNF descriptor, enter Vendor Info (Vnf Software Version, Vendor name, Vnf Descriptor Version, Version Date) and Monitor Parameters (monitor attributes). • For Virtual Links, enter Properties (connectivity type, root bandwidth, leaf bandwidth, test access), VNFC Connection Points, and Qos options • For a VDU, enter Image Requirements (VIM image, CPU, memory, disk, bandwidth, resource model, system object ID), Image Properties (HA, min/ max instances), and Monitoring Parameters

Discover VNF Descriptor Window

Use the Discover VNF Descriptor window (Figure 17-39) to create a VNF descriptor from OpenStack artifacts on a target VIM.

Access thiswindowbynavigating to the Virtual Network Function Descriptors portlet, right-clicking a descriptor, and then selecting Discover Vnfd.

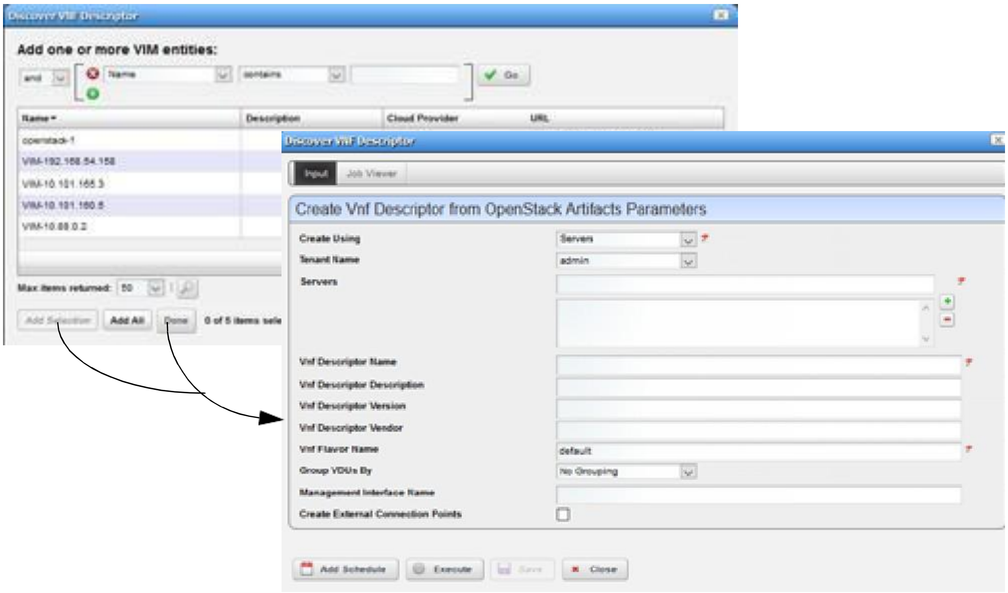

Discover VIM Descriptor

Add one or more VIM entities:

and | Name | contains

Name	Description	Cloud Provider	URL
openstack-1			
VIM-192.168.54.158			
VIM-10.101.165.3			
VIM-10.101.160.5			
VIM-10.88.0.2			

Max items returned: 50

0 of 5 items sele

Discover VNF Descriptor

Create Vnf Descriptor from OpenStack Artifacts Parameters

Create Using: Servers
Tenant Name: admin
Servers:
Vnf Descriptor Name:
Vnf Descriptor Description:
Vnf Descriptor Version:
Vnf Descriptor Vendor:
Vnf Flavor Name: default
Group VDUs By: No Grouping
Management Interface Name:
Create External Connection Points: □

Figure 17-39. Discover VNF Descriptor Window

The Discover VNF Descriptor window has the following fields and options (Table 17-44).

Table 17-44. Discover VNF Descriptor Fields/Options (Sheet 1 of 2)

Field/Option	Description
Conditional Search Parameters	Define conditional search criteria to find the appropriate VIM.
Available VIM List	Select a VIM on which to create a VNF descriptor and then click Add Selection > Done. The Create Vnf Descriptor from OpenStack Artifacts Parameters input fields are displayed.
Create Using	Select whether to use Servers or Stacks to discover a VNF descriptor.
Tenant Name	Select the target project. When you select a project, the Servers/Stacks field populates with available choices.
Servers/Stacks	Select which servers/stacks to use. This field name and available choices vary depending on your Create Using and Tenant Name selection.
Vnf Descriptor Name	Enter a name identifier for the new VNF descriptor.
Vnf Descriptor Description	Enter a detailed description for the new VNF descriptor.
Vnf Descriptor Version	Enter a version for the new VNF descriptor.
Vnf Descriptor Vendor	Enter the vendor VNF used to generate this descriptor.
Vnf Flavor Name	Enter a VNF deployment flavor name or leave the value as default.
Group VDUs By	Select whether to group VDUs by Flavor, Image, Name Pattern, or No Grouping. The default is No Grouping.

Table 17-44. Discover VNF Descriptor Fields/Options (Sheet 2 of 2)

Field/Option	Description
Management Interface Name	Enter the port through which the VDU is managed.
Create External Connection Points	Select whether to create external connection points. The default is not to create external connection points (not selected).
Add Schedule	Set the following schedule parameters used to automatically execute the Create VNF Descriptor from OpenStack Artifacts task on the specified target VIM. • Starting On sets the date and time to execute the task. • Recurrence sets whether the task executes Only Once, Only at Startup, every X minutes (Increment), or Every *X* minutes, hours, days, weekdays, weekend days, weeks, months, years. • Stopping On sets when the task stops executing. Valid values areNever, By Occurrence, or By Date and Time.
Execute	Runs the Create VNF Descriptor from OpenStack Artifacts task to create the selected descriptor on the target VIM.
Save	Preserves your created VNF descriptor configuration.
Close	Exits the Discover VNF Descriptor window.

Branch VNF Descriptor Window

Use the Branch VNF Descriptor window (Figure 17- 40) to create a descriptor with the latest/ improved version.

Accessthis windowbynavigating to the Virtual Network Function Descriptors portlet, right-clicking a descriptor, and then selecting Branch.

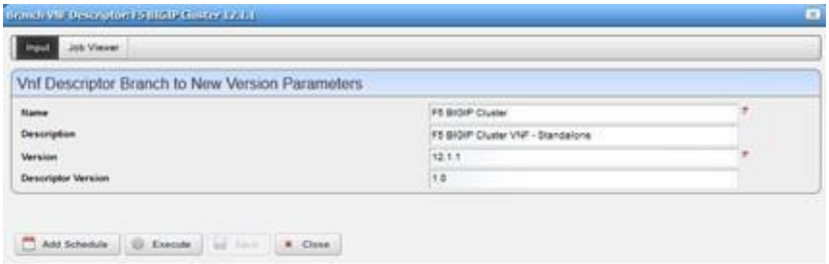

Vnf Descriptor Branch to New Version Parameters

Name: F5 BIGIP Cluster *

Description: F5 BIGIP Cluster VNF - Standalone

Version: 12.1.1 *

Descriptor Version: 1.0

Figure 17-40. Branch VNF Descriptor Window

The Branch VNF Descriptor window has the following fields and options (Table 17-45).

Table 17-45. Branch VNF Descriptor Fields/Options (Sheet 1 of 2)

Field/Option	Description
Name	Modify the displayed VNF identifier as needed.
Description	Modify the displayed VNF detailed description as needed.
Version	Modify the displayed VNF software version as needed.
Descriptor Version	Modify the displayed VNF descriptor version as needed.

Table 17-45. Branch VNF Descriptor Fields/Options (Sheet 2 of 2)

Field/Option	Description
Add Schedule	Set the following schedule parameters used to automatically execute the branch VNF descriptor to new version task. • Starting On sets the date and time to execute the task. • Recurrence sets whether the task executes Only Once, Only at Startup, every X minutes (Increment), or Every *X* minutes, hours, days, weekdays, weekend days, weeks, months, years. • Stopping On sets when the task stops executing. Valid values are Never, By Occurrence, or By Date and Time.
Execute	Runs the branch VNF descriptor to new version task.
Save	Preserves your branching configuration.
Close	Exits the Branch VNF Descriptor window.

Stage new VNF Record Window

Use the Stage new VNF Record window (Figure 17-41) to stage a VNF and then deploy the successfully staged VNF record.

Access this window by navigating to the Virtual Network Function Descriptor portlet, right clicking a descriptor, and then selecting Stage.

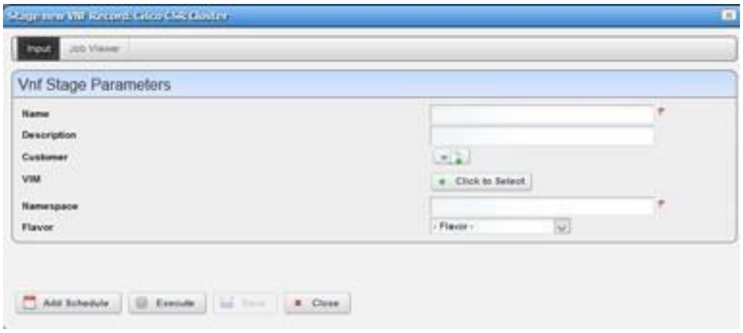

Figure 17-41. Stage new VNF Record Window

The Stage new VNF Record window has the following fields and options (Table 17-46).

Table 17-46. Stage new VNF Record Fields/Options (Sheet 1 of 2)

Field/Option	Description
Name	Enter a name that identifies the VNF you want to stage.
Description	Enter a more detailed description of the VNF you want to stage.
Customer	Select the person, group, or organization that requested the VNF. If the customer is not listed, create it.
VIM	Select the Virtualized Infrastructure Manager to which you to instantiate the VNF. Selecting a VIM populates the Namespace list.
Namespace	Select the name for the project that houses the VIM. This list is populated when you select a deployment VIM.
Flavor	Select a target deployment model (flavor). Valid values vary depending on the VNF selected.

Table 17-46. Stage new VNF Record Fields/Options (Sheet 2 of 2)

Field/Option	Description
Add Schedule	Set the following schedule parameters used to automatically execute the VNF staging task. • Starting On sets the date and time to execute the task. • Recurrence sets whether the task executes Only Once, Only at Startup, every X minutes (Increment), or Every *X* minutes, hours, days, weekdays, weekend days, weeks, months, years. • Stopping On sets when the task stops executing. Valid values are Never, By Occurrence, or By Date and Time.
Execute	Runs the VNF record stage task.
Save	Preserves your staging configuration.
Close	Exits the Stage new VNF Record window.

Virtual Network Function Records Portlet

Use the Virtual Network Function Records portlet (Figure 17-42) to maintain a list of virtualized network function (VNF) records.

From the Virtual Network Function Records portlet, operators can stage, deploy, and undeploy network services. A network designer would use this portlet to test a new or modified descriptor.

The Virtual Network Function Records portlet also provides a maximized view that includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

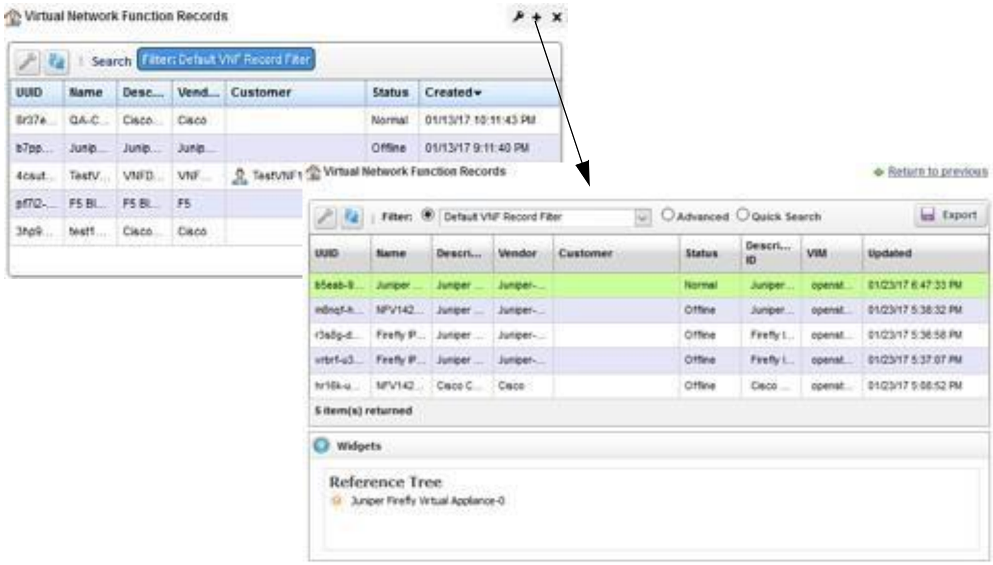

Virtual Network Function Records

Search | Filter: Default VNF Record Filter

UUID	Name	Desc...	Vend...	Customer	Status	Created
8r37e...	QA-C...	Cisco...	Cisco		Normal	01/13/17 10:11:43 PM
b7pp...	Junip...	Junip...	Junip...		Offline	01/13/17 9:11:40 PM
4csut...	TestV...	VNFD...	VNF...	TestVNF1		
pf7l2-...	F5 BI...	F5 BI...	F5			
3hp9...	test1...	Cisco...	Cisco			

Virtual Network Function Records

Return to previous

Filter: ● Default VNF Record Filter ○ Advanced ○ Quick Search Export

UUID	Name	Descri...	Vendor	Customer	Status	Descri... ID	VIM	Updated
b5eab-9...	Juniper ...	Juniper ...	Juniper-...		Normal	Juniper...	openst...	01/23/17 6:47:33 PM
m8nqf-h...	NFV142...	Juniper ...	Juniper-...		Offline	Juniper...	openst...	01/23/17 5:38:32 PM
r3a8g-d...	Firefly IP...	Juniper ...	Juniper-...		Offline	Firefly I...	openst...	01/23/17 5:36:58 PM
vrbrf-u3...	Firefly IP...	Juniper ...	Juniper-...		Offline	Firefly I...	openst...	01/23/17 5:37:07 PM
hr16k-u...	NFV142...	Cisco C...	Cisco		Offline	Cisco ...	openst...	01/23/17 5:08:52 PM

5 item(s) returned

Widgets

Reference Tree

Juniper Firefly Virtual Appliance-0

Figure 17-42. Virtual Network Function Records Portlet

Pop-Up Menu

The Virtual Network Function Records portlet provides access to the following pop-up menu options. Right-click a row to access these options (Table17-47).

Table 17-47. Virtual Network Function Records Pop-Up Menu Options (Sheet 1 of 2)

Option	Description
New	Opens the Stage VNF window, were you define a new VNF to stage and then deploy.
Edit	Opens the Editing VNF Record window, where
Details	Displays the general details, such as identification information, version information, deployment status, external references monitor parameters, health, connection points, and a reference tree. You can also view VDU information, alarms and event details, and history details by selecting the appropriate tab.
Deploy	Adds the VNF to the specified VIM by running the VNF Deploy task on the specified target, which includes processing the VNF Instantiate event for the VNF. You can also execute the Vnf Deploy action from the Actions portlet. Once deployment successfully completes, the record status shows as Instantiating. When instantiation completes, the status changes to Normal.
Undeploy	Removes the VNF from the specified VIM by running the VNF Undeploy task on the specified target, which includes processing the Terminate event for the VNF. You can also execute the Vnf Undeploy action from the Actions portlet. This action changes the record status to Offline.
Discover/Resync	Executes the Vnf Discovery task on the target VIM to identify VNF changes and performs a resync action to ensure VNFs are up to date.
Scaling	Provides access to the following scaling options: • **Scale In** executes the Scale-in event to reduce the VNF capacity. • **Scale Out** executes the Scale-out event to increase the VNF capacity. You can also execute the Vnf Scale In and Vnf Scale Out actions from the Actions portlet.
Start/Stop	Toggles between the ability to start or stop a VNF record depending on its current state. If you **stop** the record, a confirmation message is displayed. Click Stop to continue. This updates the record status to Stopped and any transactions after this point are lost. When you **start** activity, the record status returns to Normal. You can also execute the Vnf Stop and Vnf Start actions from the Actions portlet.
Resume/Suspend	Toggles between the ability to suspend a VNF’s record activity or resume a suspended VNF’s record activity. If you **suspend** the record, a confirmation message is displayed. Click Suspend to continue. This updates the record status to Suspended. Activity continues to collect and is passed through once you resume activity. When you **resume** activity, the record status returns to Normal. You can also execute the Vnf Suspend and Vnf Resume actions from the Actions portlet.

Table 17-47. Virtual Network Function Records Pop-Up Menu Options (Sheet 2 of 2)

Option	Description
Maintenance	Executes the following tasks on the selected VNF record: • **Diagnostics** runs the Diagnostics event. The user defines these events. • **Heal** runs the Heal event. • **Upgrade** runs the Upgrade event to ensure that the latest descriptor is used. • **Disaster Recovery** runs the Disaster Recovery event. The user defines these events. You can also execute the Vnf Diagnostics, Vnf Heal, Vnf Upgrade, and Vnf Disaster Recovery actions from the Actions portlet.
Delete	Removes the selected VNF record from the system. You can also execute the Vnf Delete action from the Actions portlet.
Shutdown	Gracefully shuts down VNF/VDUs, allowing any current processes to complete. A confirmation message is displayed. Click Shutdown to continue. The status remains the same. You can also execute the Vnf Shutdown action from the Actions portlet.
Audit Trail	Opens the Audit Trail Viewer window, which displays a list of actions that occurred for the selected VNF record. Select an audit record and the job record displays its status and any other job-related information.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection

Columns

Other than the general navigation and configuration options, the Virtual Network Function Records portlet includes the following columns (Table 17-48).

You can view the value for most of the hidden columns by right clicking a row and selecting Details. The other option is to add the column by clicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-48. Virtual Network Function Records Columns (Sheet 1 of 2)

Column	Description
UUID	A unique 128-bit, system generated value assigned to each VNF record created.
Name	A text that identifies the VNF record and its purpose.
Description	Details that describe the VNF’s purpose, version, *etc.*
Vendor	The vendor VNF (such as F5, Cisco, *etc.*) used to generate the VNF descriptor.
Customer	The person, group, or organization that requested the VNF.
Status	The current record status, such as Normal, Offline, Stopped, Instantiating, *etc.*
Active Operation	The current operation being performed on the VNF. Because VNF operations are long running, this lets you know when the VNF is ready to receive another operation. By default, this field is not shown.
Alarm Severity	The alarm severity for each record, such as Critical, Major, Minor, Informational, Normal, Warning. By default, this field is not shown.
Alarm Suppression Description	A text description of alarm suppression. By default, this field is not shown.

Table 17-48. Virtual Network Function Records Columns (Sheet 2 of 2)

Column	Description
Alarm Suppression Mode	An indicator that shows whether alarm suppression is on or off. By default, this field is not shown.
Created	The system-generated timestamp for the record creation. By default, this field is not shown.
Deployment Flavor	A usage deployment model. Valid values vary depending on the VNF selected. By default, this field is not shown.
Descriptor ID	The name of the descriptor from which the record was created. By default, shows only on the maximized view.
Foreign ID	A identifier of a VNF managed by an external VNFM. The foreign ID is useful when Cruz needs to do something to the VNF managed by the external VNFM. By default, this field is not shown.
Localization	This feature is not currently implemented. By default, this field is not shown.
Monitoring Parameters	Monitoring parameters provided by the descriptor. By default, this field is not shown.
Namespace	The name for the project that houses the VIM. This list is populated when you select a deployment VIM. By default, this field is not shown.
OSS ID	The identifier for the higher-level system for which the service was instantiated. In addition to being able to see the service associated with a particular OSS, the OSS ID is used to filter northbound notifications to the right place, making sure that the wrong OSS is not notified about an event that does not pertain to them. By default, this field is not shown.
Operation Start Time	The time the active operation started. By default, this field is not shown.
Updated	The system-generated timestamp for the latest record changes. By default, this field is not shown.
VIM	The Virtualized Infrastructure Manager on which to deploy the VNF record. By default, shows only on the maximized view.
VNF Address	The primary IP address associated with the VNF. By default, this field is not shown.
VNFM ID	The VNF manager name or identifier. By default, this field is not shown.
Version	The vendor VNF (such as F5, Cisco, *etc.*) version used to generate the VNF descriptor.

Stage VNF Window

Use the Stage VNF window (Figure 17-43) to define a new VNF to stage and then deploy.

Accessthiswindowbynavigatingto the Virtual Network Function Recordsportlet, right clicking a record, and then selecting New.

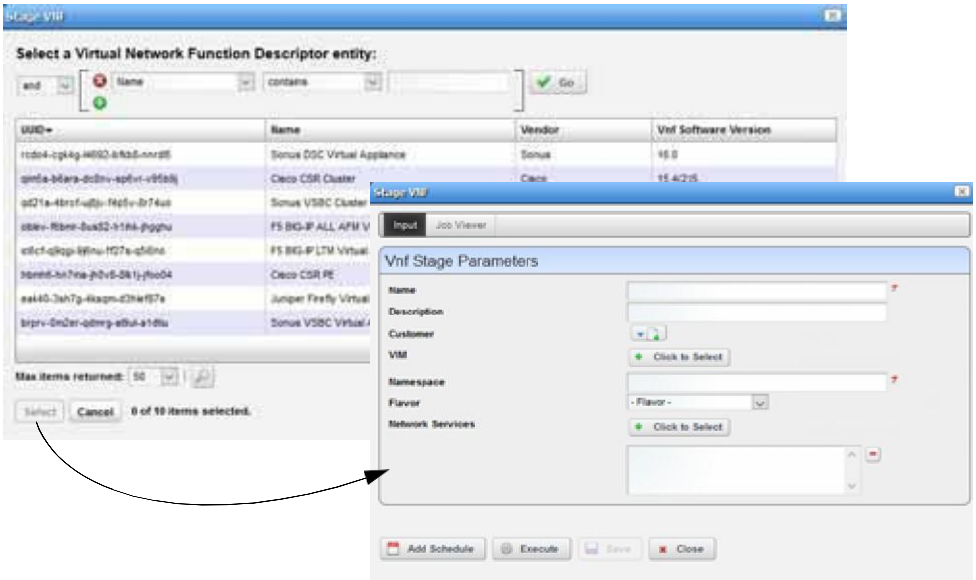

Figure 17-43. Stage VNF Window

The Stage VNF window has the following fields and option (Table 17-49) .

Table 17-49. Stage VNF Fields/Options

Field/Option	Description
Conditional Search Parameters	Define conditional search criteria to find the appropriate VNF descriptor entity.
Available VNF List	Select a VNF descriptor from which to stage a VNF record and then click Select. The Vnf Stage Parameters input fields are displayed.
Name	Enter a VNF descriptor name that identifies its purpose.
Description	Enter more details that describe the VNF's purpose, version, *etc.*
Customer	Select the person, group, or organization that requested the VNF. If the customer is not listed, create it.
VIM	Select a VIM. Selecting a VIM populates the namespace list.
Namespace	Select the project in which to deploy the VNF
Flavor	Select a deployment flavor.
Network Service	Select the network service required to stage or deploy the VNF. This field shows only when a network service is required to stage or deploy the VNF.
Add Schedule	Set the following schedule parameters used to automatically execute the VNF Stage task on the specified target. • Starting On sets the date and time to execute the task. • Recurrence sets whether the task executes Only Once, Only at Startup, every X minutes (Increment), or Every *X* minutes, hours, days, weekdays, weekend days, weeks, months, years. • Stopping On sets when the task stops executing. Valid values areNever, By Occurrence, or By Date and Time.
Execute	Runs the VNF Stage task on the specified target and reserves resources.
Save	Preserves your stage VNF configuration.
Close	Exits the Stage VNF window.

Editing VNF Record Window

Use the Editing VNF Record window (Figure 17-44) to modify the following VNF record information:

- General VNF details name, description, and monitoring parameters
- Data Values description
- Connection Points name and description
- Virtual Links name, description and QOS options
- VDUs name and description
- Data Set description

Accessthiswindowbynavigatingto the Virtual Network Function Recordsportlet, right clicking a record, and then selecting Edit.

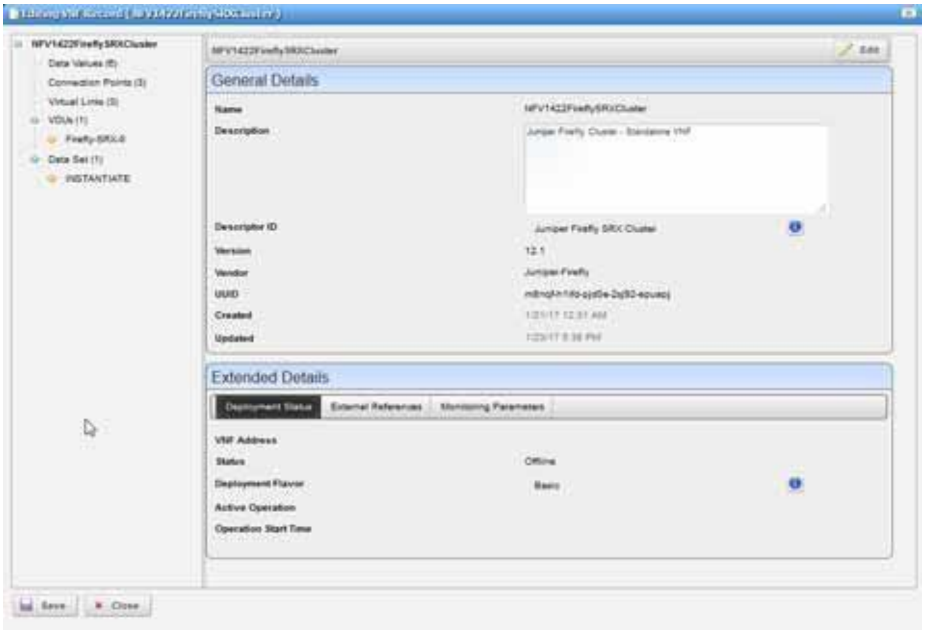

Name NFV1422FireflySRXCluster

Description Juniper Firefly Cluster - Standalone VNF

Descriptor ID Juniper Firefly SRX Cluster

Version 12.1

Vendor Juniper-Firefly

UUID m8nqf4n1i6o-ajd0e-2q93-eouaoj

Created 1/21/17 12:31 AM

Updated 1/23/17 8:38 PM

VNF Address

Status Offline

Deployment Flavor Basic

Active Operation

Operation Start Time

Figure 17-44. Editing VNF Record Window

DetailsPortlets

The Details portlets display the general details for the selected Virtualized Network Function (VNF) descriptor or record. Except where noted, the general details are the same for the VNF descriptor and VNF record.

You can also view VDU information details,descriptor deployment flavor details, connections details, alarms and event details,and history details related to the selected VNF descriptor or record.

General Details

Use the Virtualized Network Function (VNF) Genera l details (Figure 17-45) to viewidentification and version information, any monitoring parameter definitions, and a reference tree of items related to the selected VNF descriptor or record. You can also view VNF deployment status, external references, health, and connection points for the selected VNF record. For VNF descriptor connection portlets, see Connections Details on page 933.

Accessthisinformation bynavigatingto the appropriate VNFportlet, right clicking arow,and then selecting Details.

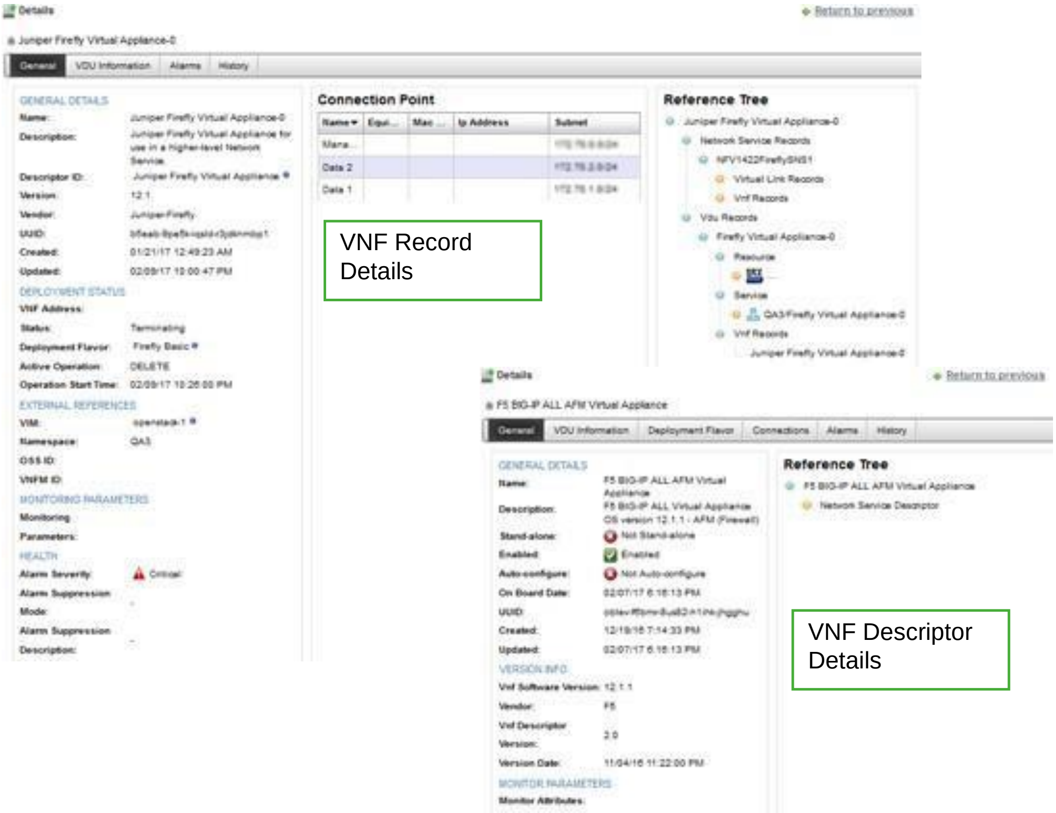

GENERAL DETAILS

Name:	Juniper Firefly Virtual Appliance-0
Description:	Juniper Firefly Virtual Appliance for use in a higher-level Network Service
Descriptor ID:	Juniper Firefly Virtual Appliance
Version:	12.1
Vendor:	JuniperFirefly
UUID:	[illegible]
Created:	01/21/17 12:49:23 AM
Updated:	02/09/17 10:00:47 PM

DEPLOYMENT STATUS

VNF Address:	
Status:	Terminating
Deployment Flavor:	Firefly Basic
Active Operation:	DELETE
Operation Start Time:	02/09/17 10:26:00 PM

EXTERNAL REFERENCES

VIM:	openstack-1
Namespace:	QA3
OSS ID:	
VNFM ID:	

MONITORING PARAMETERS

Monitoring Parameters:	

HEALTH

Alarm Severity:	Critical
Alarm Suppression Mode:	-
Alarm Suppression Description:	-

Connection Point

Name	Equi...	Mac ...	Ip Address	Subnet
Mana...				172.16.0.0/24
Data 2				172.16.2.0/24
Data 1				172.16.1.0/24

Reference Tree

- Juniper Firefly Virtual Appliance-0
  - Network Service Records
    - NFV1422FireflySNS1
      - Virtual Link Records
      - Vnf Records
  - Vdu Records
    - Firefly Virtual Appliance-0
      - Resource
      - Service
        - QA3/Firefly Virtual Appliance-0
      - Vnf Records
        - Juniper Firefly Virtual Appliance-0

GENERAL DETAILS

Name:	F5 BIG-IP ALL AFM Virtual Appliance
Description:	F5 BIG-IP ALL Virtual Appliance OS version 12.1.1 - AFM (Firewall)
Stand-alone:	Not Stand-alone
Enabled:	✓ Enabled
Auto-configure:	Not Auto-configure
On Board Date:	02/07/17 6:18:13 PM
UUID:	[illegible]
Created:	12/19/16 7:14:33 PM
Updated:	02/07/17 6:18:13 PM

VERSION INFO

Vnf Software Version:	12.1.1
Vendor:	F5
Vnf Descriptor Version:	2.0
Version Date:	11/04/16 11:22:00 PM

MONITOR PARAMETERS

Monitor Attributes:

Reference Tree

- F5 BIG-IP ALL AFM Virtual Appliance
  - Network Service Descriptor

Figure 17-45. VNF General Details Portlets

VDU Information Details

Use the Virtualized Network Function (VNF) VDU Information details (Figure 17-46) to view the following information:

- Fora Descriptor, Virtualization Deployment Unit (VDU) information (required CPU, memory, and disk) and any defined dependencies
- For a Record, VDU status and VNF virtual linkinformation

You have the option to view more details or share selected details with another user on your system.

Accessthisinformation bynavigating to the appropriate VNF portlet, right clicking arow, selecting Details, and then clicking the VDU Information tab.

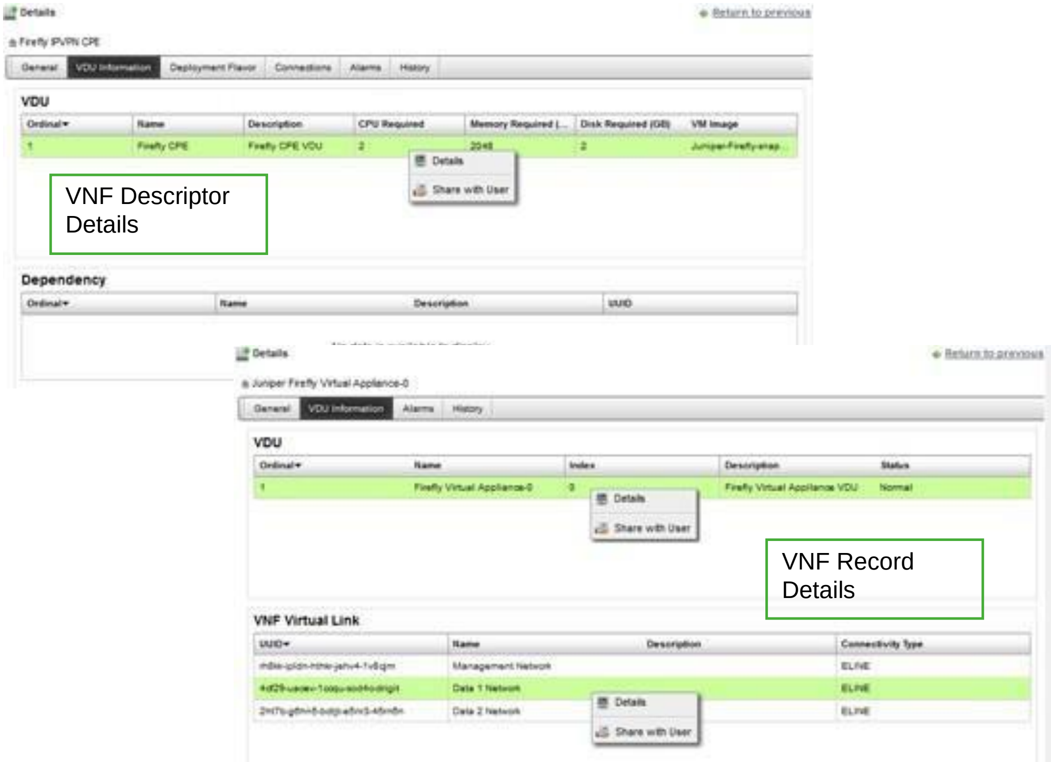

Details

Return to previous

Firefly IPVPN CPE

VDU

Ordinal▾	Name	Description	CPU Required	Memory Required (...	Disk Required (GB)	VM Image
1	Firefly CPE	Firefly CPE VDU	2	2048	2	Juniper-Firefly-snap...

VNF Descriptor Details

Dependency

Ordinal▾	Name	Description	UUID

No data is available to display

Details

Return to previous

Juniper Firefly Virtual Appliance-0

VDU

Ordinal▾	Name	Index	Description	Status
1	Firefly Virtual Appliance-0	0	Firefly Virtual Appliance VDU	Normal

VNF Record Details

VNF Virtual Link

UUID▾	Name	Description	Connectivity Type
m8ie-ipldn-htnie-jahv4-1v8qjm	Management Network		ELINE
4d29-uaoev-1ooqu-sod4odrigit	Data 1 Network		ELINE
2ni7o-g8n-6-odqi-a8vi3-46m8n	Data 2 Network		ELINE

Figure 17-46. VNF VDU Information Details Portlets

Deployment Flavor Details

Use the Virtualized Network Function (VNF) Deployment Flavor details (Figure 17-47) to viewa list of deployment flavors for the selected VNF descriptor. These deployment flavors are useful when there is a need to scale in/out the number of instances to decrease/increase VNF capacity, respectively.

Youhave the option to view more deployment flavor details or share a deployment flavor with another user on your system.

Access this information by navigating to the Virtual Network Function Descriptors portlet, right-clicking a row, selecting Details, and then clicking the Deployment Flavor tab.

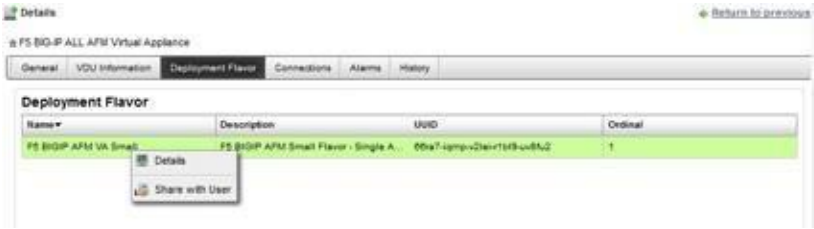

Details

Return to previous

F5 BIG-IP ALL AFM Virtual Appliance

Deployment Flavor

Name▾	Description	UUID	Ordinal
F5 BIGIP AFM VA Small	F5 BIGIP AFM Small Flavor - Single A...	66ra7-iqmp-v2iei-r1bf9-uv8fu2	1

Figure 17-47. VNF Descriptor Deployment Flavor Details Portlet

Connections Details

Usethe Virtualized Network Function (VNF) descriptor Connection details (Figure 17-48) to view more granular details about the descriptor virtual link and connection point. This information is useful when trying to troubleshoot connectivity issues.

You have the option to drill down to more details or share these details with another user.

Accessthis information bynavigating to the Virtual Network Function Descriptor portlet, right-clicking a row, selecting Details, and then clicking the Connections tab.

For VNF record connection details, see General Details on page 930. For VNF record virtual link details, see VDU Information Details on page 931.

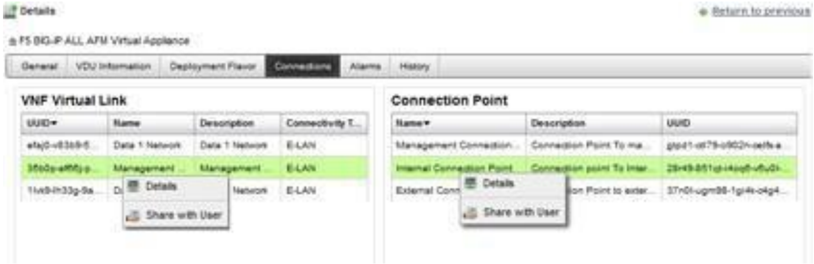

VNF Virtual Link

UUID▾	Name	Description	Connectivity T...
efaj0-v83b9-5...	Data 1 Network	Data 1 Network	E-LAN
35b0p-aff6j-p...	Management ...	Management ...	E-LAN
1lv9-lh33g-9a...	D	Network	E-LAN

Connection Point

Name▾	Description	UUID
Management Connection...	Connection Point To ma...	gtpd1-ol79-o902n-oefe-a...
Internal Connection Point	Connection point To Inter...	29r49-851qt-i4oq6-v6u0l...
External Conn	on Point to exter...	37n0l-ugm98-1gl4i-o4g4...

Figure 17-48. VNF Descriptor Connections Portlets

Alarms and Event Details

Usethe Virtualized Network Function(VNF) Alarmsdetails(Figure 17-49) to monitorand manage alarms and view event history for the selected device.

You can perform the same actions on an alarm or event history from this location as you can from the Alarms and Event History portlets, respectively. See Alarms on page 276 for details about the Alarms and Event History portlets.

Accessthisinformation bynavigating to the appropriate VNF portlet, right clicking arow, selecting Details, and then clicking the Alarms tab.

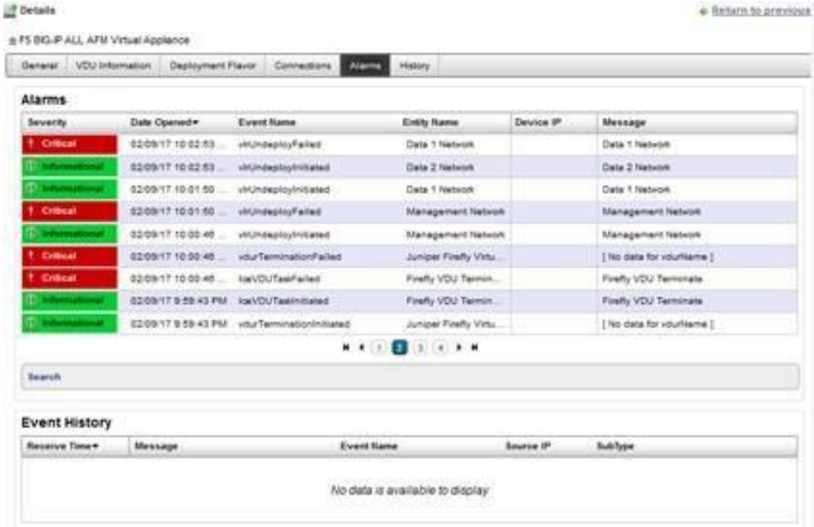

Alarms

Severity	Date Opened▾	Event Name	Entity Name	Device IP	Message
Critical	02/09/17 10:02:53 ...	vlrUndeployFailed	Data 1 Network		Data 1 Network
Informational	02/09/17 10:02:53 ...	vlrUndeployInitiated	Data 2 Network		Data 2 Network
Informational	02/09/17 10:01:50 ...	vlrUndeployInitiated	Data 1 Network		Data 1 Network
Critical	02/09/17 10:01:50 ...	vlrUndeployFailed	Management Network		Management Network
Informational	02/09/17 10:00:46 ...	vlrUndeployInitiated	Management Network		Management Network
Critical	02/09/17 10:00:46 ...	vdurTerminationFailed	Juniper Firefly Virtu...		[No data for vdurName]
Critical	02/09/17 10:00:46 ...	ksVDUTaskFailed	Firefly VDU Termin...		Firefly VDU Terminate
Informational	02/09/17 9:59:43 PM	ksVDUTaskInitiated	Firefly VDU Termin...		Firefly VDU Terminate
Informational	02/09/17 9:59:43 PM	vdurTerminationInitiated	Juniper Firefly Virtu...		[No data for vdurName]

Event History

Receive Time▾	Message	Event Name	Source IP	SubType

No data is available to display

Figure 17-49. VNF Alarms and Events Details Portlets

History Details

Use the Virtualized Network Function (VNF) History details (Figure 17-50) to viewalist ofactions (audit trail) performed on the selected VNF descriptor or record.

For a selected action, you can view the job, aging policy, delete the action, view as a portable document format (PDF), or share this information with other users on your system.

Accessthisinformationbynavigating to theappropriate VNFportlet, right clicking arow,selecting Details, and then clicking the History tab.

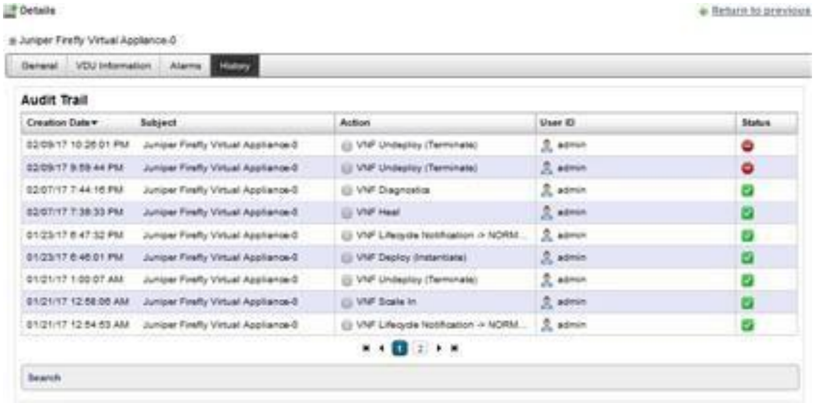

Audit Trail

Creation Date▾	Subject	Action	User ID	Status
02/09/17 10:26:01 PM	Juniper Firefly Virtual Appliance-0	VNF Undeploy (Terminate)	admin	⛔
02/09/17 9:59:44 PM	Juniper Firefly Virtual Appliance-0	VNF Undeploy (Terminate)	admin	⛔
02/07/17 7:44:16 PM	Juniper Firefly Virtual Appliance-0	VNF Diagnostics	admin	✓
02/07/17 7:38:33 PM	Juniper Firefly Virtual Appliance-0	VNF Heal	admin	✓
01/23/17 6:47:32 PM	Juniper Firefly Virtual Appliance-0	VNF Lifecycle Notification -> NORM...	admin	✓
01/23/17 6:46:01 PM	Juniper Firefly Virtual Appliance-0	VNF Deploy (Instantiate)	admin	✓
01/21/17 1:00:07 AM	Juniper Firefly Virtual Appliance-0	VNF Undeploy (Terminate)	admin	✓
01/21/17 12:58:06 AM	Juniper Firefly Virtual Appliance-0	VNF Scale In	admin	✓
01/21/17 12:54:53 AM	Juniper Firefly Virtual Appliance-0	VNF Lifecycle Notification -> NORM...	admin	✓

Figure 17-50. VNF History Details

Virtual Requirements Portlet

Use the Virtual Requirements portlet (Figure 17-51) to view the virtual domain resource requirements and usage (such as memory, CPU, and disk). Youalso have the option to modify a domain’s description.

The Virtual Requirements portlet alsoprovidesamaximized viewthatincludesadditional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

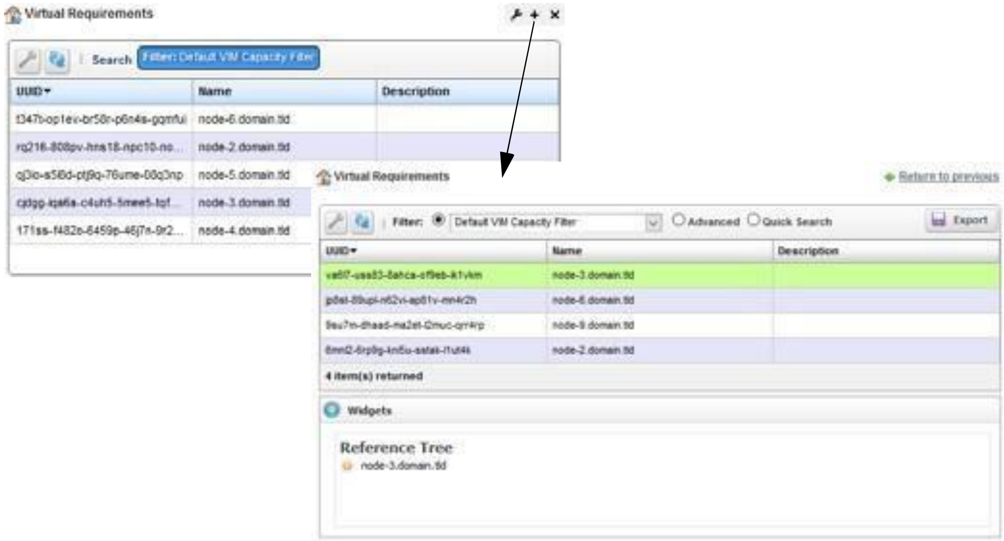
Virtual Requirements
Search
Filter: Default VIM Capacity Filter

UUID	Name	Description
f347b-op1ev-br58r-p6n4s-gqmfui	node-6.domain.tld	
rq218-808pv-hns18-npc10-no...	node-2.domain.tld	
q3io-s58d-ptj9q-76ume-08q3np	node-5.domain.tld	
cjdgg-iqa6a-c4uh5-5mee5-tqf...	node-3.domain.tld	
171ss-f482o-6459p-46j7n-9r2...	node-4.domain.tld	

Virtual Requirements
Return to previous
Filter: ● Default VM Capacity Filter ○ Advanced ○ Quick Search
Export

UUID	Name	Description
va6l7-usa83-8ahca-of9eb-k1vkm	node-3.domain.tld	
p8al-89upl-n62vi-ap81v-mn4r2h	node-6.domain.tld	
9su7n-dhaad-ma2el-l2muc-qrr4rp	node-9.domain.tld	
6mrl2-6rp9g-kn6u-aatak-l1ut4s	node-2.domain.tld	

4 item(s) returned

Figure 17-51. Virtual Requirements Portlet

Pop-Up Menu

The Virtual Requirements portlet providesaccess to the followingpop-up menu options. Right-click a row to access these options (Table 17-50).

Table 17-50. Virtual Requirements Pop-Up Menu Options

Option	Description
Edit	Opens the Edit VIM Capacity window, where you view virtual requirements information and make any needed modifications to the description.
Details	Opens the Details portlets, where you view the general details, VDU reservations, reference tree, and activities history (audit-trail). You can also modify a selected VDU’s name or description, view VDU details, or share a VDU with another user on your system.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Right clicking Note: Sharing only handles one item so it uses the first one in the selection.**

Columns

Other than the general navigationand configuration options, the Virtual Requirements portlet includes the following columns (Table 17-51).

Youcan viewthe value for most of the hidden columns byright clicking a rowand selecting Details. The other option is to add the column by clicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-51. Virtual Requirements Columns (Sheet 1 of 2)

Column	Description
UUID	The unique 128-bit, system generated value assigned to each virtual requirements record created. By default, this field is not shown.

Table 17-51. Virtual Requirements Columns (Sheet 2 of 2)

Column	Description
Name	The text that identifies the VIM domain.
Description	The hypervisor description.
CPU Free	The number of CPUs available. By default, this field is not shown.
CPU Reserved	The number of CPUs reserved. By default, this field is not shown.
CPU Unallocated	The number of CPUs not allocated. By default, this field is not shown.
CPU Used	The number of CPUs used. By default, this field is not shown.
Cloud Id	The text that identifies the VIM. By default, this field is not shown.
Created	The system-generated timestamp for the virtual requirements creation. By default, this field is not shown.
Disk Free	The disk space available in gigabytes. By default, this field is not shown.
Disk Reserved	The disk space reserved in gigabytes. By default, this field is not shown.
Disk Unallocated	The disk space not allocated in gigabytes. By default, this field is not shown.
Disk Used	The disk space used in gigabytes. By default, this field is not shown.
Host Id	The OpenStack identifier. By default, this field is not shown.
IP Address	The address where the VIM is located. By default, this field is not shown.
Last Synchronization	The most recent VIM synchronization timestamp. By default, this field is not shown.
Memory Free	The memory available in megabyte s. By default, this field is not shown.
Memory Reserved	The memory reserved in megabytes. By default, this field is not shown.
Memory Unallocated	The memory not allocated in megabytes. By default, this field is not shown.
Memory Used	The memory used in megabytes. By default, this field is not shown.
Total CPU	The total number of CPUs (used, free, reserved, and unallocated). By default, this field is not shown.
Total Disk	The total gigabytes of disk space (used, free, reserved, and unallocated). By default, this field is not shown.
Total Memory	The total megabytes of memory (used, free, reserved, and unallocated). By default, this field is not shown.
Type	Information from the OpenStack system describing the type of hypervisor. By default, this field is not shown.
Updated	The system-generated timestamp for the latest virtual requirements changes. By default, this field is not shown.

Virtual Reservations Portlet

Use the Virtual Reservations portlet (Figure 17-52) to verifythat a VDU's resourceswerereserved after you stage a VNF record.

From the Virtual Reservation portlet, asystemadministrator maintains the reservations, a network designer implements reservations, and operators (if they the proper permissions) view to verify resources were reserved after VNF staging.

The Virtual Reservations portlet alsoprovidesamaximized viewthat includes additional filtering and the ability to export the table to a PDF, Excel, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

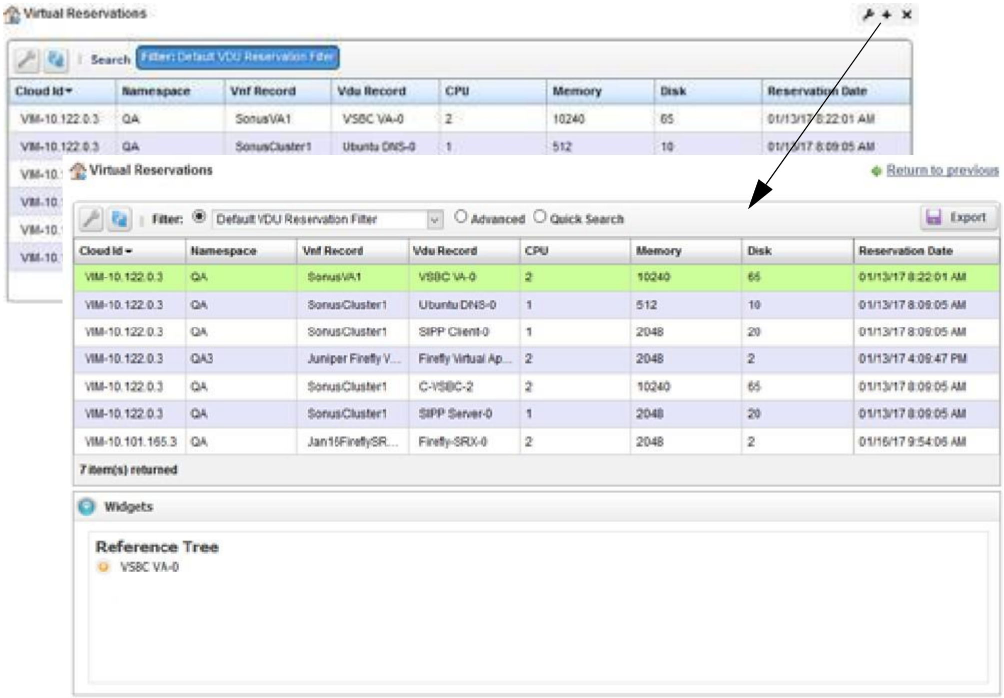

Cloud Id	Namespace	Vnf Record	Vdu Record	CPU	Memory	Disk	Reservation Date
VIM-10.122.0.3	QA	SonusVA1	VSBC VA-0	2	10240	65	01/13/17 8:22:01 AM
VIM-10.122.0.3	QA	SonusCluster1	Ubuntu DNS-0	1	512	10	01/13/17 8:09:05 AM
VIM-10.							
VIM-10.							
VIM-10.							
VIM-10.							

Cloud Id	Namespace	Vnf Record	Vdu Record	CPU	Memory	Disk	Reservation Date
VIM-10.122.0.3	QA	SonusVA1	VSBC VA-0	2	10240	65	01/13/17 8:22:01 AM
VIM-10.122.0.3	QA	SonusCluster1	Ubuntu DNS-0	1	512	10	01/13/17 8:09:05 AM
VIM-10.122.0.3	QA	SonusCluster1	SIPP Client-0	1	2048	20	01/13/17 8:09:05 AM
VIM-10.122.0.3	QA3	Juniper Firefly V...	Firefly Virtual Ap...	2	2048	2	01/13/17 4:09:47 PM
VIM-10.122.0.3	QA	SonusCluster1	C-VSBC-2	2	10240	65	01/13/17 8:09:05 AM
VIM-10.122.0.3	QA	SonusCluster1	SIPP Server-0	1	2048	20	01/13/17 8:09:05 AM
VIM-10.101.165.3	QA	Jan15FireflySR...	Firefly-SRX-0	2	2048	2	01/16/17 9:54:06 AM

7 item(s) returned

Figure 17-52. Virtual Reservations Portlet

Pop-Up Menu

The Virtual Reservations portlet providesaccess to the following pop-up menu options. Right-click a row to access these options (Table 17-52).

Table 17-52. Virtual Reservations Pop-Up Menu Options

Option	Description
Edit	Opensthe Editing VDU Reservation window, where youcan change the VDU name and description for the selected reservation.
Details	Displays the general VDU reservation details, such as identification information, CPU, memory, disk, date, *etc*. for the selected reservation.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note:** Sharing only handles one item so it uses the first one in the selection.

Columns

Other than the general navigation and configuration options, the Virtual Reservations portlet includes the following columns (Table 17-53).

You can view the value for most of the hidden columns by right clicking a row and selecting Details. The other option is to add the column byclicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-53. Virtual Reservations Columns

Column	Description
Cloud Id	The Virtualized Infrastructure Manager (VIM) name.
Namespace	The name for the project that houses the VIM.
Vnf Record	The VNF record name.
Vdu Record	The VDU record name.
CPU	The number of CPUs reserved for the selected VDU.
Memory	Displays how much memory is reserved for the selected VDU.
Disk	Displays how much disk space is reserved for the selected VDU.
Reservation Date	The date the resources were reserved.
Created	The timestamp when the reservation record was created. By default, this field is not shown.
Description	Details about the reserved VDU, such as function, version, purpose, *etc*. By default, this field is not shown.
Name	The text that identifies the reserved VDU. By default, this field is not shown.
UUID	The unique 128-bit, system generated value assigned to each virtual reservation record created. By default, this field is not shown.
Updated	The timestamp for the latest reservation record changes. By default, this field is not shown.

Virtualized Infrastructure Managers Portlet

Use the Virtualized Infrastructure Managers portlet (Figure 17-53) to maintain a list of virtualized infrastructure managers(VIMs) to which you deploynetwork services, physical network functions (PNFs) and virtualized network functions (VNFs).

From the Virtualized Infrastructure Managers portlet, a system administrator creates and maintains VIMs and operators (if they have the proper permissions) view VIM resources as part of their network monitoring process.

The Virtualized Infrastructure Managers portlet also provides a maximized view that includes additional filtering and the ability to export the table to a portable document format (PDF), Excel format, or CSV format.

Both views have the same pop-up menu options and by default the same columns. See Expanded Portlet on page 130 for a description of its standard options provided.

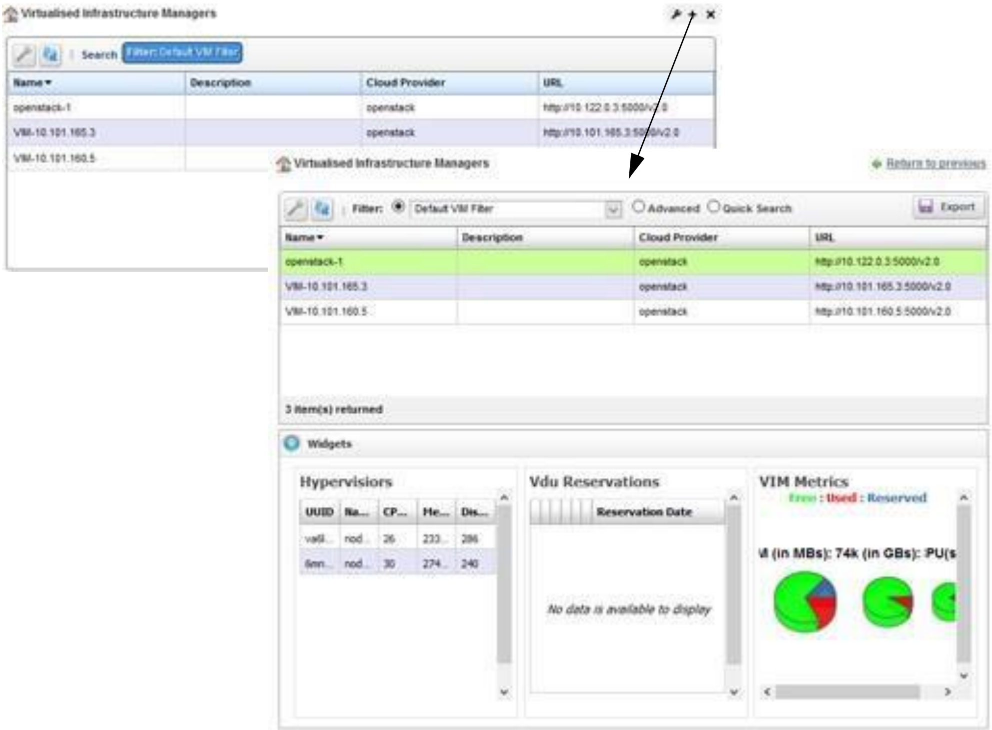

Virtualised Infrastructure Managers

Search Filter: Default VM Filter

Virtualised Infrastructure Managers

Filter: ● Default VM Filter ○ Advanced ○ Quick Search

Name	Description	Cloud Provider	URL
openstack-1		openstack	http://10.122.0.3:5000/v2.0
VIM-10.101.165.3		openstack	http://10.101.165.3:5000/v2.0
VIM-10.101.160.5		openstack	http://10.101.160.5:5000/v2.0

3 item(s) returned

Widgets

Hypervisiors

UUID	Na...	CP...	Me...	Dis...
va8...	nod...	26	233...	286
6mn...	nod...	30	274...	240

Vdu Reservations

Reservation Date

No data is available to display

VIM Metrics

Free : Used : Reserved

M (in MBs): 74k (in GBs): PU(s

Figure 17-53. Virtualized Infrastructure Managers Portlet

Pop-Up Menu

The Virtualized Infrastructure Managers portlet providesaccess to the followingpop-up menu options. Right-click a row to access these options (Table 17-54).

Table 17-54. Virtualized Infrastructure Managers Pop-Up Menu Options (Sheet 1 of 2)

Option	Description
New	Opens the Editing VIM window, where you create a VIM.
Edit	Opens the Editing VIM window, where you modify the selected VIM's description.
Details	Displays the general details, such as identification information, version information, deployment status, external references monitor parameters, health, connection points, and a reference tree. You can also view alarms and event details and history details selecting the appropriate tab.
Resync	Runs the VIM Resync task on the target VIM to update capacity information. You can also execute the VIM Resync action from the Actions portlet.
Topology	Opens the Topology portlet, where you define a multi-layered, customizable topology view of your network to help track network devices state. See Topology Portlet on page 235 for the Topology portlet configuration details.
Delete	Removes the selected VIM from the system. You can also execute the VIM Delete action from the Actions portlet.

Table 17-54. Virtualized Infrastructure Managers Pop-Up Menu Options (Sheet 2 of 2)

Option	Description
Import/Export	Provides the following actions when available for the selected image: • Import retrieves a file containing XML VIM descriptions. Some imports can come from a URL. • Export Selection exports the selected VIM description to an XML file. • Export All exports all VIM descriptions to an XML file. Click Download Export File to specify where to save the file. The Import/Export option is useful as a backup or to share descriptors with other projects.
Share with User	Opens the Share with User window, where you select a user to which you want to share the selected asset and then enter a message. You can share with colleagues existing on your system. **Note: Sharing only handles one item so it uses the first one in the selection.**

Columns

Other than the general navigation and configuration options, the Virtualized Infrastructure Managers portlet includes the following columns (Table 17-55).

You can view the value for most of the hidden columns by right clicking a row and selecting Details. The other option is to add the column byclicking the Settings tool, selecting the Columns tab, clicking Show for the appropriate column, and the applying the change.

Table 17-55. Virtualized Infrastructure Managers Columns (Sheet 1 of 2)

Column	Description
Name	The text that identifies the VIM.
Description	The optional details provided about the VIM, such as its purpose.
Cloud Provided	The cloud environment where the VIM resides.
URL	The address used to access the VIM. For example: `http://`*`ipAddress`*`:5000/v2.0`
Alarm Severity	The severity of alarms against the VIM. By default, this field is not shown.
Alarm Suppression Mode	An indicator that shows whether alarms are suppressed. By default, this field is not shown.
CPU Free	The number of CPUs available. By default, this field is not shown.
CPU Reserved	The number of CPUs reserved. By default, this field is not shown.
CPU Unallocated	The number of CPUs not allocated. By default, this field is not shown.
CPU Used	The number of CPUs used. By default, this field is not shown.
Created	The VIM creation timestamp. By default, this field is not shown.
Disk Free	The disk space available in gigabytes. By default, this field is not shown.
Disk Reserved	The disk space reserved in gigabytes. By default, this field is not shown.
Disk Unallocated	The disk space not allocated in gigabytes. By default, this field is not shown.
Disk Used	The disk space used in gigabytes. By default, this field is not shown.
Last Synchronization	The most recent VIM synchronization timestamp. By default, this field is not shown.
Memory Free	The memory available in megabyte s. By default, this field is not shown.
Memory Reserved	The memory reserved in megabytes. By default, this field is not shown.

Table 17-55. Virtualized Infrastructure Managers Columns (Sheet 2 of 2)

Column	Description
Memory Unallocated	The memory not allocated in megabytes. By default, this field is not shown.
Memory Used	The memory used in megabytes. By default, this field is not shown.
Priority	A relative number indicating which VIM should be chosen over another, where multiple VIMs support a VNF deployment. By default, this field is not shown.
Private Network Space	The private network space identifier in the following format: `123.17.0.0/16` This is the overall address space used to create smaller subnet private networks. By default, this field is not shown.
Shared Network Space	The shared network space identifier in the following format: `123.18.0.0/16` This is the overall address space used to create smaller subnet shared networks. By default, this field is not shown.
Total CPU	The total number of CPUs (used, free, reserved, and unallocated). By default, this field is not shown.
Total Disk	The total gigabytes of disk space (used, free, reserved, and unallocated). By default, this field is not shown.
Total Memory	The total megabytes of memory (u sed, free, reserved, and unallocated). By default, this field is not shown.
UUID	A unique 128-bit, system generated value assigned to each VIM created. By default, this field is not shown.
Updated	The latest VIM changes timestamp. By default, this field is not shown.

Editing VIM Window

Use the Editing VIMwindow (Figure 17-54) to create and maintain Virtualized Infrastructure Managers (VIMs).

Accessthiswindowbynavigatingto the Virtualized Infrastructure Managersportlet, right clicking a VIM or the portlet, and then selecting Edit or New. You can also access this window from the details reference tree, by right clicking a VIM, and then selecting Edit.

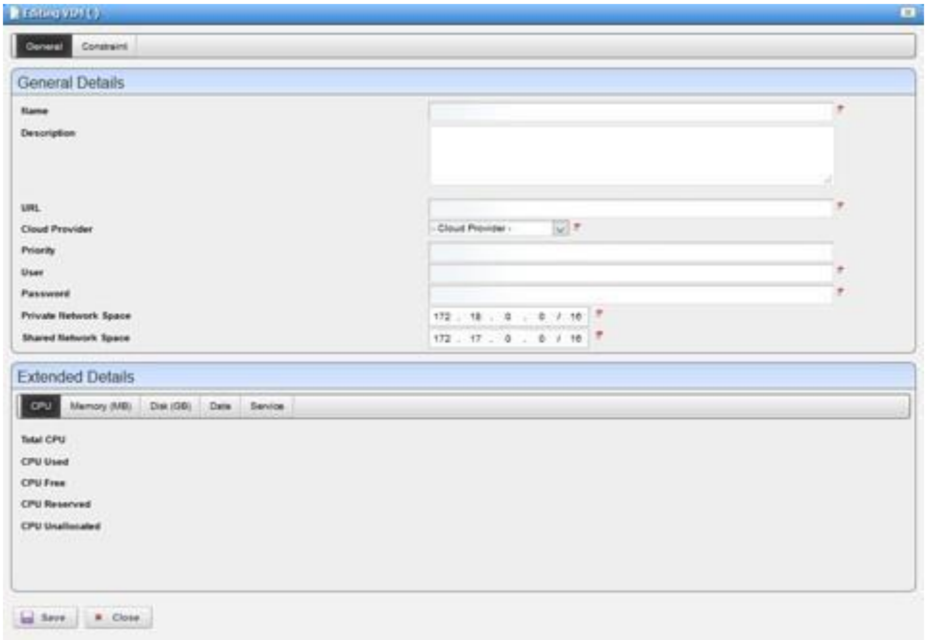

Private Network Space 172 . 18 . 0 . 0 / 16

Shared Network Space 172 . 17 . 0 . 0 / 16

Figure 17-54. Editing VIM Window

The Editing VIM window has the following fields and option (Table 17-56).

Table 17-56. Editing VIM Fields/Options (Sheet 1 of 2)

Field/Option	Description
Name	Enter required text that identifies the VIM.
Description	Enter an optional description that provides more information about the VIM, such as its purpose.
URL	Enter the address used to access the VIM. For example: `https://`*`ipAddress`*`:5000/v2.0`
Cloud Provider	Select the provider where the VIM resides.
Priority	Enter a relative number indicating which VIM should be chosen over another, where multiple VIMs support a VNF deployment. By default, this field is not shown.
User	Enter the user account ID that is used to access the VIM (OpenStack).
Password	Enter the password assigned the provided user ID.
Private Network Space	Modify appropriately. This field is automatically populated in the following format when you select the New menu option: `123.17.0.0/16` This is the overall address space used to create smaller subnet private networks.
Shared Network Space	Modify appropriately. This field is automatically populated in the following format when you select the New menu option: 123.18.0.0/16 This is the overall address space used to create smaller subnet shared networks. By default, this field is not shown.
Extended Details	Displays details about the VIM, such as CPU, memory, and disk information as well as the synchronized, created, and updated dates, and service UUID.

Table 17-56. Editing VIM Fields/Options (Sheet 2 of 2)

Field/Option	Description
Constraints	Displays the list of defined constraints, which determine where to place network services or VNFs if the VIM is not specified during network service or VNF descriptor staging process. Modify the list as needed by adding, deleting, or modifying constraints.

Details Portlets

The Details portlets display the general details for the selected Virtualized Infrastructure Manager (VIM).

You can also view alarms and event details and history details related to the selected VIM.

Accessthe detailsportlets bynavigating to the Virtualized Infrastructure Managers portlet, right-clicking a VIM, and then selecting Details.

General Details

Use the Virtualized Infrastructure Manager (VIM) General details (Figure 17-55) to view identification and version information, CPU, memory (MB), disk (GMB), activity dates, service information, health, constraints, hypervisors, and a reference tree of items related to the selected VIM.

From the Hypervisor details, you have the following menu options:

- **Edit** to view VIM capacity or change the VIM' s description from the Editing VIM Capacity window
- **Details** allows you to view more details about a selected hypervisor and any audit trail history
- **Share with User** allows you to share the selected asset with colleagueson your system

From the Vim Images details, you have the following menu options:

- **Details** allows you to view details and the reference tree for the selected image
- **Share with User** allows you to share the selected asset with other users on your system

From the Reference Tree, you have the following menu options:

- **Edit** whenyouright-clickacapacityelement, opensthe Editing VIMCapacitywindow, where you view VIM capacity or change the VIM'sdescription
  Edit when you right-click a VIM, opens the Editing VIM window, where you modify the selected VIM's description
- **Details** allows you to view more details about a selected hypervisor and any audit trail history
- **Share with User** allows you to share the selected asset with colleagues on your system

Access this information by navigating to the Virtualized Infrastructure Managers portlet, right-clicking a row, and then selecting Details.

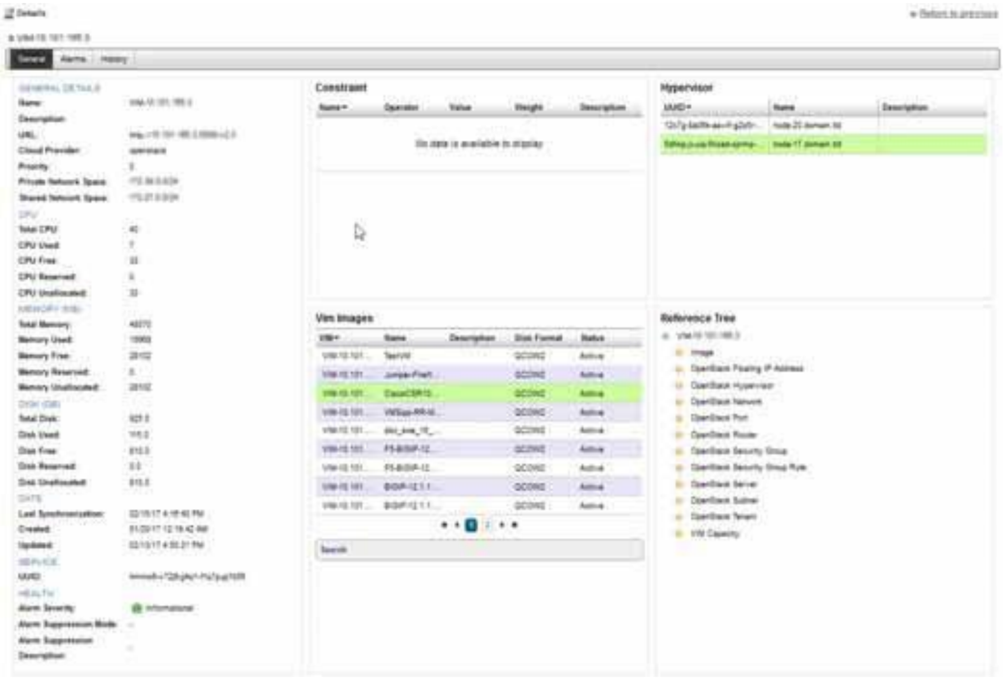

Details

Return to previous

VIM-10.101.165.3

GENERAL DETAILS

Name:	VIM-10.101.165.3
Description:	
URL:	http://10.101.165.3:5000/v2.0
Cloud Provider:	openstack
Priority:	0
Private Network Space:	172.30.0.0/24
Shared Network Space:	172.21.0.0/24
CPU	
Total CPU:	40
CPU Used:	7
CPU Free:	33
CPU Reserved:	0
CPU Unallocated:	33
MEMORY (MB)	
Total Memory:	48370
Memory Used:	15950
Memory Free:	28102
Memory Reserved:	0
Memory Unallocated:	28102
DISK (GB)	
Total Disk:	927.0
Disk Used:	115.0
Disk Free:	812.0
Disk Reserved:	0.0
Disk Unallocated:	812.0
DATE	
Last Synchronization:	02/15/17 4:18:40 PM
Created:	01/20/17 12:19:42 AM
Updated:	02/13/17 4:50:21 PM
SERVICE	
UUID:	[illegible]
HEALTH	
Alarm Severity:	Informational
Alarm Suppression Mode:	--
Alarm Suppression Description:	--

Constraint

Name	Operator	Value	Weight	Description
No data is available to display				

Hypervisor

UUID	Name	Description
[illegible]	node-20.domain.tld	
[illegible]	node-17.domain.tld	

Vim Images

VIM	Name	Description	Disk Format	Status
VIM-10.101...	TestVM		QCOW2	Active
VIM-10.101...	Juniper-FireF...		QCOW2	Active
VIM-10.101...	CiscoCSR10...		QCOW2	Active
VIM-10.101...	VMSup-RR-M...		QCOW2	Active
VIM-10.101...	[illegible]		QCOW2	Active
VIM-10.101...	F5-BIGIP-12...		QCOW2	Active
VIM-10.101...	F5-BIGIP-12...		QCOW2	Active
VIM-10.101...	BIGIP-12.1.1...		QCOW2	Active
VIM-10.101...	BIGIP-12.1.1...		QCOW2	Active

Reference Tree

- VIM-10.101.165.3
  - Image
  - OpenStack Floating IP Address
  - OpenStack Hypervisor
  - OpenStack Network
  - OpenStack Port
  - OpenStack Router
  - OpenStack Security Group
  - OpenStack Security Group Rule
  - OpenStack Server
  - OpenStack Subnet
  - OpenStack Tenant
  - VIM Capacity

Figure 17-55. VIM General Details Portlets

AlarmsandEventDetails

Use the Virtualized Infrastructure Manager (VIM) Alarms details (Figure 17-56) to monitor and manage alarms and view event history for the selected VIM.

You can perform the same actions on an alarm or event history from this location as you can from the Alarms and Event History portlets, respectively. See Alarms on page 276 for details about the Alarms and Event History portlets.

Access this information by navigating to the Virtualized Infrastructure Managers portlet, right-clicking a row, selecting Details, and then clicking Alarms.

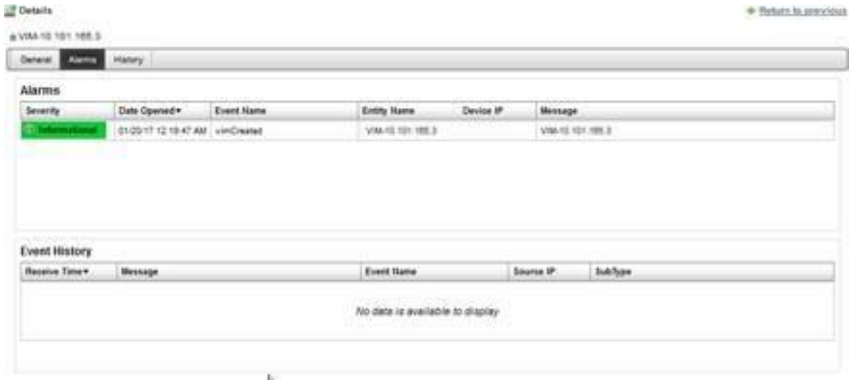

Details

Return to previous

VIM-10.101.165.3

Alarms

Severity	Date Opened	Event Name	Entity Name	Device IP	Message
Informational	01/20/17 12:19:47 AM	vimCreated	VIM-10.101.165.3		VIM-10.101.165.3

Event History

Receive Time	Message	Event Name	Source IP	SubType
No data is available to display				

Figure 17-56. Alarms and Event History Portlets

History Details

Use the Virtualized Infrastructure Manager (VIM) History details (Figure 17-57) to viewa list of actions (audit trail) performed on the selected VIM.

For a selected action, you can view the job, aging policy, delete the action, view as a portable document format (PDF), or share this information with other users on your system.

Access this information by navigating to the Virtualized Infrastructure Managers portlet, right-clicking a row, selecting Details, and then clicking History.

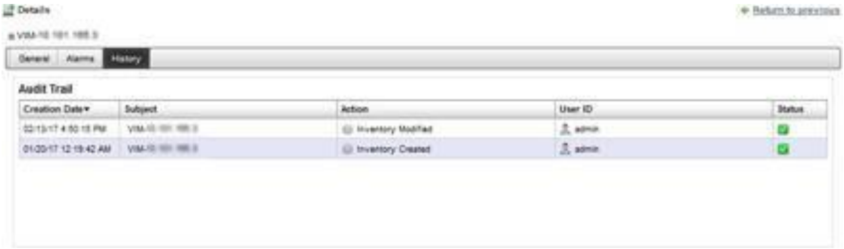

Audit Trail

Creation Date▾	Subject	Action	User ID	Status
02/13/17 4:50:15 PM	VIM-[redacted]	Inventory Modified	admin	✓
01/20/17 12:19:42 AM	VIM-[redacted]	Inventory Created	admin	✓

Figure 17-57. VIM History Details Portlet

Managing Virtualized Network Functions

From the application, you can manage virtualized network functions (VNFs) by staging and deploying VNF resources (instantiating) and then managing those resources by monitoring network health, updating VNFs, and adding or removing VNFs as needed to resolve issues or expansion.

See Network Virtualization Portlets on page 8 66 for a detailed description of each portlet an d the different action s that you can perform.

- Instantiating a Complex Virtual Network
- Instantiating a Standalone VNF Record
- Monitoring Network Health
- Updating Virtual Network Functions
- Undeploying Virtual Network Functions

Instantiating a Complex Virtual Network

Thissection illustrates how to instantiate acomplex Virtualized Network Function (VNF) record. The example environment has a Session Boarder Controller, a client Session Initiation Protocol (SIP) server, and a DNS with internal networks that connects the virtual networks.

Instantiate a virtual network by:

- Reviewing a System’s Current State
- Reviewing a VNF Descriptor
- Staging a Complex VNF Record
- Verifying Reserved Resources
- Deploying a Complex VNF Record
- Viewing Alarms and Notifications
- Verifying a Complex VNF Instantiation
- Reviewing the VNF Record Status

Reviewing a System’s Current State

The Virtualized Infrastructure Managers (VIMs) portlet liststhe OpenStack controllers and VIM instances. Before you instantiate a Virtualized Network Function (VNF), it is good to k now the system’s current state and the OpenStack controllersand VIM instances available to which you can deploy services.

The Cruz application provides many management tools from which to access the network health and configuration, such as:

- Topology
- Resource Management
- Performance Monitoring
- Infrastructure Capacity

Reviewthe system’s current state from the Cruz Virtualized Infrastructure Managers portlet as follows.

NOTE:

TheVirtualized Infrastructure Managersportletlocationisdependentonyourcompany’s CruzNFV installation and configuration.

1 Select a VIM or OpenStack controller.

2 Right-click > Details.

Theselected instance’s details display CPU, memory, and diskcapacity, serviceand health status, and lists any constraints and hypervisor records.

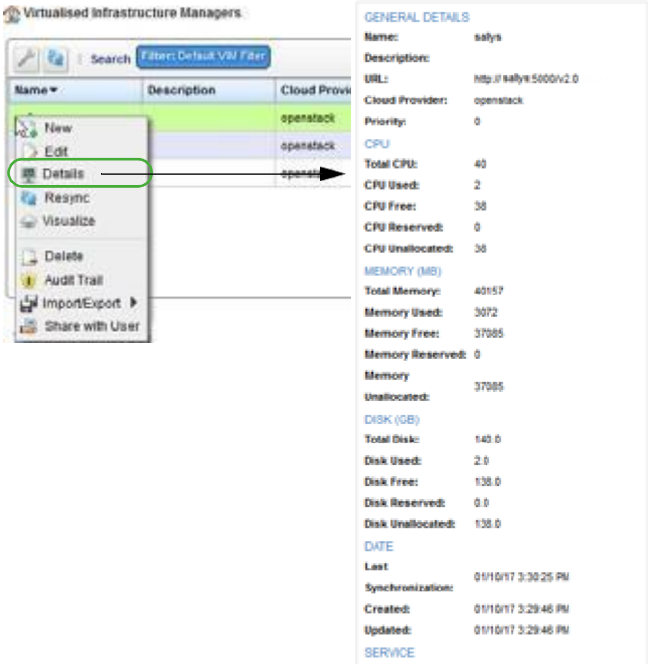

GENERAL DETAILS	
Name:	sallys
Description:	
URL:	http://sallys:5000/v2.0
Cloud Provider:	openstack
Priority:	0
CPU	
Total CPU:	40
CPU Used:	2
CPU Free:	38
CPU Reserved:	0
CPU Unallocated:	38
MEMORY (MB)	
Total Memory:	40157
Memory Used:	3072
Memory Free:	37085
Memory Reserved:	0
Memory Unallocated:	37085
DISK (GB)	
Total Disk:	140.0
Disk Used:	2.0
Disk Free:	138.0
Disk Reserved:	0.0
Disk Unallocated:	138.0
DATE	
Last Synchronization:	01/10/17 3:30:25 PM
Created:	01/10/17 3:29:46 PM
Updated:	01/10/17 3:29:46 PM
SERVICE	

3 Select a hypervisor server.

4 Right-click > Details.

Theselected hypervisor details aredisplayed, such as CPU, memory, and diskcapacityand any related VDU reservations.

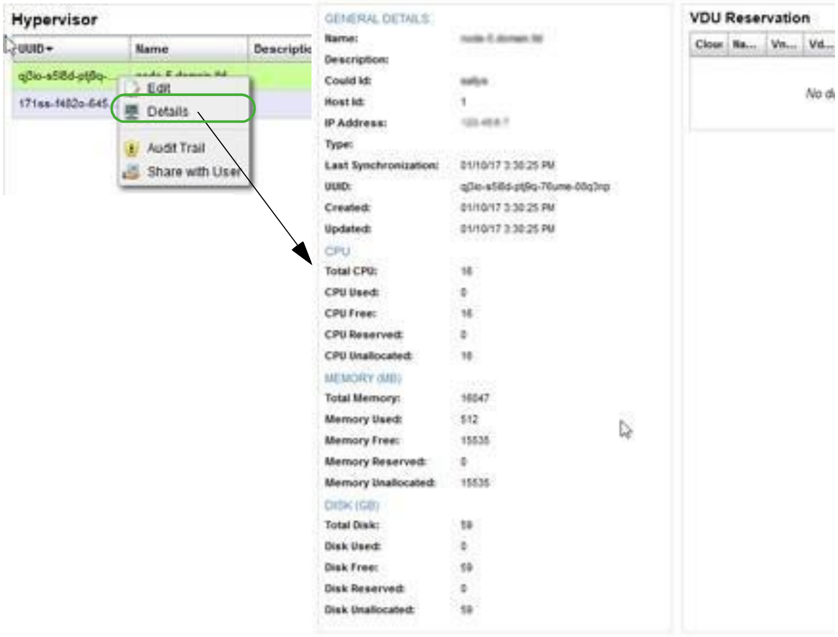

GENERAL DETAILS

Name: [redacted]

Description:

Could Id: [redacted]

Host Id: 1

IP Address: [redacted]

Type:

Last Synchronization: 01/10/17 3:30:25 PM

UUID: q3io-s58d-ptj9q-70ume-00q3np

Created: 01/10/17 3:30:25 PM

Updated: 01/10/17 3:30:25 PM

CPU

Total CPU: 16

CPU Used: 0

CPU Free: 16

CPU Reserved: 0

CPU Unallocated: 16

MEMORY (MB)

Total Memory: 16047

Memory Used: 512

Memory Free: 15535

Memory Reserved: 0

Memory Unallocated: 15535

DISK (GB)

Total Disk: 59

Disk Used: 0

Disk Free: 59

Disk Reserved: 0

Disk Unallocated: 59

5 Go to the OpenStack dashboard to view information about the server to which you are going to deploy, such as the network topology, networks, instances, and so on.

6 Return to the Cruz application.

Next review the VNF descriptor you plan to deploy to gain an understanding of its functionality.

Reviewing a VNF Descriptor

The Cruz Virtual Network Function (VNF) Descriptors portlet provides a list of the VNF descriptors available. Before using a descriptor to deploy a VNF, review the descriptor to gain a better understanding of its functionality.

Review a VNF descriptor from the Virtual Network Function Descriptors portlet as follows.

1 Select a VNF descriptor.

2 Right-click > Edit.

The Editing Virtual Network Function Descriptor window is displayed.

3 Review the descriptor details, such as VDU components used to connect internal networks to the VNF, lifecycle events (terminate/instantiate), deployment flavors (models), and so on. It needs to be complete and at a deployable state.

Note that when instantiating a complex VNF, the stand-alone option is **not** selected.

4 Click Close when finished.

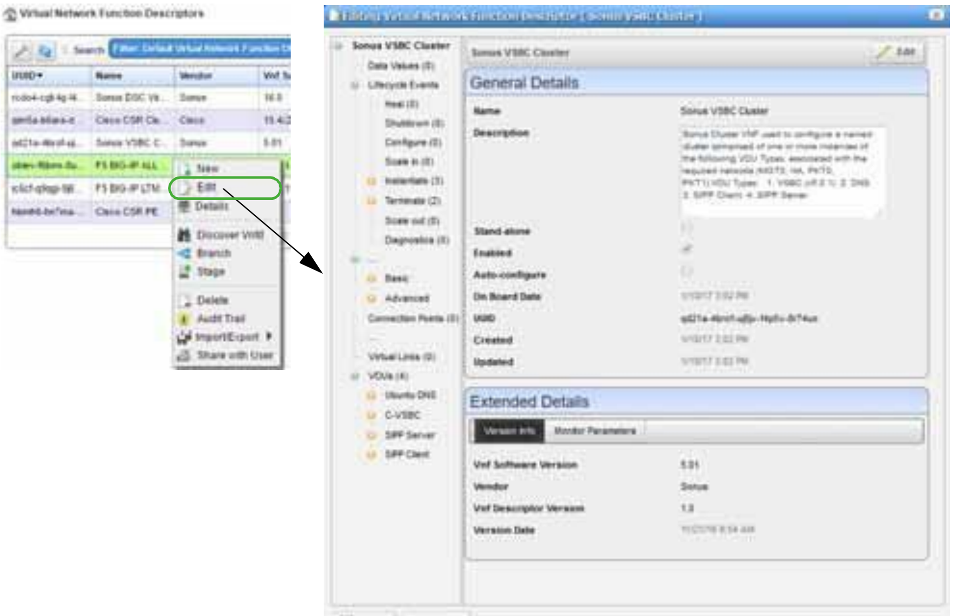

UUID▾	Name	Vendor	Vnf S
rcdo4-cg8-4g-4i...	Sonus DSC V8...	Sonus	16.0
qmfa-bfiara-d...	Cisco CSR Clu...	Cisco	15.4.2
qd21a-4bndf-q...	Sonus VSBC C...	Sonus	5.01
[illegible]	F5 BIG-IP ALL		
iclcf-qlsgp-58...	F5 BIG-IP LTM		
hbmh6-hn7ma...	Cisco CSR PE		

Editing Virtual Network Function Descriptor [Sonus VSBC Cluster]

Sonus VSBC Cluster

General Details

Name: Sonus VSBC Cluster

Description: Sonus Cluster VNF used to configure a named cluster comprised of one or more instances of the following VDU Types, associated with the required networks (MGT0, HA, PKT0, PKT1) VDU Types: 1. VSBC (v5.0.1) 2. DNS 3. SIPP Client 4. SIPP Server

Stand-alone: □

Enabled: ✓

Auto-configure: □

On Board Date: 1/19/17 3:02 PM

UUID: qd21a-4bndf-q8js-hlp5v-8rT4ux

Created: 1/19/17 3:02 PM

Updated: 1/19/17 3:02 PM

Extended Details

Vnf Software Version: 5.01

Vendor: Sonus

Vnf Descriptor Version: 1.0

Version Date: 11/27/16 8:54 AM

Now you are ready to stage a VNF record.

Staging a Complex VNF Record

The staging process creates and stages a VNF record based on the selected descriptors information found and then reserves resources. The staging process only reserves resources, it does not deploy the VNF record to the OpenStack controller.

Stage a complex VNF record from the Virtual Network Function Records portlet as follows.

NOTE:

The portlet’s location depends on your company’s configuration.

1 Right-click > New.

The Stage VNF window displays a list of descriptor entities.

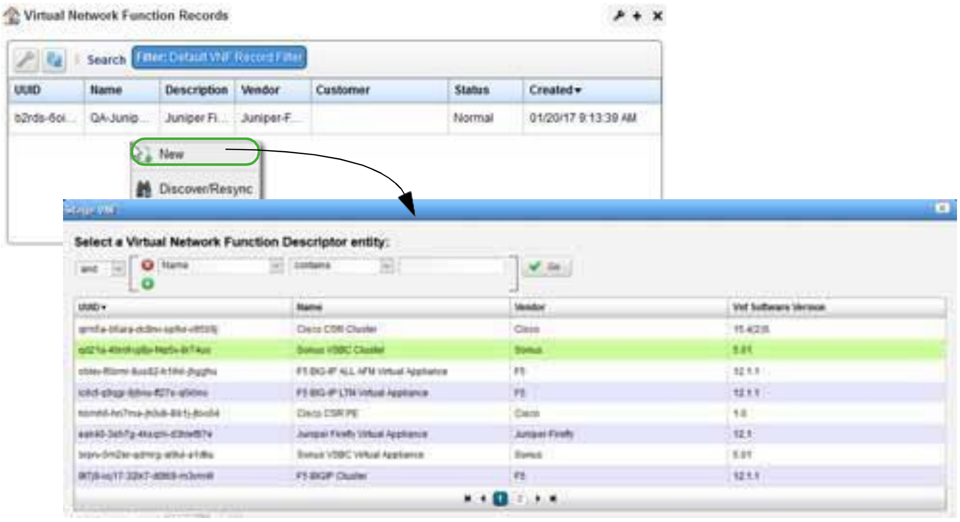
Virtual Network Function Records
Search
Filter: Default VNF Record Filter
New
Discover/Resync
Stage VNF
and
Name
contains
Go

UUID	Name	Description	Vendor	Customer	Status	Created▾
b2rds-6oi...	QA-Junip...	Juniper Fi...	Juniper-F...		Normal	01/20/17 9:13:39 AM

Select a Virtual Network Function Descriptor entity:

UUID▾	Name	Vendor	Vnf Software Version
qmfa-bfiara-dcdnv-sjdlvi-v8t5t9	Cisco CSR Cluster	Cisco	15.4(2)S
qd21a-4bndf-q8js-hlp5v-8rT4ux	Sonus VSBC Cluster	Sonus	5.01
obtev-f6brni-6us82-h1hhl-jhgghu	F5 BIG-IP ALL AFM Virtual Appliance	F5	12.1.1
iclcf-qlsgp-8jbnu-f27s-q5t0ns	F5 BIG-IP LTM Virtual Appliance	F5	12.1.1
hbmh6-hn7ma-jh0s8-8h1j-j6vd4	Cisco CSR PE	Cisco	1.0
aah40-3sh7g-4ksqzh-d3hef67e	Juniper Firefly Virtual Appliance	Juniper Firefly	12.1
bnrv-5m2er-admrg-a8kd-a1dfu	Sonus VSBC Virtual Appliance	Sonus	5.01
8t7j8-iaj17-32k7-d869-m3vmi8	F5 BIGIP Cluster	F5	12.1.1

2 Select the service to instantiate.

3 Click Select.

The Vnf Stage Parameters window is displayed.

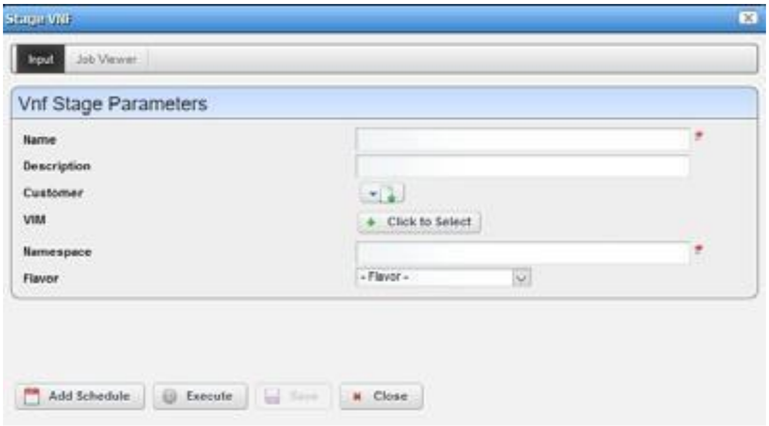

4 Enter a name for the VNF and a description that includes specifics on the services.

5 Specify the parameters, such as the VIM (Virtualized Infrastructure Manager) on which to deploy.

NOTE:

If you do not specify the VIM, the system determines the appropriate VIM. The Name, VIM, Namespace, and Flavor fields are required.

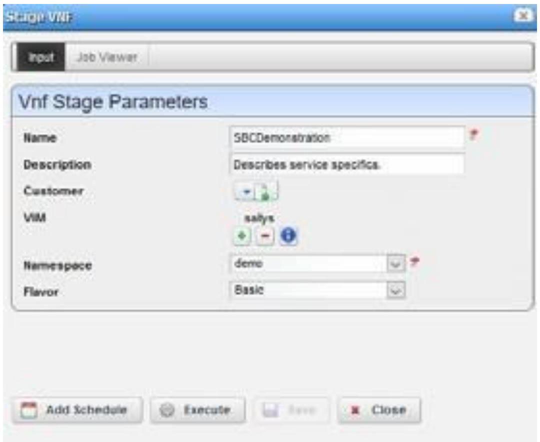

Name: SBCDemonstration

Description: Describes service specifics.

VIM: sallys

Namespace: demo

Flavor: Basic

6 Click Execute.

Upon completion, the stage results are displayed.

If staging is successful, the record is created and staged based on the descriptor information found and resourcesarereserved. However, nothing is deployed to the OpenStack controller.

If staging fails, the reason for the failure is displayed.

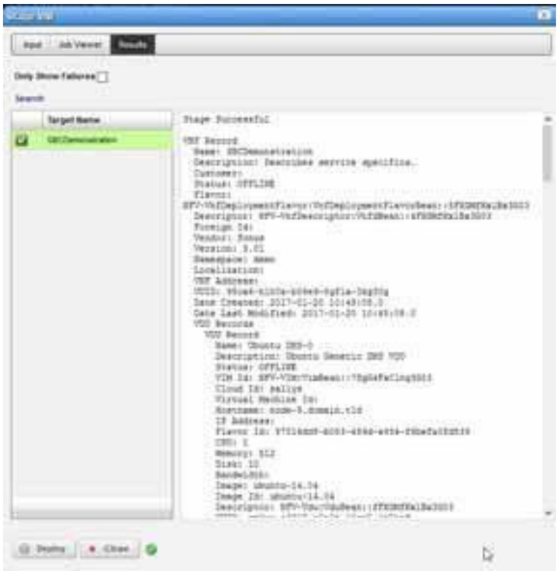

```
Stage Successful

VNF Record
  Name: SBCDemonstration
  Description: Describes service specifica...
  Customer:
  Status: OFFLINE
  Flavor:
SFV-VnfDeploymentFlavor:VnfDeploymentFlavorBean::[illegible]
  Descriptor: SFV-VnfDescriptor:VnfdBean::[illegible]
  Foreign Id:
  Vendor: Sonus
  Version: 5.01
  Namespace: demo
  Localization:
  VNF Address:
  UUID: [illegible]
  Date Created: 2017-01-20 10:48:08.0
  Date Last Modified: 2017-01-20 10:48:08.0
  VDU Records
    VDU Record
      Name: Ubuntu DNS-0
      Description: Ubuntu Generic DNS VDU
      Status: OFFLINE
      VIM Id: SFV-Vim:VimBean::[illegible]
      Cloud Id: [illegible]
      Virtual Machine Id:
      Hostname: node-5.domain.tld
      IP Address:
      Flavor Id: [illegible]
      CPU: 1
      Memory: 512
      Disk: 10
      Bandwidth:
      Image: ubuntu-14.04
      Image Id: ubuntu-14.04
      Descriptor: SFV-Vdu:VduBean::[illegible]
```

7 Click Deploy once the staging operation successfully completes.

Otherwise, optionally verify that the resources were reserved and then deploy the VNF record.

Verifying Reserved Resources

Verifythat the resourceswerereserved from the Virtualized Infrastructure Managers portlet as follows. You can also view the capacity details.

NOTE:

The portlet’s location depends on your company’s configuration.

1 Select a VIM.

2 Right-click > Details.

The Details are displayed. Notice the reserved CPU, memory, and disk values.

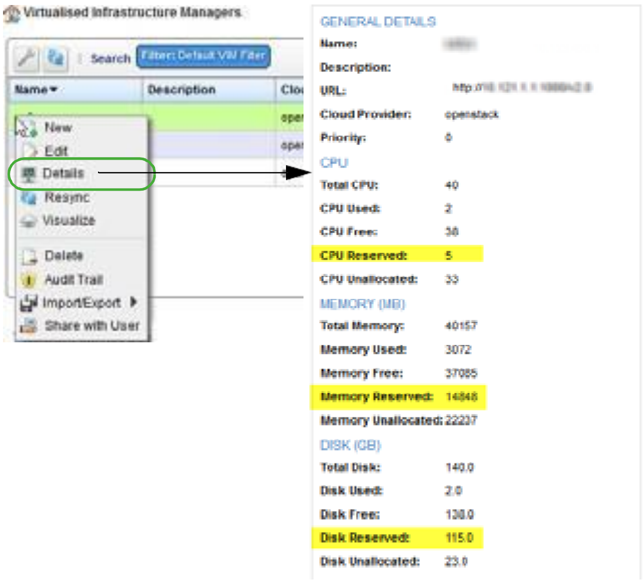

GENERAL DETAILS

Name: [redacted]

Description:

URL: http:// [redacted]

Cloud Provider: openstack

Priority: 0

CPU

Total CPU: 40

CPU Used: 2

CPU Free: 38

CPU Reserved: 5

CPU Unallocated: 33

MEMORY (MB)

Total Memory: 40157

Memory Used: 3072

Memory Free: 37085

Memory Reserved: 14848

Memory Unallocated: 22237

DISK (GB)

Total Disk: 140.0

Disk Used: 2.0

Disk Free: 138.0

Disk Reserved: 115.0

Disk Unallocated: 23.0

3 Go to the Hypervisor list.

4 Select a hypervisor.

5 Right-click > Details.

The VDU Reservation details show the reservations for the virtual networks to instantiate.

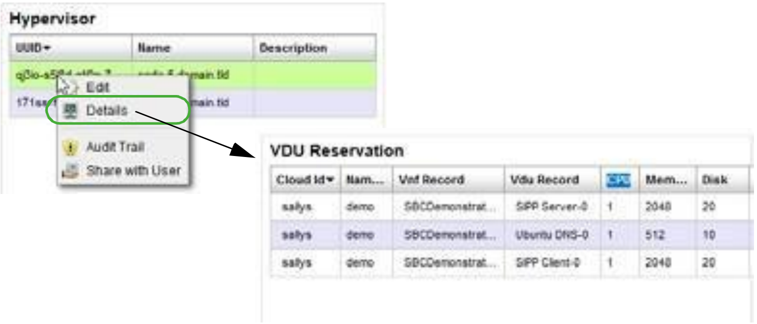

VDU Reservation

Cloud Id	Nam...	Vnf Record	Vdu Record	CPU	Mem...	Disk
sallys	demo	SBCDemonstrat...	SIPP Server-0	1	2048	20
sallys	demo	SBCDemonstrat...	Ubuntu DNS-0	1	512	10
sallys	demo	SBCDemonstrat...	SIPP Client-0	1	2048	20

6 Select a VDU.

7 Right-click > Details.

The VDU Reservation details are displayed.

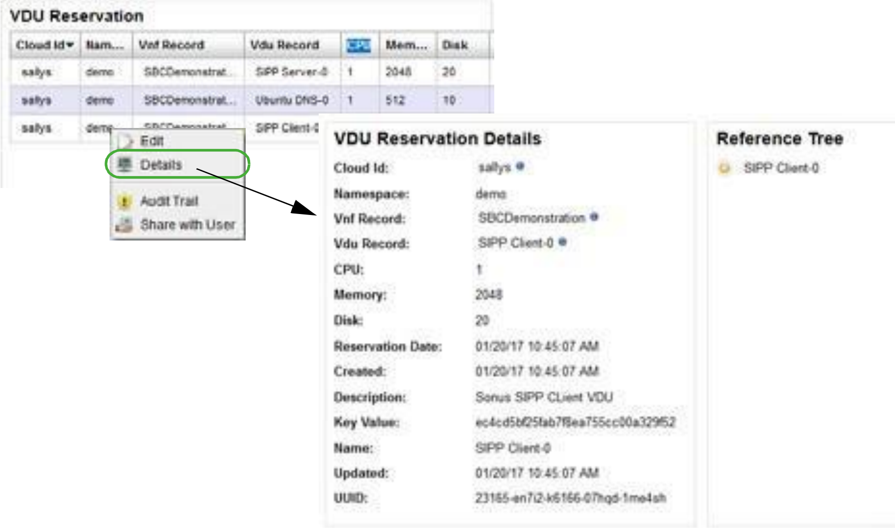

VDU Reservation

Cloud Id	Nam...	Vnf Record	Vdu Record	CPU	Mem...	Disk
sallys	demo	SBCDemonstrat...	SIPP Server-0	1	2048	20
sallys	demo	SBCDemonstrat...	Ubuntu DNS-0	1	512	10
sallys	demo	SBCDemonstrat...	SIPP Client-0			

VDU Reservation Details

Cloud Id:	sallys
Namespace:	demo
Vnf Record:	SBCDemonstration
Vdu Record:	SIPP Client-0
CPU:	1
Memory:	2048
Disk:	20
Reservation Date:	01/20/17 10:45:07 AM
Created:	01/20/17 10:45:07 AM
Description:	Sonus SIPP CLient VDU
Key Value:	ec4cd5bf25fab7f8ea755cc00a329f52
Name:	SIPP Client-0
Updated:	01/20/17 10:45:07 AM
UUID:	23165-en7i2-k6166-07hqd-1me4sh

Reference Tree

SIPP Client-0

Once you optionally verified the reserved resources, deploy the VNF record.

Deploying a Complex VNF Record

The deployment process sets up the network connectivity for a staged VNF record. Deploy a complex VNF record from the Virtual Network Functions Records portlet as follows.

NOTE:

The portlet's location depends on your company's configuration.

1 Select the VNF record you staged.

2 Right-click > Deploy.

The default Vnf Deploy Parameters are displayed.

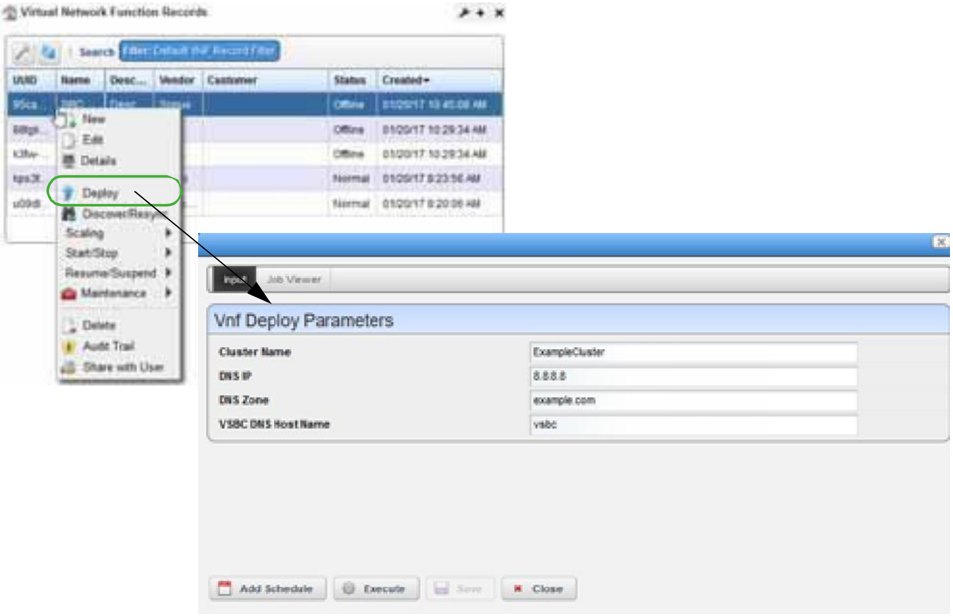

Virtual Network Function Records

Search Filter: Default VNF Record Filter

UUID	Name	Desc...	Vendor	Customer	Status	Created
95ca...	SBC	Desc...	Sonus		Offline	01/20/17 10:45:06 AM
88tgi...					Offline	01/20/17 10:29:34 AM
k3fw-...					Offline	01/20/17 10:29:34 AM
kps3t...					Normal	01/20/17 8:23:56 AM
u09di...					Normal	01/20/17 8:20:06 AM

Vnf Deploy Parameters

Cluster Name	ExampleCluster
DNS IP	8.8.8.8
DNS Zone	example.com
VSBC DNS Host Name	vsbc

3 Change default parameters as needed.

4 Click Execute.

The Job Viewer displays the deployoperation results, such as the steps taken to instantiate the networks needed to talk to each other and the information on the individual Virtualization Deployment Units (VDUs).

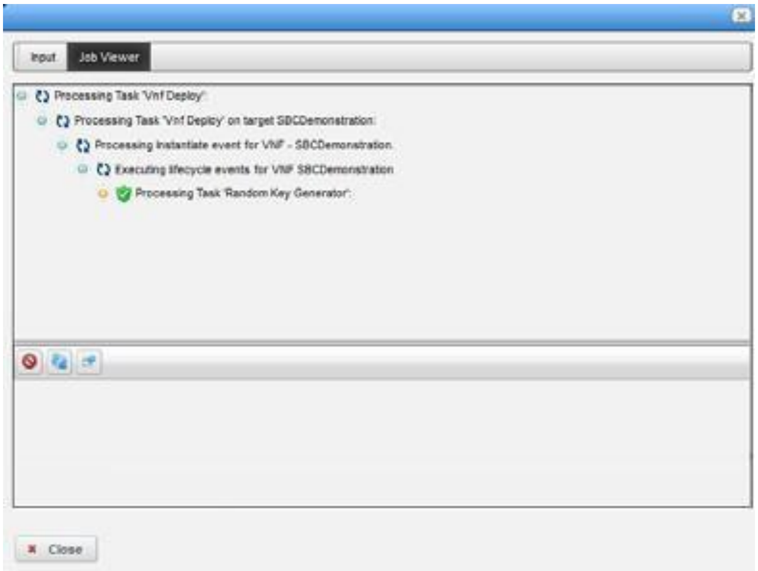

```
Processing Task 'Vnf Deploy':
    Processing Task 'Vnf Deploy' on target SBCDemonstration:
        Processing Instantiate event for VNF - SBCDemonstration.
            Executing lifecycle events for VNF SBCDemonstration
                Processing Task 'Random Key Generator':
```

5 Click Close.

Once deployed, you can view VNF information provided, verifythe instantiated VNF, and review the VDUR status.

Viewing Alarms and Notifications

The Alarms page allows you to view the alarms, notifications, and event history provided for the selected VNF.

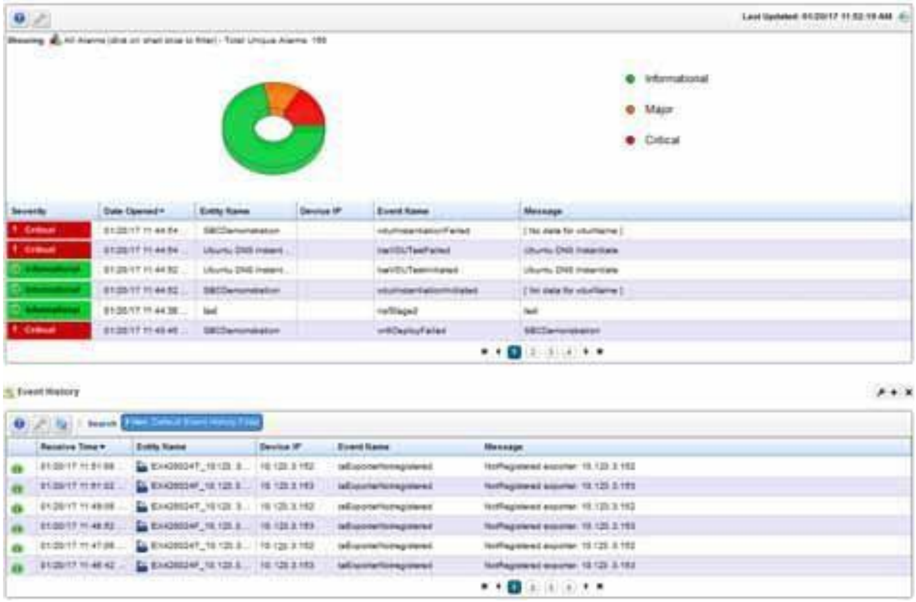

Last Updated: 01/20/17 11:52:19 AM
Showing: All Alarms (click on chart slice to filter) - Total Unique Alarms: 198
Informational
Major
Critical

Severity	Date Opened	Entity Name	Device IP	Event Name	Message
Critical	01:20:17 11:44:54 ...	SBCDemonstration		vduInstantiationFailed	[No data for vduName]
Critical	01:20:17 11:44:54 ...	Ubuntu DNS Instant...		lceVDUTestFailed	Ubuntu DNS Instantiate
Informational	01:20:17 11:44:52 ...	Ubuntu DNS Instant...		lceVDUTerminated	Ubuntu DNS Instantiate
Informational	01:20:17 11:44:52 ...	SBCDemonstration		vduInstantiationInitiated	[No data for vduName]
Informational	01:20:17 11:44:38 ...	test		nsStaged	test
Critical	01:20:17 11:43:46 ...	SBCDemonstration		vnfDeployFailed	SBCDemonstration

Receive Time	Entity Name	Device IP	Event Name	Message
01:20:17 11:51:08 ...	EX4200-24T_10.120.3...	10.120.3.152	isExporterNotregistered	NotRegistered exporter: 10.120.3.152
01:20:17 11:51:02 ...	EX4200-24F_10.120.3...	10.120.3.153	isExporterNotregistered	NotRegistered exporter: 10.120.3.153
01:20:17 11:49:08 ...	EX4200-24T_10.120.3...	10.120.3.152	isExporterNotregistered	NotRegistered exporter: 10.120.3.152
01:20:17 11:48:52 ...	EX4200-24F_10.120.3...	10.120.3.153	isExporterNotregistered	NotRegistered exporter: 10.120.3.153
01:20:17 11:47:08 ...	EX4200-24T_10.120.3...	10.120.3.152	isExporterNotregistered	NotRegistered exporter: 10.120.3.152
01:20:17 11:46:42 ...	EX4200-24F_10.120.3...	10.120.3.153	isExporterNotregistered	NotRegistered exporter: 10.120.3.153

Next verify the instantiated VNF and then review the VDUR status

Verifying a Complex VNF Instantiation

Verify the VNF instantiation from the Functional Resources page as follows.

1 Verify that the Managed Resources portlet lists the VNF record that you created.

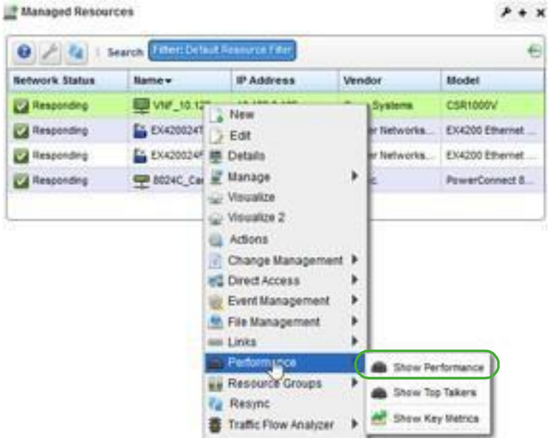

Network Status	Name	IP Address	Vendor	Model
Responding	VNF_10.12...	10.1...	...Systems	CSR1000V
Responding	EX4200-24T		...r Networks...	EX4200 Ethernet ...
Responding	EX4200-24F		...r Networks...	EX4200 Ethernet ...
Responding	8024C_Ca...		...c	PowerConnect 8...

2 Select a virtual network.

3 Right-click > Performance > Show Performances to see that you are collecting call rate data.

The Performance Dashboard is displayed.

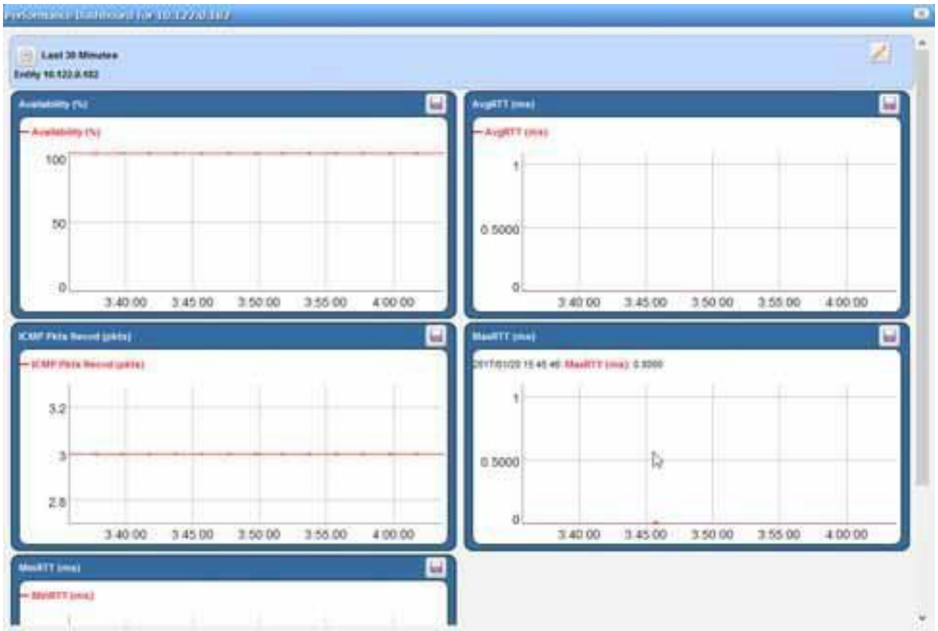

4 Navigate to the Virtual Network Function Records portlet.

5 Verify that the record status is Normal.

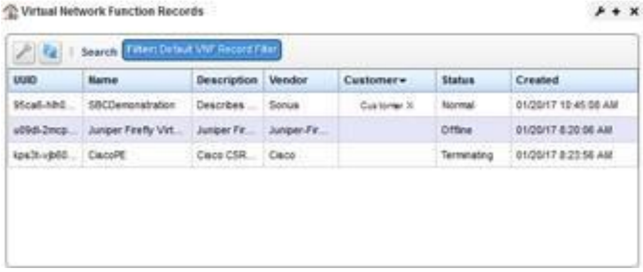
Virtual Network Function Records

Search Filter: Default VNF Record Filter

UUID	Name	Description	Vendor	Customer	Status	Created
95ca8-hh0...	SBCDemonstration	Describes ...	Sonus	Customer X	Normal	01/20/17 10:45:08 AM
a09dl-2mcp...	Juniper Firefly Virt...	Juniper Fir...	Juniper-Fir...		Offline	01/20/17 8:20:06 AM
kpa3t-vjb60...	CiscoPE	Cisco CSR...	Cisco		Terminating	01/20/17 8:23:56 AM

The last step is to review the VNF record status.

Reviewing the VNF Record Status

Review the VNF Record status from the Virtual Network Function Records portlet. By default, the Statuscolumn showseach VNFrecord'sstatus. However, youcanalsoreviewall VNFcomponent's status from the Details portlets as follows.

1 Select a VNF record.

2 Right-click > Details.

The General details are displayed. Notice that the status is normal.

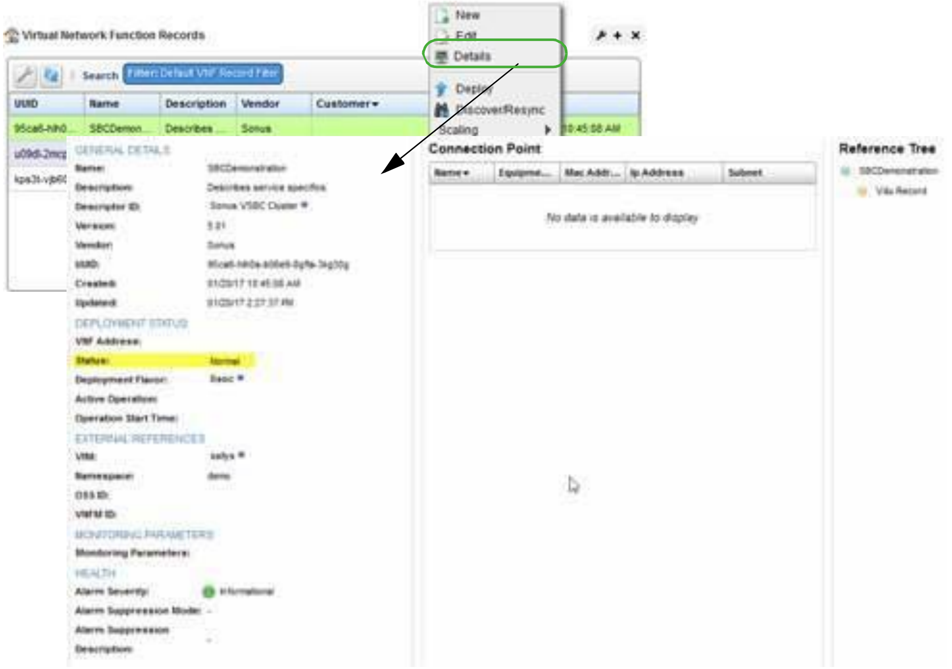
Virtual Network Function Records
Search
Filter: Default VNF Record Filter
New
Edit
Details
Deploy
Discover/Resync
Scaling

UUID	Name	Description	Vendor	Customer
95ca6-nh0...	SBCDemon...	Describes ...	Sonus	
u09di-2mcj				
kpa3t-vjb6				

GENERAL DETAILS

Name:	SBCDemonstration
Description:	Describes service specifics
Descriptor ID:	Sonus VSBC Cluster
Version:	5.21
Vendor:	Sonus
UUID:	95ca6-nh0a-a06e6-0gfe-3xg30g
Created:	01/20/17 10:45:58 AM
Updated:	01/20/17 2:27:37 PM

DEPLOYMENT STATUS

VNF Address:	
Status:	Normal
Deployment Flavor:	Basic
Active Operation:	
Operation Start Time:	

EXTERNAL REFERENCES

VIM:	sallys
Namespace:	demo
OSS ID:	
VNFM ID:	

MONITORING PARAMETERS

Monitoring Parameters:	

HEALTH

Alarm Severity:	● Informational
Alarm Suppression Mode:	-
Alarm Suppression Description:	-

Connection Point

Name	Equipme...	Mac Addr...	Ip Address	Subnet

No data is available to display

Reference Tree

- SBCDemonstration
  - Vdu Record

a. Click the VDU Information tab.
VDU and VNF Virtual Links portlets are displayed and show that the VDUs came up successfully.

b. Select a VDU.

c. Right-click > Details.
The General details for the selected VDU are displayed.
The virtual machine ID is used to manage the VDU. The equipment ID is used to interact with the virtual network in Cruz as if it were a device.

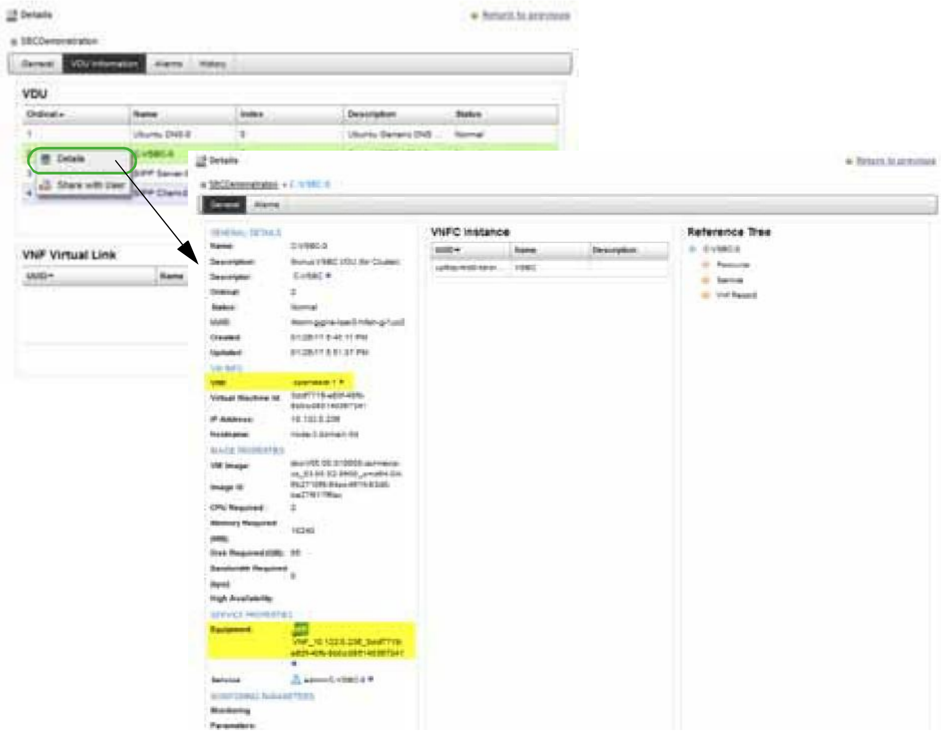

Details

SBCDemonstration

General | VDU Information | Alarms | History

VDU

Ordinal	Name	Index	Description	Status
1	Ubuntu DNS-0	3	Ubuntu Generic DNS ...	Normal

Details

SBCDemonstration ▸ C-VSBC-0

General | Alarms

GENERAL DETAILS

Name:	C-VSBC-0
Description:	Sonus VSBC VDU (for Cluster)
Descriptor:	C-VSBC
Ordinal:	2
Status:	Normal
UUID:	[illegible]
Created:	01/26/17 5:41:11 PM
Updated:	01/26/17 5:51:37 PM
VIM:	openstack 1
Virtual Machine Id:	[illegible]
IP Address:	10.122.0.236
Hostname:	[illegible]
VM Image:	[illegible]
Image Id:	[illegible]
CPU Required:	2
Memory Required (MB):	10240
Disk Required (GB):	65
Bandwidth Required (bps):	0
High Availability:	
Equipment:	VNF_10.122.0.236_[illegible]
Service:	[illegible]
Monitoring Parameters:	

VNFC Instance

UUID	Name	Description
[illegible]	VSBC	

Reference Tree

3 View the VNF connection instance details.

a. Select a VNF connection.

b. Right-click > Details.

The VNF Connection Point Instance list is displayed.

c. Select a connection instance.

d. Right-click > Details.

The VNF's General connection point details are displayed.

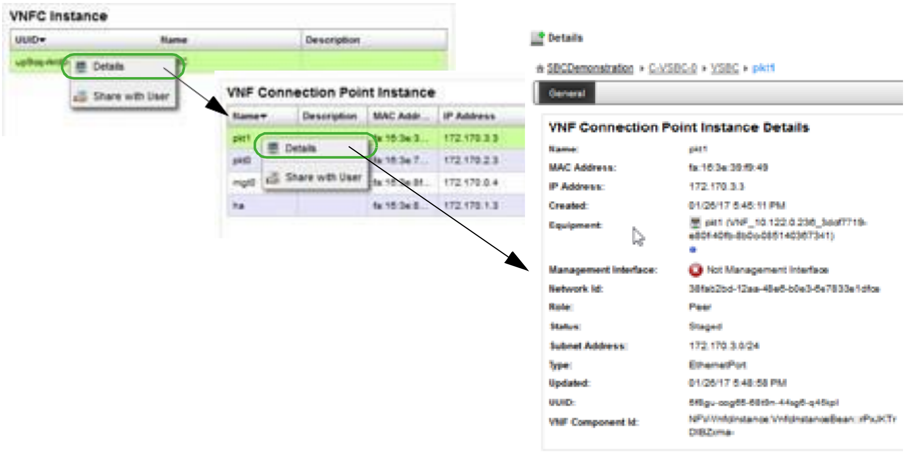

VNFC Instance

UUID	Name	Description
[illegible]	[illegible]	

VNF Connection Point Instance

Name	Description	MAC Addr...	IP Address
pkt1		fa:16:3e:3...	172.170.3.3
pkt0		fa:16:3e:7...	172.170.2.3
mgt0		fa:16:3e:8f...	172.170.0.4
ha		fa:16:3e:8...	172.170.1.3

Details

SBCDemonstration ▸ C-VSBC-0 ▸ VSBC ▸ pkt1

General

VNF Connection Point Instance Details

Name:	pkt1
MAC Address:	fa:16:3e:39:f9:49
IP Address:	172.170.3.3
Created:	01/26/17 5:45:11 PM
Equipment:	pkt1 (VNF_10.122.0.236_bddf7719-e80f-40fb-8b0c-085140367341)
Management Interface:	Not Management Interface
Network Id:	38fab2bd-12aa-48e6-b0e3-6e7833e1dfce
Role:	Peer
Status:	Staged
Subnet Address:	172.170.3.0/24
Type:	EthernetPort
Updated:	01/26/17 5:48:58 PM
UUID:	5f8gu-oog65-68t9n-44og6-q45kpl
VNF Component Id:	NPVVnfdInstance:VnfdInstanceBean: rPxJKTr D8Zcma-

4 View the resource details.

a. Click the NFVI Resources tab.

The Virtualized Infrastructure Managers portlet is displayed.

b. Select a resource.

c. Right-click > Details.

The General details are displayed. Notice the change in CPU usage.

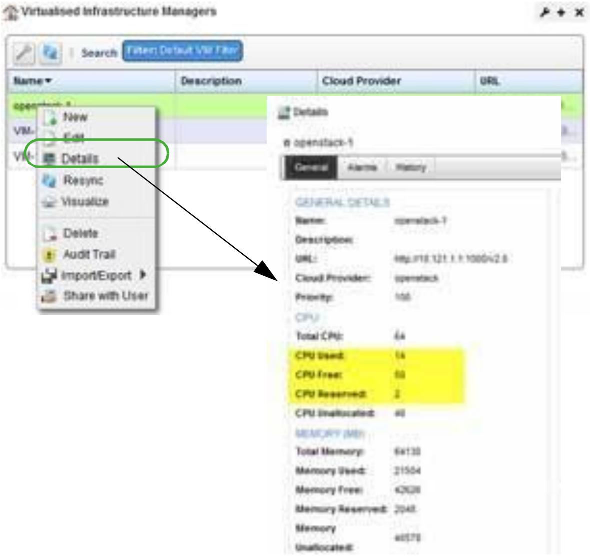

GENERAL DETAILS

Name: openstack-1

Description:

URL: http://10.121.1.1:1000/v2.0

Cloud Provider: openstack

Priority: 100

CPU

Total CPU: 64

CPU Used: 14

CPU Free: 50

CPU Reserved: 2

CPU Unallocated: 48

MEMORY (MB)

Total Memory: 64130

Memory Used: 21504

Memory Free: 42626

Memory Reserved: 2048

Memory Unallocated: 40578

d. Select a hypervisor.

e. Right-click > Details.

Note that there are no more reservations and the CPU, memory, and disk usage is higher.

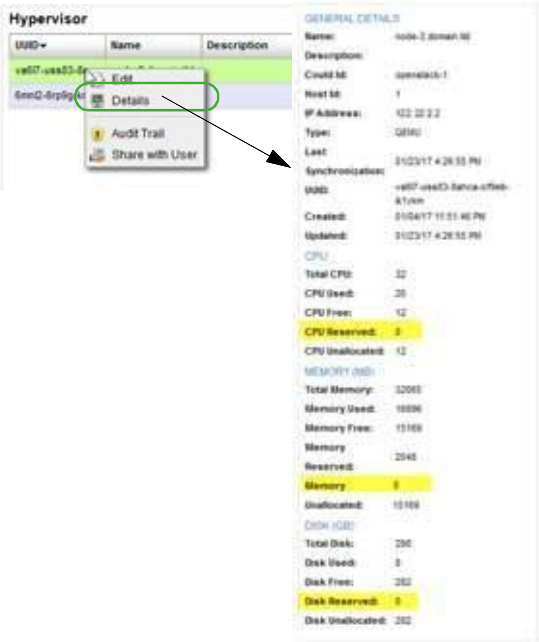

GENERAL DETAILS

Name:	node-3.domain.tld
Description:	
Could Id:	openstack-1
Host Id:	1
IP Address:	122.22.2.2
Type:	QEMU
Last Synchronization:	01/23/17 4:26:55 PM
UUID:	va6f7-usa83-8ahca-cffeb-a7ckm
Created:	01/04/17 11:51:46 PM
Updated:	01/23/17 4:26:55 PM

CPU

Total CPU:	32
CPU Used:	20
CPU Free:	12
CPU Reserved:	0
CPU Unallocated:	12

MEMORY (MB)

Total Memory:	32065
Memory Used:	16896
Memory Free:	15169
Memory Reserved:	2048
Memory Unallocated:	0 15169

DISK (GB)

Total Disk:	290
Disk Used:	8
Disk Free:	282
Disk Reserved:	0
Disk Unallocated:	282

5 View network topology and compute instances from the OpenStack dashboard.

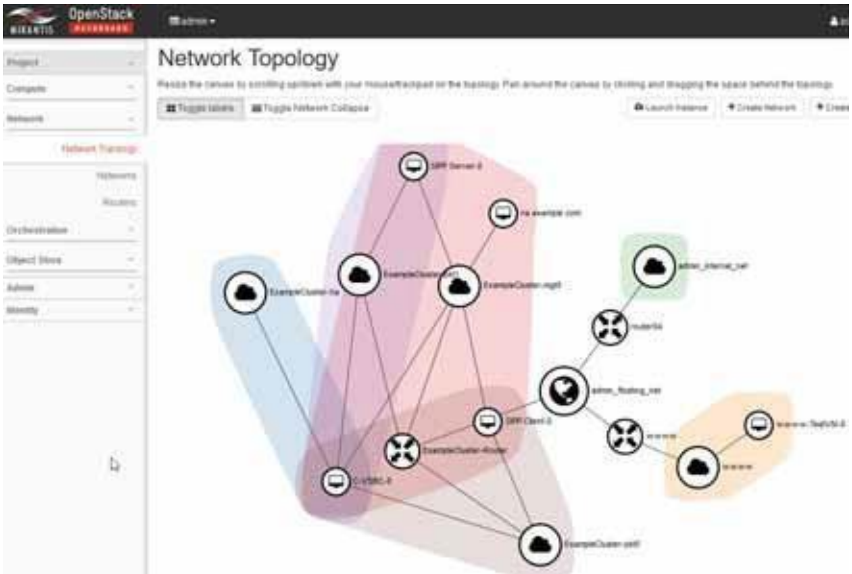

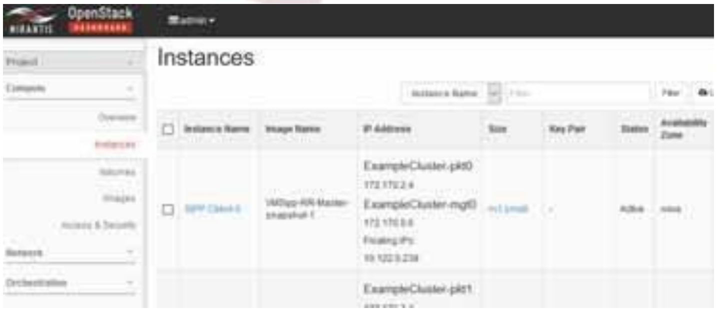

Instances

□	Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone
□	SIPP-Client-0	VMSpp-RR-Master-snapshot-1	ExampleCluster-pkt0 172.172.2.4 ExampleCluster-mgt0 172.170.0.6 Floating IPs: 10.122.0.238	m1.small	-	Active	nova
			ExampleCluster-pkt1				

Once you instantiate a complex VNF (stage and deploy), you can perform dailyoperations on the VNF, such as monitoring alarms, updating virtual network functions (scale in/out), and undeploying virtual network functions .

Instantiating a Standalone VNF Record

This section shows how to instantiate a standalone Virtualized Network Function (VNF) record. The example standalone VNF descriptor used contains all lifecycle events (LCEs) required to create a set of Virtualization Deployment Units (VDUs) and contains the networks required to interconnect VDUs in the OpenStack environment. In this example, only a VNF record is required because there are no other VNFs required to interconnect with the VNF descriptor, and because the same target OpenStack project space contains all the VDUs.

Instantiate a standalone VNF record by:

- Staging a Standalone VNF Record
- Deploying a Standalone VNF Record
- Verifying VM Network Connections

Staging a Standalone VNF Record

The staging process creates and stages a Virtualized Network Function (VNF) record based on the selected descriptor’s information found and then reserves resources. The staging process does not deploy the VNF record to the OpenStack controller.

Stage a standalone VNF record from the Virtual Network Function Records portlet as follows.

NOTE:

The portlet’s location depends on your company’s configuration.

1 Right-click > New.

The Stage VNF window displays a list of descriptor entities.

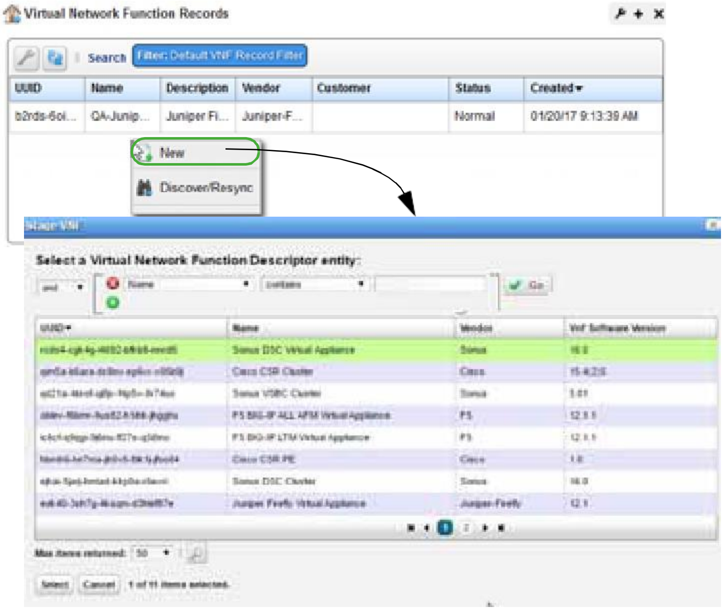
Virtual Network Function Records
Search
Filter: Default VNF Record Filter
New
Discover/Resync
Stage VNF
Go
Select
Cancel

UUID	Name	Description	Vendor	Customer	Status	Created▾
b2rds-6oi...	QA-Junip...	Juniper Fi...	Juniper-F...		Normal	01/20/17 9:13:39 AM

Select a Virtual Network Function Descriptor entity:

and ▾ Name ▾ contains ▾

UUID▾	Name	Vendor	Vnf Software Version
[illegible]	Sonus DSC Virtual Appliance	Sonus	16.0
[illegible]	Cisco CSR Cluster	Cisco	15.4.2.S
[illegible]	Sonus VSBC Cluster	Sonus	5.01
[illegible]	F5 BIG-IP ALL AFM Virtual Appliance	F5	12.1.1
[illegible]	F5 BIG-IP LTM Virtual Appliance	F5	12.1.1
[illegible]	Cisco CSR PE	Cisco	1.0
[illegible]	Sonus DSC Cluster	Sonus	16.0
[illegible]	Juniper Firefly Virtual Appliance	Juniper-Firefly	12.1

Max items returned: 50 ▾

1 of 11 items selected.

2 Select the service you want to instantiate.

3 Click Select.

The Vnf Stage Parameters window is displayed.

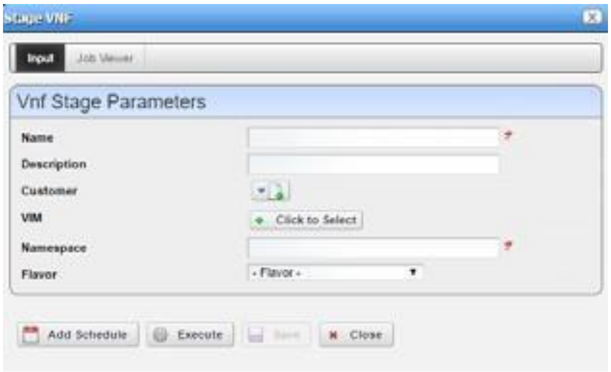

Vnf Stage Parameters

Name	
Description	
Customer	
VIM	Click to Select
Namespace	
Flavor	- Flavor -

4 Enter a name for the VNF and a description that includes specifics on the services.

5 Specify the parameters, such as the VIM (Virtualized Infrastructure Manager) on which to deploy.

The Name, VIM, Namespace, and Flavor fields are required.

NOTE:

If you do not specify the VIM, the system determines the appropriate VIM. The Name, VIM, Namespace, and Flavor fields are required.

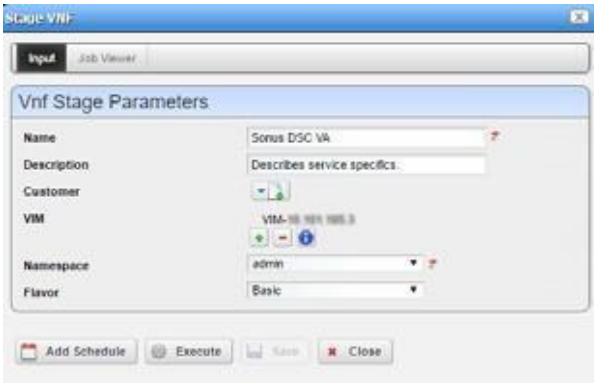

Vnf Stage Parameters

Name	Sonus DSC VA
Description	Describes service specifics.
Customer	
VIM	VIM-[illegible]
Namespace	admin
Flavor	Basic

6 Click Execute.

Upon completion, the stage results are displayed.

If staging is successful, the record is created and staged based on the descriptor information found and resourcesarereserved. However, nothing is deployed to the OpenStack controller.

If staging fails, the reason for the failure is displayed.

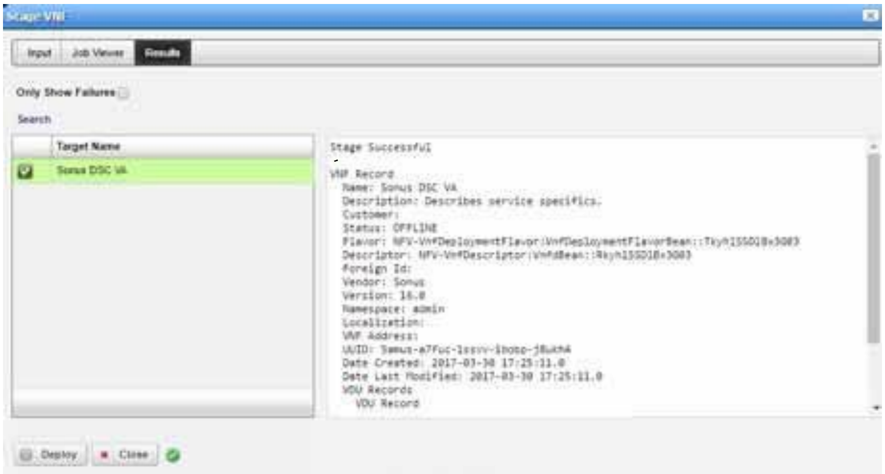

```
Stage Successful
-
VNF Record
  Name: Sonus DSC VA
  Description: Describes service specifics.
  Customer:
  Status: OFFLINE
  Flavor: NFV-VnfDeploymentFlavor:VnfDeploymentFlavorBean::Tkyh15SD18x3083
  Descriptor: NFV-VnfDescriptor:VnfdBean::Rkyh15SD18x3083
  Foreign Id:
  Vendor: Sonus
  Version: 16.0
  Namespace: admin
  Localization:
  VNF Address:
  UUID: Sonus-a7fuc-1ssvv-lhobo-j8ukh4
  Date Created: 2017-03-30 17:25:11.0
  Date Last Modified: 2017-03-30 17:25:11.0
  VDU Records
    VDU Record
```

Now that the stage process successfully completed, it is time to deploy the VNF record.

Deploying a Standalone VNF Record

Deployastandalone Virtualized Network Function (VNF) record from the Stage VNF Results panel or the VNF record from the Virtual Network Function Records portlet as follows. The example instructions here deploy a VNF from the Stage VNF Results panel.

Deploy a standalone VNF record from the Stage VNF Results panel as follows.

1 Click Deploy.

The default VNF deploy parameter values are displayed.

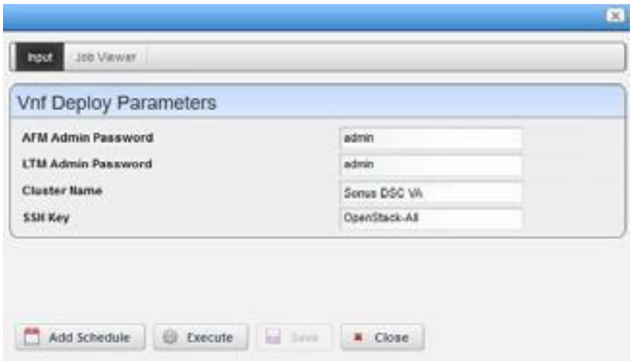

Vnf Deploy Parameters

AFM Admin Password	admin
LTM Admin Password	admin
Cluster Name	Sonus DSC VA
SSH Key	OpenStack-All

2 Change the default parameters as needed.

3 Click Execute.

The Job Viewer displays the deployoperation results, suchas the steps taken to instantiate the networks needed to talk to each other and the information on the individual Virtualization Deployment Units (VDUs).

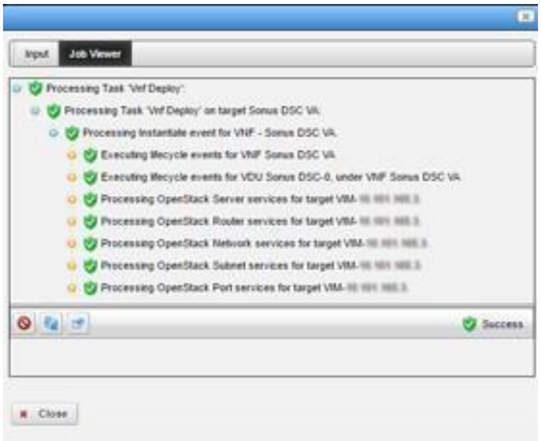

Processing Task 'Vnf Deploy':
Processing Task 'Vnf Deploy' on target Sonus DSC VA.
Processing Instantiate event for VNF - Sonus DSC VA.
Executing lifecycle events for VNF Sonus DSC VA
Executing lifecycle events for VDU Sonus DSC-0, under VNF Sonus DSC VA
Processing OpenStack Server services for target VIM-[redacted]
Processing OpenStack Router services for target VIM-[redacted]
Processing OpenStack Network services for target VIM-[redacted]
Processing OpenStack Subnet services for target VIM-[redacted]
Processing OpenStack Port services for target VIM-[redacted]

Success

After successfully deploying the VNF record, make sure that the VM’s network connections are up and okay.

Verifying VM Network Connections

Verify that the VM’s network connections are up and okay from the OpenStack dashboard as follows.

1 Select Project > Compute > Instances.

A list of instances and their information is displayed.

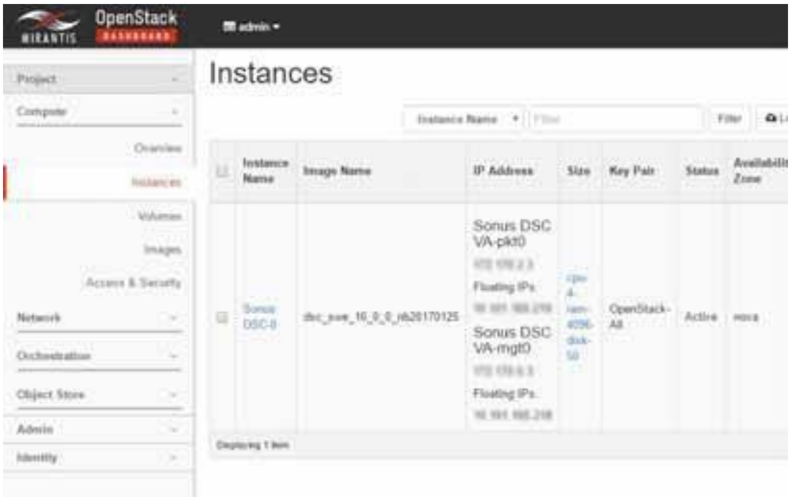

Instances

□	Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availabili Zone
□	Sonus DSC-0	dsc_swe_10_0_0_rb20170125	Sonus DSC VA-pkt0 [redacted] Floating IPs: [redacted] Sonus DSC VA-mgt0 [redacted] Floating IPs: [redacted]	cpu-4-ram-4096-disk-50	OpenStack-All	Active	nova

Displaying 1 item

2 Verify that the resource records were deployed.

a. Wait 5 minutes.

b. Go to the Managed Resources portlet.

The list shows the created Virtualized Network Function (VNF) records deployed to the OpenStack system. If the firmware and software versions show, the device is ready for SNMP requests.

NOTE:

If you do not have resource management implemented, ignore step h.

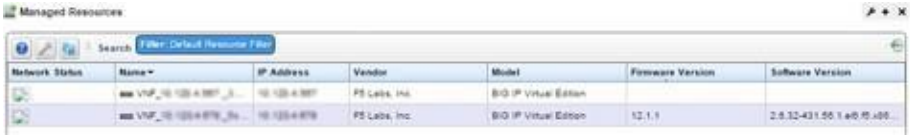

Managed Resources

Search Filter: Default Resource Filter

Network Status	Name	IP Address	Vendor	Model	Firmware Version	Software Version
	VNF_[illegible]	[illegible]	F5 Labs, Inc.	BIG IP Virtual Edition		
	VNF_[illegible]	[illegible]	F5 Labs, Inc	BIG IP Virtual Edition	12.1.1	2.6.32-431.56.1.el6.f5.x86...

When the system detects that the device is up, the system automatically runs the Resync and other Discovery Profile tasks for the VNF record. If successful, the VNF resource shows up as a managed resource.

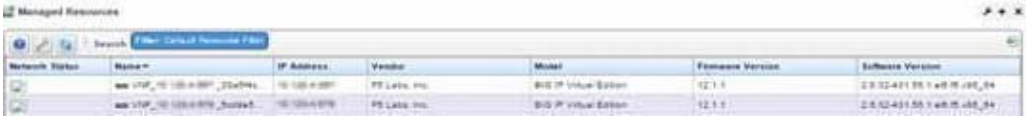

Managed Resources

Search Filter: Default Resource Filter

Network Status	Name	IP Address	Vendor	Model	Firmware Version	Software Version
	VNF_[illegible]	[illegible]	F5 Labs, Inc.	BIG IP Virtual Edition	12.1.1	2.6.32-431.56.1.el6.f5.x86_64
	VNF_[illegible]	[illegible]	F5 Labs, Inc.	BIG IP Virtual Edition	12.1.1	2.6.32-431.56.1.el6.f5.x86_64

3 Check networkconnectivity to each VNF descriptor using its management interface, the assigned floating IP address, and both ping and ssh commands.

Once you instantiate a VNF (stage and deploy), you can perform dailyoperations on the VNF, such as monitoring alarms, updating virtual network functions (scale in/out), and undeploying virtual network functions.

Monitoring Network Health

The Cruz application has many tools to monitor network health, determine whether any action is required, and then take the appropriate action, such as stage and deploy new virtual services or devices, scale network usage in/out, or undeploying virtual services or devices.

Here are a few of the tasks to monitor network health:

- Monitoring Alarms
- Viewing Managed Resources Performance
- Viewing VNF Record Details
- Viewing VIM Details
- Viewing Resource Monitors and Thresholds
- Reviewing Event Rules

Monitoring Alarms

Monitor alarms from the Alarms portlet. Here are the actions that you can perform on alarms:

- Take ownership of those alarms you plan to address (Acknowledge Alarm).
- Assign selected alarms to another user on the system to resolve (Assign User).
- Remove resolved alarms from the list and mark them for database archiving (Clear Alarm).
- Remove many alarms that are old, low severity, or both at the same time (Clear Group of Alarms). Forexample, you have many information alarms that are over a week old and do not want to clear them one at a time. This action requires a filter for the alarm group and the **action is irreversible**.

For a detailed description of this portlet and its options, see Alarms Portlet on page 276.

Monitor alarms from the Cruz application as follows.

1 Navigate to the Alarms portlet.

2 Review alarms whose severity is Warning, Minor, Major, or Critical.

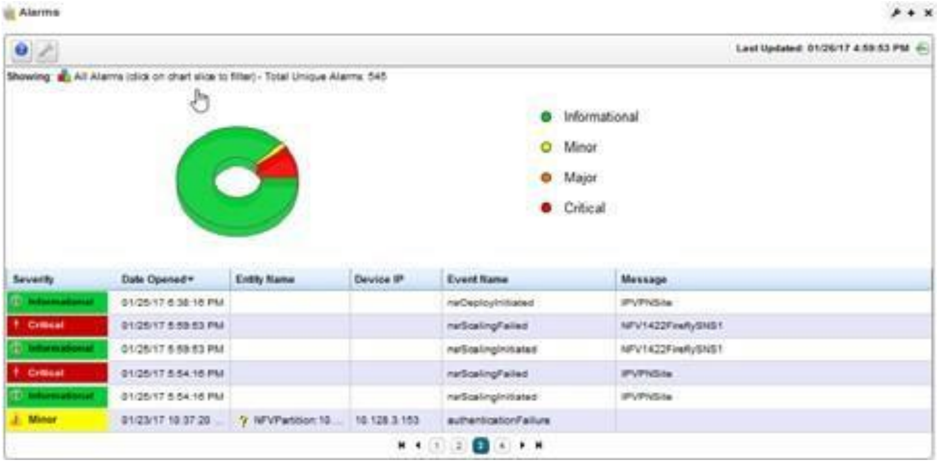

Alarms

Last Updated: 01/26/17 4:59:53 PM

Showing: All Alarms (click on chart slice to filter) - Total Unique Alarms: 545

Severity	Date Opened	Entity Name	Device IP	Event Name	Message
Informational	01/25/17 6:38:16 PM			nsDeployInitiated	IPVPNSite
Critical	01/25/17 5:59:53 PM			nsScalingFailed	NFV1422FireflySNS1
Informational	01/25/17 5:59:53 PM			nsScalingInitiated	NFV1422FireflySNS1
Critical	01/25/17 5:54:16 PM			nsScalingFailed	IPVPNSite
Informational	01/25/17 5:54:16 PM			nsScalingInitiated	IPVPNSite
Minor	01/23/17 10:37:20 ...	NFVPartition:10...	10.128.3.153	authenticationFailure	

3 Acknowledge those alarms that you plan to address.

Address alarms from the most sever state to the least sever state (Critical, Major, Minor, and Warning).

4 Address alarm issues appropriately.

Viewing Managed Resources Performance

A managed resource’s Performance Dashboard shows patterns for CPU and memory utilization, packet counts, RTT (round-trip time) measurements, and so on. Viewing VIM Details on page 969 also shows resource usage (memory, CPU, and disk). For a detailed description of this portlet and its options, see Performance Dashboard on page 419.

View a managed resource’s performance from the Managed Resources portlet as follows.

1 Select a resource.

2 Right-click > Performance > Show Performance.

The Performance Dashboard displays data patterns based on the defined properties, entities, and view attributes defined.

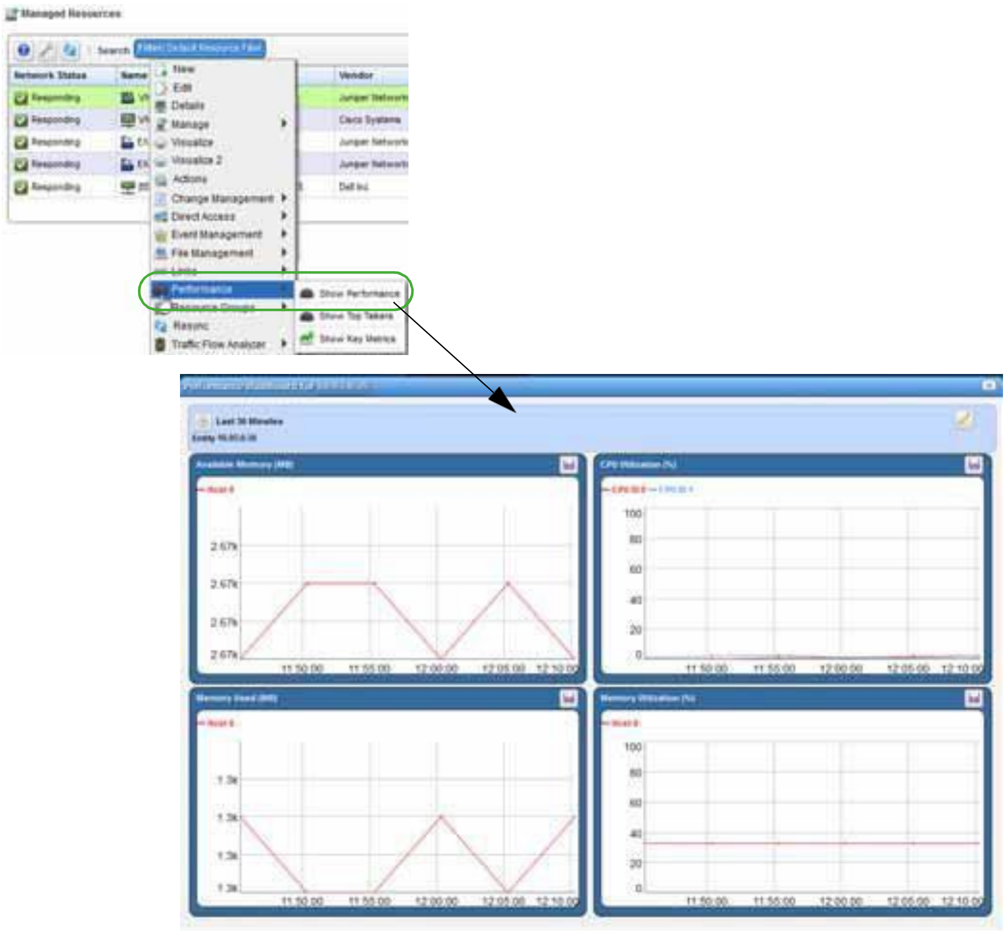

Viewing VNF RecordDetails

Viewa VNF record’s details to see its general information, deployment status, external references, monitoring parameters, and health. The VNF record details also provide information about connection points, VDUs, VNF virtual links, alarms, event history, and actions audit trail.

View VNF record details from the Virtual Network Function Records portlet as follows.

1 Select a record.

2 Right-click > Details.

The Generaldetailsaredisplayed, whichincludesgeneraldetails, deploymentstatus, external references, monitoring parameters, health, connection points, and a reference tree.

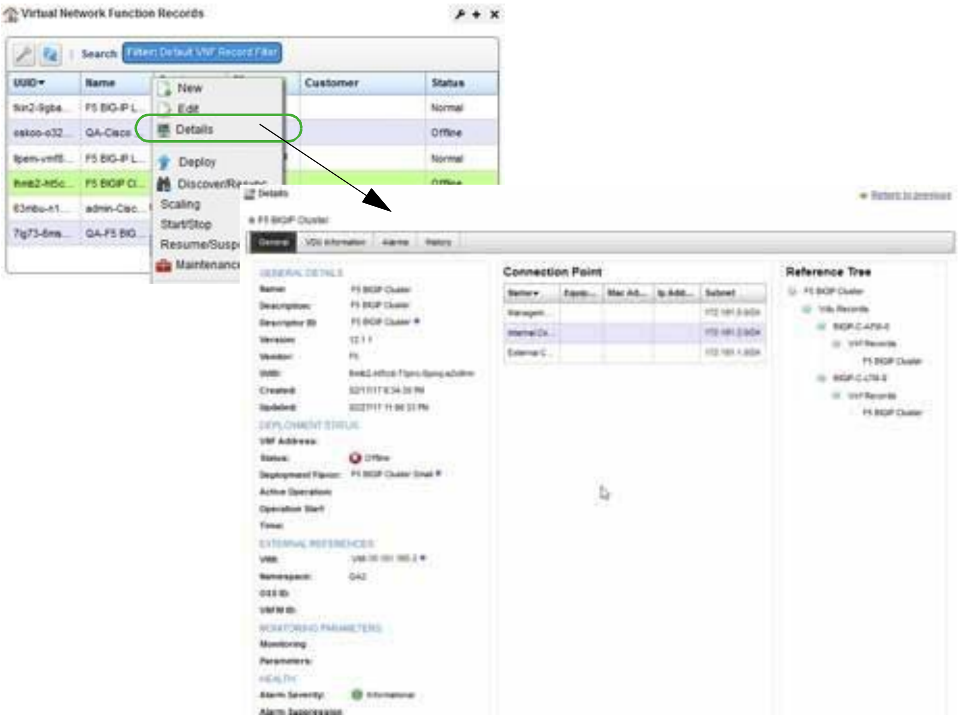

UUID	Name	Customer	Status
9in2-9gba...	F5 BIG-IP L...		Normal
eskoo-o32...	QA-Cisco...		Offline
9pem-vmf8...	F5 BIG-IP L...		Normal
8mb2-h5c...	F5 BIGIP Cl...		Offline
63mbu-n1...	admin-Cisc...		
7g73-6ms...	QA-F5 BIG...		

3 Click the VDU Information tab.

A list of related VDUs and VNF virtual links are displayed.

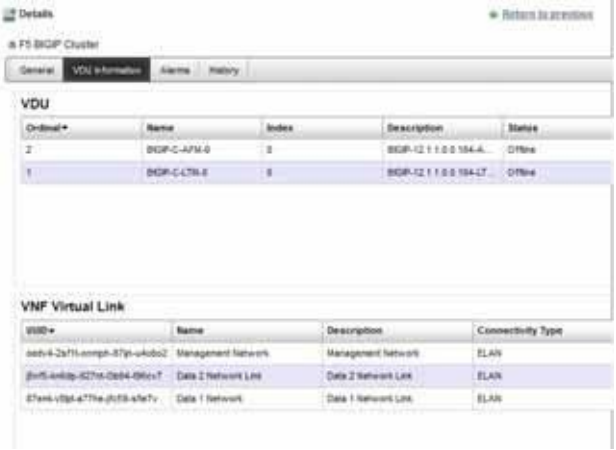

VDU

Ordinal	Name	Index	Description	Status
2	BIGIP-C-AFM-0	0	BIGIP-12.1.1.0.0.184-A...	Offline
1	BIGIP-C-LTM-0	0	BIGIP-12.1.1.0.0.184-LT...	Offline

VNF Virtual Link

UUID	Name	Description	Connectivity Type
sedv4-2a7h-onmph-87pi-u4obo2	Management Network	Management Network	ELAN
jbnf5-kn8dp-827nt-0b84-0l6cvT	Data 2 Network Link	Data 2 Network Link	ELAN
87ank-v0pt-a77he-jfc68-sfe7v	Data 1 Network	Data 1 Network Link	ELAN

4 Click Alarms or History.

The Alarms details provides a list of alarms and a list of event history. The History details is an audit trail of actions.

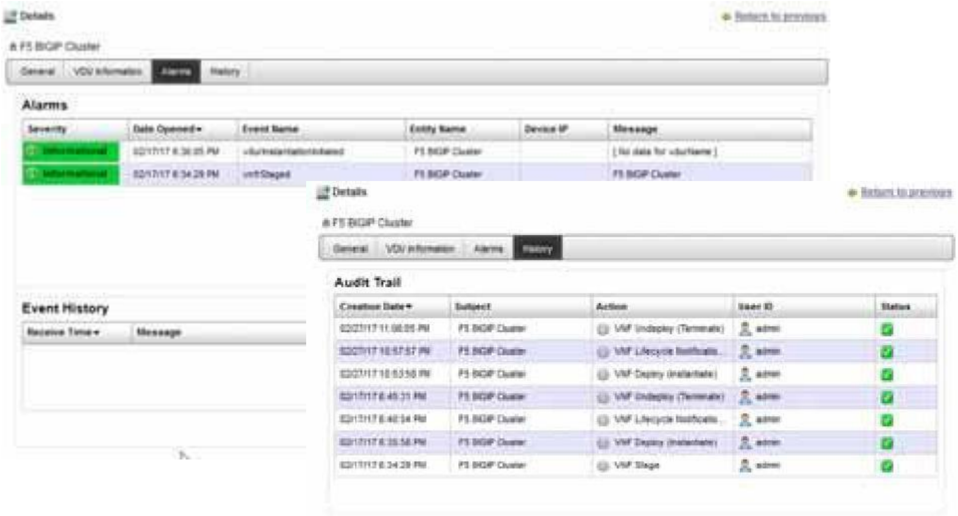

Alarms

Severity	Date Opened▾	Event Name	Entity Name	Device IP	Message
Informational	02/17/17 6:38:05 PM	vduInstantiationInitiated	F5 BIGIP Cluster		[No data for vduName]
Informational	02/17/17 6:34:28 PM	vnfStaged	F5 BIGIP Cluster		F5 BIGIP Cluster

Event History

Receive Time▾	Message

Audit Trail

Creation Date▾	Subject	Action	User ID	Status
02/27/17 11:08:05 PM	F5 BIGIP Cluster	VNF Undeploy (Terminate)	admin	✓
02/27/17 10:57:57 PM	F5 BIGIP Cluster	VNF Lifecycle Notificatio...	admin	✓
02/27/17 10:53:58 PM	F5 BIGIP Cluster	VNF Deploy (Instantiate)	admin	✓
02/17/17 6:45:31 PM	F5 BIGIP Cluster	VNF Undeploy (Terminate)	admin	✓
02/17/17 6:40:54 PM	F5 BIGIP Cluster	VNF Lifecycle Notificatio...	admin	✓
02/17/17 6:35:58 PM	F5 BIGIP Cluster	VNF Deploy (Instantiate)	admin	✓
02/17/17 6:34:28 PM	F5 BIGIP Cluster	VNF Stage	admin	✓

Viewing VIM Details

The Virtualized Infrastructure Managers (VIM) details includes resource information (CPU, memoryand diskusage), anyconstraints, images, hypervisors, reference tree, alarms, event history, and action audit trail. Viewing Managed Resources Performance on page 966 shows usage patterns.

Youcan alsoviewinformation about the VIM from your OpenStacksystemifyou have permissions to access that system.

View VIMdetailsfrom the Virtualized Infrastructure Managersportletand the OpenStacksystem as follows.

1 Select a resource.

2 Right-click > Details.

The General details are displayed. Notice the amount of CPU, memory, and disk being used.

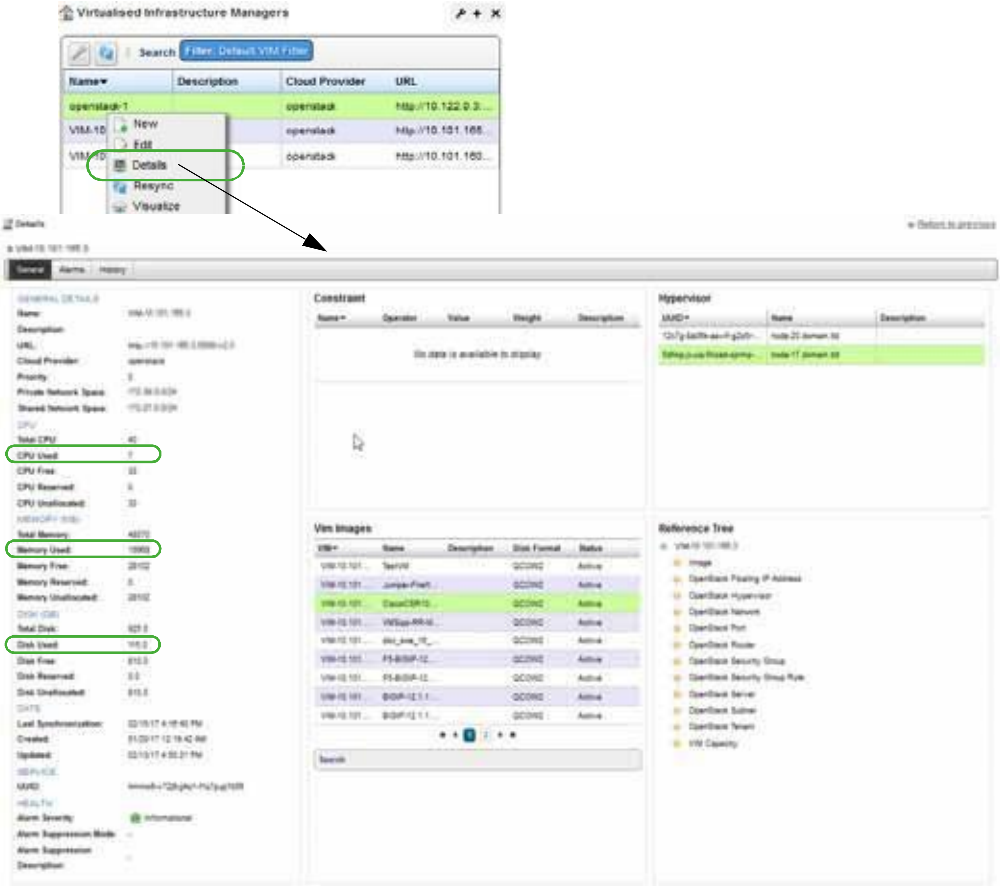

Virtualised Infrastructure Managers

Search Filter: Default VIM Filter

Name▾	Description	Cloud Provider	URL
openstack-1		openstack	http://10.122.0.3...
VIM-10		openstack	http://10.101.166...
VIM-10		openstack	http://10.101.160...

3 Go to the OpenStack dashboard.

4 Select Project > Network > Network Topology.

5 Review the network topology.

Notice the networks and VDUs deployed.

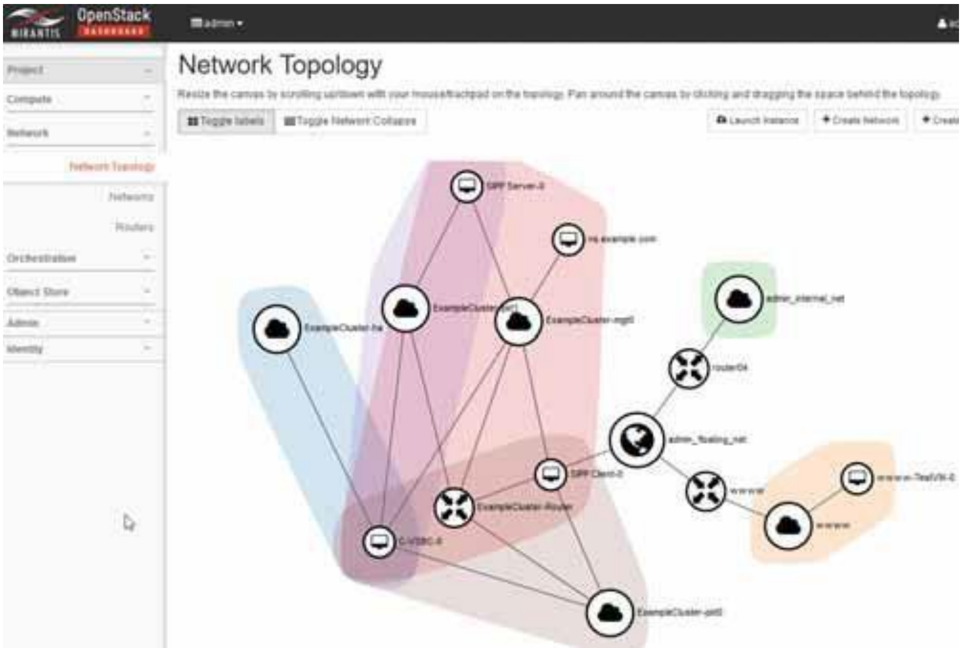

Network Topology

Resize the canvas by scrolling up/down with your mouse/trackpad on the topology. Pan around the canvas by clicking and dragging the space behind the topology.

6 Select Project > Compute > Instances.

A list of instances aredisplayed.

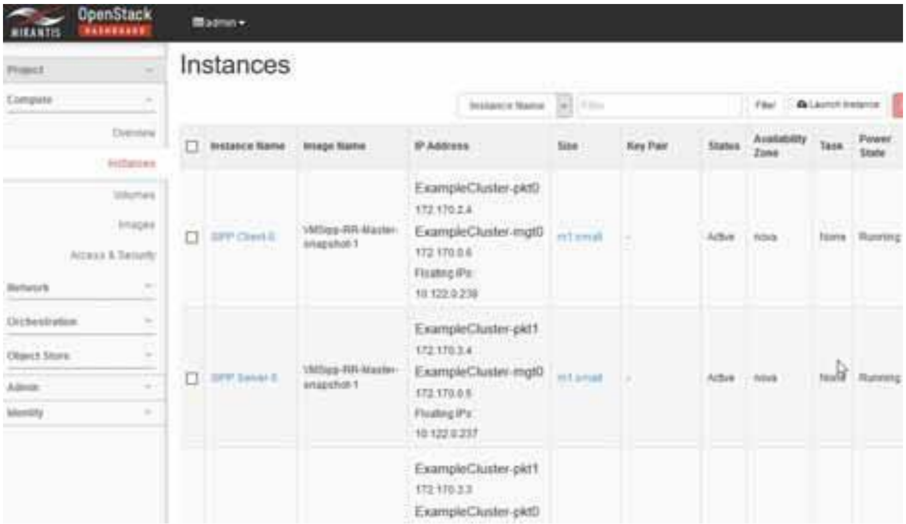

Instances

☐	Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State
☐	SIPP Client-0	VMSipp-RR-Master-snapshot-1	ExampleCluster-pkt0 172.170.2.4 ExampleCluster-mgt0 172.170.0.6 Floating IPs: 10.122.0.238	m1.small	-	Active	nova	None	Running
☐	SIPP Server-0	VMSipp-RR-Master-snapshot-1	ExampleCluster-pkt1 172.170.3.4 ExampleCluster-mgt0 172.170.0.5 Floating IPs: 10.122.0.237	m1.small	-	Active	nova	None	Running
			ExampleCluster-pkt1 172.170.3.3 ExampleCluster-pkt0						

7 Return to the Cruz application.

Viewing Resource Monitors and Thresholds

Look at your resource monitor configuration to understand the parameters set that could generate alerts and determine whether changes are needed as the network changes. Resource monitor configuration settings include:

- General Monitor and Retention options that are common to all different monitor types.
- Monitor Options that specify monitor targets and options specific to a monitor type.
- Calculated Metrics, where you create attributes that are calculated from existing monitor attributes.
- Thresholds, where you set threshold intervals on attributes in the monitor.
- Inventory Mapping, where you associate predefined inventory metrics with a monitored attribute to normalize the attribute if an VNF does not report metrics in a way that matches the monitored attribute's name or format. Available metrics include CPU utilization percentage, memory utilization percentage, bandwidth utilization percentage, and so on.
- Conditions, where you define conditions.

For a detailed description of this portlet and its options, see Resource Monitors on page 361.

View resource monitors and threshold information from the Resource Monitors portlet as follows.

1 Select a resource monitor.

2 Right-click > Edit Monitor.

The General monitor and retention options currently monitoring the call rate are displayed.

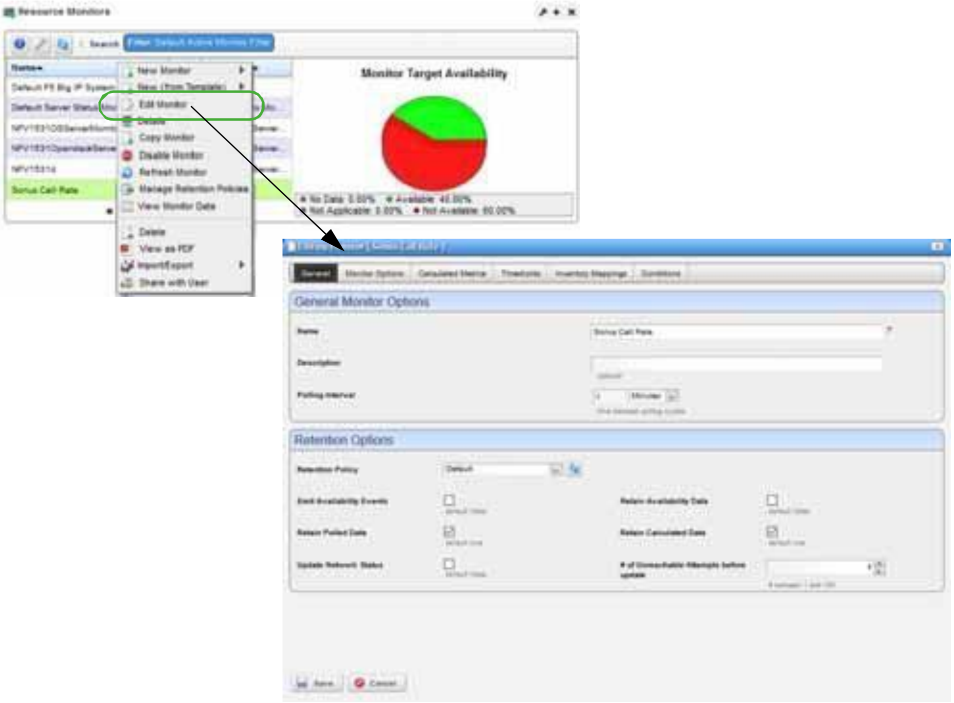

Name: Sonus Call Rate

Retention Policy: Default

3 Click the Monitor Options tab.

Notice the monitor's entities and SNMP attributes.

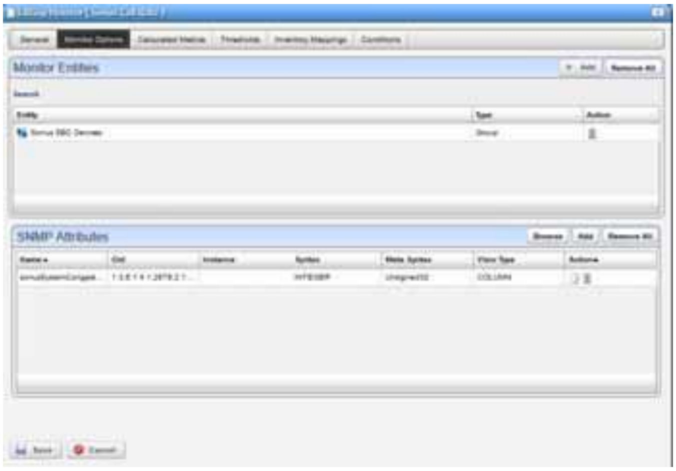

Monitor Entities

Search

Entity	Type	Action
Sonus SBC Devices	Group	

SNMP Attributes

Name	Oid	Instance	Syntax	Meta Syntax	Value Type	Actions
sonusSystemCongest...	1.3.6.1.4.1.2879.2.1...		INTEGER	Unsigned32	COLUMN	

4 Click the Thresholds tab.

The Configured Attributes are displayed.

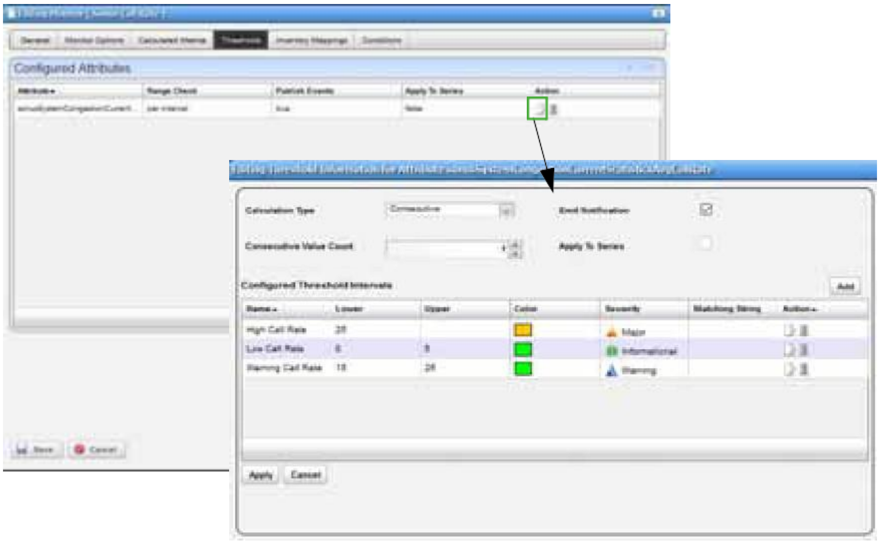

Configured Attributes

Attribute	Range Check	Publish Events	Apply To Series	Action
sonusSystemCongestionCurrent...	per-interval	true	false	

Calculation Type: Consecutive

Event Notification ☑

Consecutive Value Count: 1

Apply To Series ☐

Configured Threshold Intervals

Name	Lower	Upper	Color	Severity	Matching String	Action
High Call Rate	25		🟧	Major		
Low Call Rate	6	9	🟩	Informational		
Warning Call Rate	15	25	🟩	Warning		

5 Click the Edit action.

The configured threshold intervals are displayed, such as high, low, and warning call rates.

6 Review other configurations by clicking the appropriate tab.

7 Cancel out of the monitoring windows when you are done.

Reviewing Event Rules

Event processing rules define how the system responds to events. By default, seeded rules exist. However, if you have permission to, you can create your own rules (New Protocol Translation, Stream Based Correlation, Event Definition Override, or Automation), copy or modify existing rules, delete them, or Import rules from and Export rules to files. For a detailed description of this portlet and its options, see Automation and Event Processing Rules on page 288.

Review event rules from the Automation and Event Processing Rules portlet as follows.

1 Search for the processing rules.

The Event Processing Rulesportlet displays the searchresults.

2 Select the rule.

3 Right-click > Edit.

The definition for the selected rule is displayed.

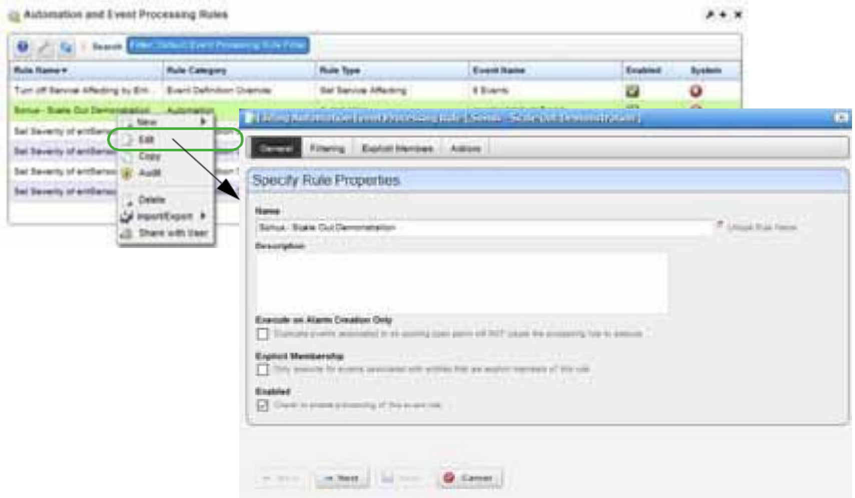

4 Review the event definition, filtering conditions, and actions to take.

A scale out scales the VNF to the next highest deployment flavor.

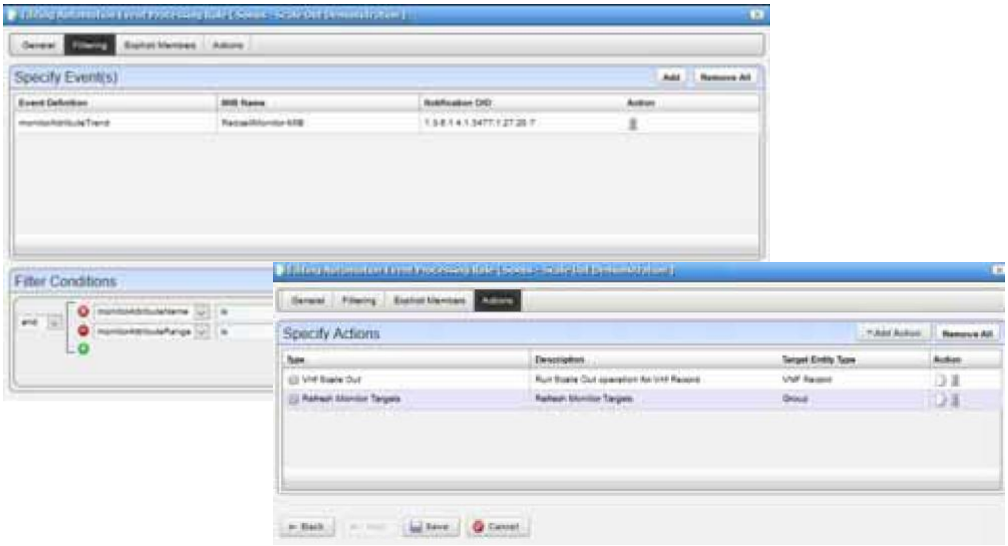

Updating Virtual Network Functions

Youshould update your virtualized network functions (VNFs) on a regular basis to make sure that you always have the latest definition and changes from the target VIM (Virtualized Infrastructure Manager.

Update a VNF from the Virtual Network Function Records portlet as follows.

1 Select a record.

2 Right-click > Discover/Resync.

The Discover VNF Records window is displayed.

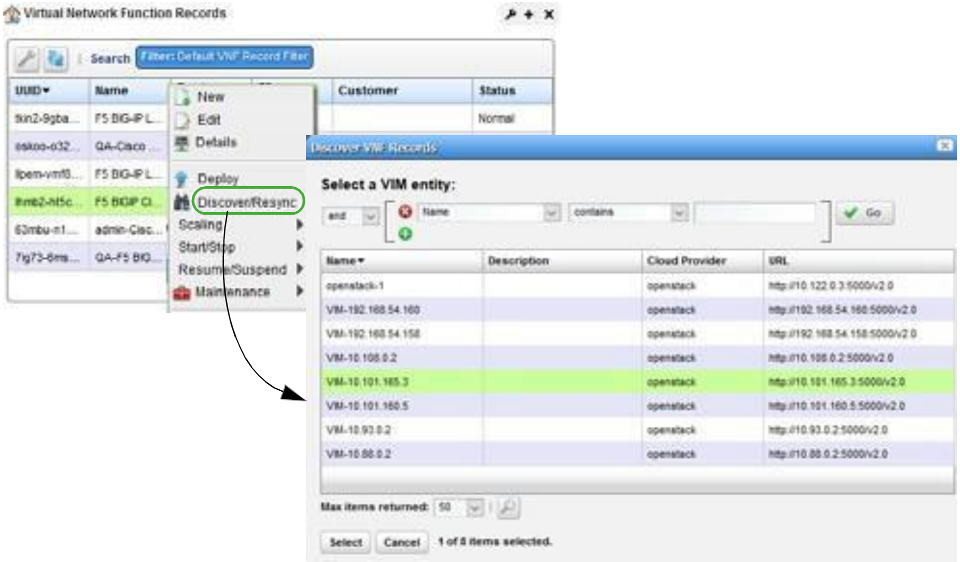

UUID	Name	Customer	Status
tkn2-9gba...	F5 BIG-IP L...		Normal
eskoo-032...	QA-Cisco ...		
lpem-vmf0...	F5 BIG-IP L...		
thmb2-ht5c...	F5 BIGIP Cl...		
63mbu-n1...	admin-Cisc...		
7ig73-6ms...	QA-F5 BIG...		

Select a VIM entity:

and | Name | contains

Name	Description	Cloud Provider	URL
openstack-1		openstack	http://10.122.0.3:5000/v2.0
VIM-192.168.54.160		openstack	http://192.168.54.160:5000/v2.0
VIM-192.168.54.158		openstack	http://192.168.54.158:5000/v2.0
VIM-10.108.0.2		openstack	http://10.108.0.2:5000/v2.0
VIM-10.101.165.3		openstack	http://10.101.165.3:5000/v2.0
VIM-10.101.160.5		openstack	http://10.101.160.5:5000/v2.0
VIM-10.93.0.2		openstack	http://10.93.0.2:5000/v2.0
VIM-10.88.0.2		openstack	http://10.88.0.2:5000/v2.0

Max items returned: 50

1 of 8 items selected.

3 Select the appropriate VIM entity.

4 Click Select.

The Vnf Discovery Parameters are displayed.

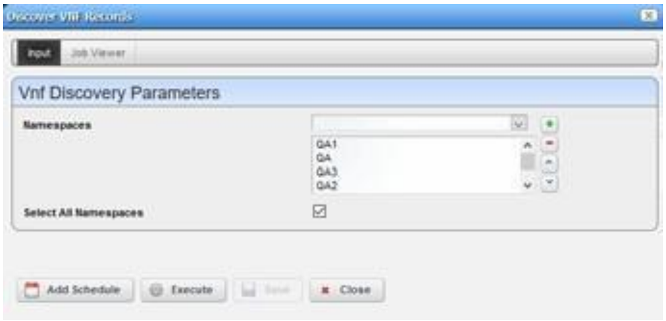

Vnf Discovery Parameters

Namespaces: QA1, QA, QA3, QA2

Select All Namespaces ☑

5 Select the namespaces.

6 Click Execute.

The Job Viewerdisplaystheexecutionprogress, anyerrors, and success/failure, followed bythe Results information, such as VDU, VNFC, and ot her record information. If you VNF is up to date, the following message is displayed.

```
No vnfs/changes found on VIM.
```

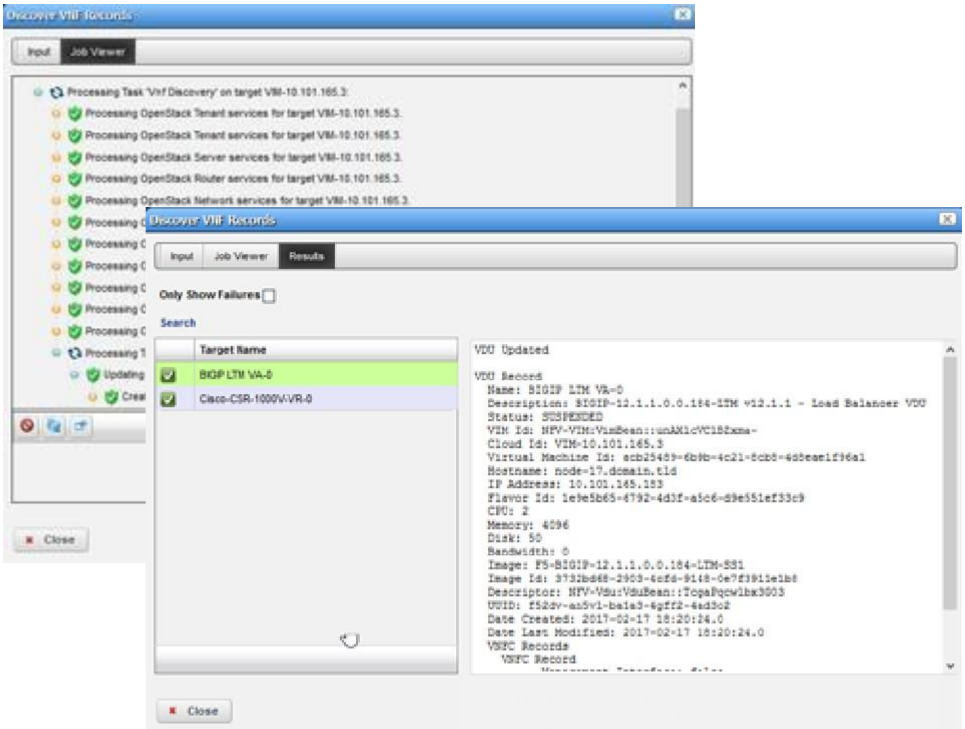

```
Processing Task 'Vnf Discovery' on target VIM-10.101.165.3:
  Processing OpenStack Tenant services for target VIM-10.101.165.3.
  Processing OpenStack Tenant services for target VIM-10.101.165.3.
  Processing OpenStack Server services for target VIM-10.101.165.3.
  Processing OpenStack Router services for target VIM-10.101.165.3.
  Processing OpenStack Network services for target VIM-10.101.165.3.
  Processing C
  Processing C
  Processing C
  Processing C
  Processing C
  Processing C
  Processing T
    Updating
      Creat
```

	Target Name
☑	BIGIP LTM VA-0
☑	Cisco-CSR-1000V-VR-0

```
VDU Updated

VDU Record
  Name: BIGIP LTM VA-0
  Description: BIGIP-12.1.1.0.0.184-LTM v12.1.1 - Load Balancer VDU
  Status: SUSPENDED
  VIM Id: NFV-VIM:VimBean::unAX1cVC182xma-
  Cloud Id: VIM-10.101.165.3
  Virtual Machine Id: ecb25489-6b9b-4c21-8cb8-4d8eae1f96a1
  Hostname: node-17.domain.tld
  IP Address: 10.101.165.183
  Flavor Id: 1e9e5b65-4792-4d3f-a5c6-d9e551ef33c9
  CPU: 2
  Memory: 4096
  Disk: 50
  Bandwidth: 0
  Image: F5-BIGIP-12.1.1.0.0.184-LTM-SS1
  Image Id: 3732bd48-2903-4cfd-9148-0e7f3911e1b8
  Descriptor: NFV-Vdu:VduBean::TogaPqcw1bx3003
  UUID: f52dv-an5v1-ba1a3-4gff2-4ad3o2
  Date Created: 2017-02-17 18:20:24.0
  Date Last Modified: 2017-02-17 18:20:24.0
  VNFC Records
    VNFC Record
```

7 Close the Discover VNF Records window.

Undeploying Virtual Network Functions

Undeploy virtualized network functions (VNFs) from the Virtual Network Function Records portlet as follows.

1 Select a record.

2 Right-click > Undeploy.

A confirmation message is displayed.

3 Click Undeploy.

The Job Viewer displays the steps taken and the results.

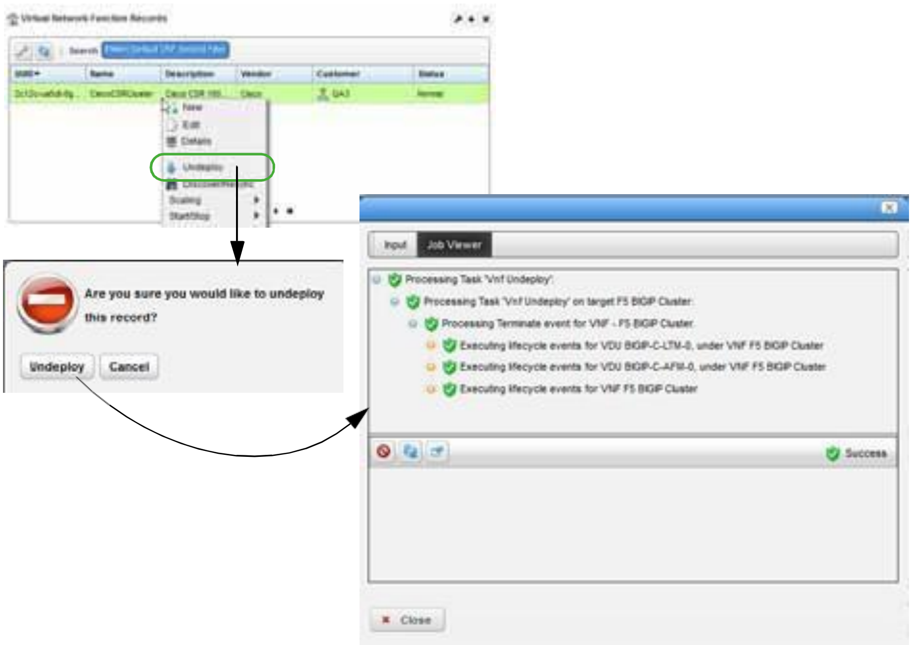

Are you sure you would like to undeploy this record?

Processing Task 'Vnf Undeploy'
Processing Task 'Vnf Undeploy' on target F5 BIGIP Cluster
Processing Terminate event for VNF - F5 BIGIP Cluster
Executing lifecycle events for VDU BIGIP-C-LTM-0, under VNF F5 BIGIP Cluster
Executing lifecycle events for VDU BIGIP-C-AFM-0, under VNF F5 BIGIP Cluster
Executing lifecycle events for VNF F5 BIGIP Cluster

Success

a. Click Close.

b. Verify the record’s status is now Offline.

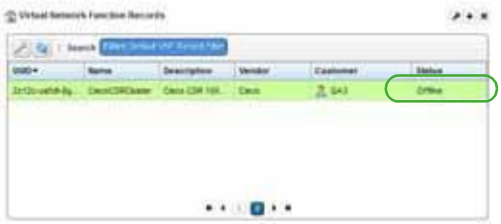

A

Ports Used

The following reference material is provided in this section:

StandardPortAssignments

Initial installation scans the following ports, and reports any conflicts with them:

Database: 3306 or user-configured database host, if using MySQL server.

Application server: 8089, 8162, 8489 [HTTPS], 8082

Web Portal: 8080, 8443 [HTTPS]

SNMP: 161, 162

Syslog: 514

When installation encounters a conflict with any of the above ports, a panel appears displaying a warning and the port[s] in conflict. You can then elect to continue since you can change the application ports after installation. If installation encounters no port conflicts, then no panel appears.

NOTE:

The installation scans TCP ports to detect potential conflicts. It does not scan UDP port conflicts including SNMP Ports 161 and 162. No SNMP or other applications should bind to UDP ports 161 and 162 since such bindings interfere with the application. If this conflict exists, the following error appears (with others):

```
FATAL ERROR - Initializing SNMP TrapListener
```

You may also configure network ports’ availability on firewalls. Sometimes, excluding applications from firewall interference is all that is needed (see Ports and Application To Exclude from Firewall on page 990). If you have remote mediation servers, see Remote Mediation Ports on page 985.

The following are some of the standard port assignments for installed components. These are often configurable, even for “standard” services like FTP or HTTP, with alterations to the files mentioned, so these are the default, typical or expected port numbers rather than guaranteed assignments. Also, see Protocol Flows on page 985 for more about network connections. The JBoss directory number may vary with your package’s version; *.* appears rather than actual numbers.

Destination Ports	Service	Files	Notes
3306	Database		... or user-configured database host, if using MySQL server.
8089, 8162, 8489 [HTTPS], 8082	Application server		
8080, 8443 [HTTPS]	Web Portal:		
HTTP/S (Web Client)			
8089[4]	oware.webservices.port	[user.root]\oware\lib\owweb services.properties	appserver. **Note:** this port was 80 in some previous versions.
8489[4, 5, 7]	org.apache.coyote.tomcat4.CoyoteConnector (Apache)	[user.root]\oware\jboss-\*.*\server\oware\deploy\jbossweb-tomcat41.sar\META-INF\ jboss-service.xml	app/medserver, jmx console, and web services, including Axis2
Other Ports			

Destination Ports	Service	Files	Notes
n/a[5](ICMP)	ping		MedSrv - > NtwkElement, NtwkElement - > MedSrv, ICMP ping for connection monitoring.
20[4, 5, 7] (TCP)	FTP Data Port	n/a Configurable in File Servers portlet editor	(Internally configurable), "MedSrv - > FTPSrv NtwkElement - > FTPSrv" medserver[1]
21[4, 5, 7] (TCP)	FTP Control Port	n/a	(Internally Configurable) "MedSrv - > FTPSrv NtwkElement - > FTPSrv" medserver[1]
22[4, 5, 7] (TCP)	SSH	n/a	MedSrv - > NtwkElement, secure craft access medserver[1]
23[4, 5, 7] (TCP)	Telnet	n/a	MedSrv - > NtwkElement, non-secure craft access medserver[1]
25 [4,5, 7] (TCP)	com.dorado.mbeans.OWEmailMBean (mail)	Configurable in the SMTP configuration editor in the Common Setup Tasks portlet.	AppSrv - > SmtpRelay, communication channel to email server from Appserver
69[4, 5, 7] (UDP)	TFTP	n/a	(Configurable internally), MedSrv - > TFTPSrv NtwkElement - > TFTPSrvmedserver [1]
161[4, 5, 7] (UDP)	com.dorado.media tion.snmp.request.listener.port (SNMP), oware.media tion.snmp.trap.forwarding.source.port	[user.root]\owareapps\ezmediation\lib\owmediation.properties	MedSrv - > NtwkElement, SNMP request listener and trap forwarding source medserver[1]
162[4, 5] (UDP)	oware.media tion.snmp.trap.forwarding.des tination.port (SNMP)	[user.root]\owareapps\ezmediation\lib\ezmediation.properties change this property: com.dorado.snmp.trap.listener.binding =0.0.0.0/162	NtwkElement - > MedSrv, SNMP trap forwarding destination port, medserver [1]
514[4, 5] (UDP)	com.dorado.mediation.syslog.port (syslog)	To change the syslog port, add com.dorado.mediation.syslog.port =[new port number] to owareapps\installprops\lib\installed.properties	NtwkElement - > MedSrv (mediation syslog port) medserver[1]
1098[4, 5, 7] (TCP)	org.jboss.naming.NamingService (JBOSS)	[user root]\oware\jboss-*.*\owareconf\jboss-root-service.xml	AppSrv - > MedSrv MedSrv - > AppSrv user client - >AppSrv user client - >MedSrv, (JBOSS naming service), app/medserver

Destination Ports	Service	Files	Notes
1099[4, 5, 7] (TCP)	org.jboss.naming.NamingService (JBOSS)	[user.root]\oware\jboss-*.*\owareconf\jboss-root-service.xml	MedSrv - > AppSrv, user client - > AppSrv, user client -> MedSrv, (JBOSS naming service & OWARE context server URL), app/medserver
1099[2, 4, 5, 7] (TCP)	OWARE.CONTEXT.SERVER.URL	[user.root]\oware apps\install props\lib\installed.properties [user.root]\oware apps\install props\medserver\lib\installed.properties	MedSrv - > AppSrv, user client - > AppSrv. user client - > MedSrv. (JBOSS naming service & OWARE context server URL) client medserver[1]
1100-1101	org.jboss.ha.jndi.HANamingService,	[user.root]/oware/jboss-*.*/server/all/deploy/cluster-service.xml	
1103[4, 5] (UDP)	jnp.reply.discoveryPort (JNP)	[user.root]\oware\lib\owappserver.properties	AppSrv - > MedSrv, AppSrv - > user client, (JNP reply discovery port), app/medserver
1123[4, 5] (UDP)	jnp.discoveryPort (JNP)	[user.root]\oware\lib\owappserver.properties	MedSrv - > AppSrv, user client - > AppSrv, (JNP discovery port), app/medserver
1521[4, 7] (TCP)	com.dorado.jdbc.database_name.oracle (JDBC)	[user.root]\oware apps\install props\lib\installed.properties	AppSrv - >OracleDBSrv, (JDBC database naming [Oracle]) database
1812[4, 7] (TCP)	RADIUS port (note, RADIUS is not supported in Cruz)	[user.root]\oware\jboss-*.*\server\oware\conf\login-config.xml	AppSrv - > RADIUS Srv, Appserver (RADIUS client login enabled– optional)
3306[4, 7] (TCP)	com.dorado.jdbc.database_name.mysql	[user.root]\oware apps\install props\lib\installed.properties	AppSrv - > MySQLSrv, (JDBC database naming [MySQL]) appserver)
3100[4, 5, 7] (TCP) 3200[4, 5, 7]	org.jboss.ha.jndi.HANaming Service (JBOSS)	[user.root]\oware\jboss-*.*\owareconf\cluster-service.xml	AppSrv - > AppSrv, user client - > AppSrv AppSrv - > MedSrv MedSrv - > AppSrv user client - > AppSrv user client - > MedSrv (JBOSS HA JNDI HA Naming service [1100 is stub] app/medserver

Destination Ports	Service	Files	Notes
3355[4] - application & mediation servers 8082 - portal	Direct access	Override application server port with this property: com.dorado.mediation.socket.relay.listen.port =3355	For both, the relay increments from the default until lit can bind to an open port.
4444	org.jboss.invocation.jrmp.server.JRMPInvoker	[user.root]/oware/jboss-*.*/server/all/conf/jboss-service.xml, RMIObjectPort, jboss:service = invoker,type =jrmp	
4445[4, 5, 7] (TCP)	org.jboss.invocation.pooled.server.PooledInvoker (JBOSS)	[user.root]\oware\jboss-*.*\owareconf\jboss–root-service.xml	AppSrv - > MedSrv MedSrv - > AppSrv user client - > AppSrv user client - > MedSrv, app/medserver
4446[4, 5, 7] (TCP)	org.jboss.invocation.jrmp.server.JRMPInvoker (JBOSS)	[user.root]\oware\jboss-*.*\owareconf\jboss–root-service.xml	(AppSrv -> AppSrv, AppSrv -> MedSrv, MedSrv - > AppSrv, user client - > AppSrv, user client - > MedSrv) app/medserver
5988, 5989	WBEM Daemon (5989 is the secure port) defaults		You can add ports and daemons in monitored services. These are only the default. WBEM requires one port, and only one, per daemon.
6500-10[4, 5, 7] (TCP)	JBOSS	Specify such connections in the ezmediation/lib/ezmediation.properties file.	user client - > MedSrv (user client to mediation server cut-through)
7800[2](TCP)	org.jboss.ha.framework.server.ClusterPartition (JBOSS)	[user.root]\oware\conf\cluster-service.xml	disabled - see UDP for same, (JBOSS HA frame work server cluster partition) TCP only
8009 (TCP)	org.mortbay.http.ajp.AJP13Listener	[user.root]\oware\jboss-*.*\server\oware\deploy\jbossweb-tomcat41.sar\META-INF\ jboss-service.xml	Obsolete — appserver
8083 (TCP)	org.jboss.web.WebService (JBOSS)	[user.root]\oware\jboss-*.*\owareconf\jboss–root-service.xml	Used by JBoss web service, appserver
8093[4, 5. 7] (TCP)	org.jboss.mq.il.uil2.UILServerILService	[user.root]\oware\jboss-*.*\owareconf\uil2-service.xml	MedSrv - > AppSrv, user client - > AppSrv (JBOSS mq iluil2 UIL Server-IL Server), app/medserver (Jboss JMS)
8443[2,4, 5, 7]	org.apache.coyote.tomcat4.CoyoteConnector	[user.root]\oware\jboss-*.*\server\oware\deploy\jbossweb.sar\META-INF\ jboss-service.xml	user client - > AppSrv (Apache Coyote Tomcat4 Coyote connector), appserver. This is the default HTTPS port for the web portal.

Destination Ports	Service	Files	Notes
9001[4, 6, 7] (UDP)	mediation.listener.mul ti cast.intercomm.port	[user.root]\lib\owmediation listeners.properties	MedSrv < - > MedSrv (mediation listener multicast intercommunications port) medserver[3]
9996, 6343 (UDP)	Traffic Flow Analysis	trafficanalyzer.ocp	You must configure the router to send flow reports to the Cruz server on UDP port 9996 for Netflow/jFlow and 6343 for sflow by default.
31310[4, 6, 7] (TCP)	JBoss		AppSrv - > AppSrv
45566[4, 5] (UDP)	org.jboss.ha.frame work.server.ClusterPart ition	[user.root]\jboss-*.*\owareconf \cluster-service.xml	AppSrv - > Multicast, (JBoss HA frame work server cluster partition), UDP only
54027[4,7]	Process Monitor	[user.root]\oware\lib\pmstar tup.dat	mgmt client - > AppSrv, mgmt client - > MedSrv (process monitor local client for server stop/start/status) app/medserver

[1] Remote mediation servers or application servers behaving as though they were mediation servers (single host installation).

[2] Unused in standardconfiguration.

[3] Client does not connect to medserver on this port.

[4] This port isconfigurable.

[5] Firewall Impacting

[6] The most likely deployment scenarios will have all servers co-resident at the same physical location; as such, communications will not traverse through a firewall

[7] Bidirectional

To operate through a firewall, you may need to override default port assignments.

Ifyoucluster your installation, you must disable multicast forcommunication through firewalls(to mediation servers or clients). Refer to the *Cruz Installation Guide* for more information.

NOTE:

To configure ports, open their file in a text editor and search for the default port number. Edit that, save the file and restart theapplication server andclient. Makesure youchangeports onallaffected machines.

Note that mediation service also establishes a socket connection to client on ports 6500 to 6510 for cut through. Specify such port connections in the `ezmediation/lib/ezmediation.properties file`. (As always, best practice is to override when specifying properties.)

Finding Port Conflicts

You can find ports in use with the following command line:

```
netstat -a -b -o | findstr [port number]
```

Use this command to track down port conflicts if, for example, installation reports one. Best practice is to run Cruz on its own machine to avoid such conflicts.

Remote Mediation Ports

Youmustopenthefollowingportsbetweenapplicationserversandremote mediationservers: 8443, 3306, 3200, 7800, 8009, 8080, 9001, 31310, 45566.

Protocol Flows

The following network protocol flows represent the application's interactions with Network Devices (for example: Dell Powerconnect switches). The (N) in these lines identifies dynamic port assignments. Often, Cruz establishes several comm unication flows to a spec ified static port so N can represent several dynamic ports. This list also outlines alternative flows for JBoss JMS activation.

NOTE:

This does not identify time service flows like ntp that can manage the time on the servers.

The following were changes to a standard installation done for the sake of measuring the protocol flows. In the J2EE Naming Service: the RMIPort was changed to 31310. Also, `owappserver.properties` (turns off mediation v2 servic es on application server) was changed: mediation true->false. This essentially disables mediation on the application server.

The following is the installation that produced the listed protocol flows: Full Application Server Installation, Custom Mediation Installation, toggling off 2 (MySQL) and 5 (App Server). The client was a simple client installation.

Application Server to Mediation Server Flows

J2EE

TCP Med Svr (N) - > App Svr (1098)

TCP Med Svr (N) <-> App Svr (1098)

TCP Med Svr (N) - > App Svr (1099)

TCP Med Svr (N) <-> App Svr (1099)

TCP Med Svr (N) - > App Svr (4446)

TCP Med Svr (N) <-> App Svr (4446)

TCP Med Svr (N) - > App Svr (4445)

TCP Med Svr (N) <-> App Svr (4445)

JBoss JMS enabled:

TCP Med Svr (N) - > App Svr (8093)

TCP Med Svr (N) < - > App Svr (8093)

Application Server to Application Server in Application Server Cluster

IGMPApp Svr-A/B-> 230.13.13.13 (Multicastaddressassigned perapplication cluster)

UDP App Svr-A/B (45566) -> 230.13.13.13 (45566)

IGMP App Svr-A/B - > 230.0.0.253

UDP App Svr-A/B (1123) - > 230.0.0.253 (1123)

UDP App Svr-A/B (1103) - > 230.0.0.5 (1103)

TCP App Svr-A (1100) <- App Svr-B (N)

TCP App Svr-A (1100) <-> App Svr-B (N)

TCP App Svr-A (31310) <- App Svr-B (N)

TCP App Svr-A (31310) <-> App Svr-B (N)

(Dynamic Port statically defined to be 31310 in the `cluster-service.xml` property file)

TCP App Svr-A (4446) <- App Svr-B (N)

TCP App Svr-A (4446) <-> App Svr-B (N)

TCP App Svr-A (2507) <- App Svr-B (N)

TCP App Svr-A (2507) <-> App Svr-B (N)

TCP App Svr-A (2508) <- App Svr-B (N)

TCP App Svr-A (2508) <-> App Svr-B (N)

TCP App Svr-B (N) -> App Svr-A (8080)

TCP App Svr-B (N) <-> App Svr-A (8080)

Application Server to Oracle Database Server

Optionally configured

TCP App Svr (N) -> Oracle DB Svr (1521)

TCP App Svr (N) <-> Oracle DB Svr (1521)

Application Server to MySQL Database Server

Embedded Database

TCP App Svr (N) - > MySQL Svr (3306)

TCP App Svr (N) <-> MySQL Svr (3306)

Mediation Server to Application Server Flows

J2EE

TCP App Svr (N) - > Med Server (4446)

TCP App Svr (N) <-> Med Server (4446)

TCP App Svr (N) - > Med Server (4445)

TCP App Svr (N) <-> Med Server (4445)

TCP App Svr (N) - > Med Server (1098)

TCP App Svr (N) <-> Med Server (1098)

Mediation Server uses 230.0.0.223:1123 to discover the application server cluster.

UDP Med Server (N) -> 230.0.0.223 (1123) (This multicast address is configurable)

UDP Med Svr (1123) - > 230.0.0.253 (1123)

Mediation Server to Mediation Server Flows

Mediation Server to Mediation Server cluster pair flows use the same ports to communicate between each other as those in the section Application Server to Application Server in Application Server Cluster on page 986. In addition, HA Trap processing use the following configurable multicast flow:

IGMP Med Svr-A/B - > 226.0.0.226

UDP Med Svr-A (9001) -> Med Svr-B (9001)

UDP Med Svr-A (9001) <-> Med Svr-B (9001)

Mediation Server to Network Element Flows

Telnet

TCP Med Server (N) - > Network Element (23)

TCP Med Server (N) < - > Network Element (23)

SSHv1/SSHv2

TCP Med Server (N) - > Network Element (22)

TCP Med Server (N) < - > Network Element (22)

FTP (mediation server FTPs files to and from the FTP server)

TCP Med Server (N) - > FTP/TFTP Svr (21)

TCP Med Server (N) < - > FTP/TFTP Svr (21)

TCP Med Server (N) <- FTP/TFTP Svr (20)

TCP Med Server (N) <- > FTP/TFTP Svr (20)

TFTP

Not applicable

SNMP

UDP Med Server (162) < - Network Element (N) (trap receipt)

UDP Med Server (N) - > Network Element (161) (get/set)

UDP Med Server (N) < - Network Element (161)

ICMP

No ports are involved with ICMP, but you must a llow ICMP traffic from the application/ mediation server and devices (and back).

Syslog

UDP Med Server (514) <- Network Element (514) (syslog messages) (TCP is possible but not implemented)

Mediation Server to FTP/TFTP Server

TCP (N) -> FTP/TFTP Svr (21) (ftp-control)

TCP (N) < - > FTP/TFTP Svr (21)

TCP (N) <- FTP/TFTP Svr (20) (ftp-data)

TCP (N) < - > FTP/TFTP Svr (20)

TCP Med Svr (N) -> FTP/TFTP Svr (69) (Testing “File Server”)

TCP Med Svr (N) < - > FTP/TFTP Svr (M)

Mediation Server to Trap Forwarding Destination

IP Trap Forwarding

Network Element (161) - > Trap Forwarding Receiver (statically defined N ex. 162 UDP)

NOTE:

The forwarded trap actually has the IP address of the Network Element, not the Med Server.

Network Element to FTP/TFTP Server

FTP

Network Element (N) - > FTP/TFTP Svr (21)

Network Element (N) <- > FTP/TFTP Svr (21)

Network Element (N) < - FTP/TFTP Svr (20)

Network Element (N) <-> FTP/TFTP Svr (20)

Network Element (N) -> FTP/TFTP Svr (69)

Network Element (N) <-> FTP/TFTP Svr (M)

Devices should have connectivity to the external FTP/TFTP server. *M* means we recommend installing external file serverson mediation serversforaperformance improvement. Youcan also use the internal FTP/TFTP server in Windows environments.

Client to Application Server

J2EE

TCP RC clt (N) - > App Svr (1099)

TCP RC clt (N) <-> App Svr(1099)

TCP RC clt (N) - > App Svr (1098)

TCP RC clt (N) <-> App Svr(1098)

TCP RC clt (N) - > App Svr (4446)

TCP RC clt (N) <-> App Svr(4446)

TCP RC clt (N) - > App Svr (4445)

TCP RC clt (N) <-> App Svr(4445)

IGMP RC clt(N) - > 230.0.0.5

UDP RC clt(1123) -> 230.0.0.253 (1123)

UDP RC clt(1103) <- AppSvr(1103)

TCP RC clt (N) - > App Svr (1100)

TCP RC clt (N) <-> App Svr (1100)

JBoss JMS enabled:

TCP RC clt (N) - > App Svr (8093)

TCP RC clt (N) < - > App Svr (8093)

Client to Mediation Server (Direct Access, or Cut thru)

Telnet/SSHv1/SSHv2 Cut - through

RC clt (N) - > Med Svr (1099)

RC clt (N) <- > Med Svr (1099)

RC clt (N) - > Med Svr (1098)

RC clt (N) <-> Med Svr (1098)

RC clt (N) - > Med Svr (4446)

RC clt (N) <-> Med Svr (4446)

RC clt (N) - > Med Svr (4445)

RC clt (N) <-> Med Svr (4445)

RC clt (6500) <- Med Svr (N) (6500 represents ports 6500-6510)

RC clt (6500) <- > Med Svr (N)

Email Network Element Config Differences

If email from the application server is turned on, then the following port must be opened between the application and email server:

TCP App Svr (N) - > smtp relay (25)

TCP App Svr (N) < - > smtp relay (25)

JBoss Management Access

The J2EE server has port 8080 open to allow web browsers access to the JBoss Management console. If you want to access this capability then the system browsing the jmx console must have access.

Mgmt client (N) - > App Server (8080)

To access the Mediation Servers:

Mgmt client (N) - > Med Server (8080)

Ports and Application To Exclude from Firewall

Exclude `java.exe`, tcp port 21 and udp port 69 from firew all interference to let the application function. The java process to exclude from firewall blocking is `<Installdir>\oware3rd\jdk[version number]\jre\bin\java.exe`.

If you have distributed the database functions then you must allow the database process to communicate with your machine through your firewall as well. The embedded database process is `mysqld-max-nt.exe` (in Windows, the path is `<installdir>oware3rd\mysql\[version number]\bin\mysql-max-nt.exe. Consult your DBA for Oracle processes, if applicable.`

Firewall Configuration

Example Linux firewall configuration (from `iptables-save > my-config-file`):

```
    -A INPUT -p tcp -m state --state NEW -m tcp --dport 21 -j ACCEPT
```

```
-A INPUT -p tcp -m state --state NEW -m tcp --dport 22 -j ACCEPT
-A INPUT -p udp -m state --state NEW -m udp --dport 69 -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 161 -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 162 -j ACCEPT
-A INPUT -p udp -m state --state NEW -m udp --dport 162 -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 514 -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 1099  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 1100  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 1101  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 3306  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 8089  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 5900  -j ACCEPT
-A INPUT -p udp -m state --state NEW -m udp --dport 5900  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 6343  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 8080  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 8089  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 8082  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 8083  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 8119  -j ACCEPT
-A INPUT -p udp -m state --state NEW -m udp --dport 8162  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 8162  -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 9996  -j ACCEPT
-A INPUT -p udp -m state --state NEW -m udp --dport 9996  -j ACCEPT
```

Add any newlines to the firewallfile `/etc/sysconfig/iptables`, and restart thefirewall service.

Installed Third Party Applications

This software includes the following applications. License information follows in parenthesis:

ant (licensed under GNU Lesser General Public License[LGPL] http://www.gnu.org/licenses/)

cygwin (LGPL)

expect (Public Domain)

J Free Charts (LGPL)

Jasper Reports (LGPL)

JBoss (see directory name for version) (LGPL)

JDK (Open Source)

JLoox — (Commercially Licensed to th is vendor, not re-distributable)

MySQL — (GNU Public License [GPL])

Open SSH — includes OpenSSL (LGPL)

OpenLDAP — (OpenLDAP License: http://www.openldap.org/software/release/license.html)

Perl — (LGPL)

TCL — (LGPL)

Tomcat — (Apache License**)**

Liferay Portal—Liferay Community Edition (LGPL)

The LGPL, GPL and Apache licenses let us redistribute the above listed open source components, but the EULA for all Cruz products prohibits redistribution of any package or component of the software. Consult the product’s EULA.htm file, typically in the `InstData` directory of the installation source for more information.

Windows Management Instrumentation Ports

Windows Management Instrumentation uses the following ports:

Protocol or Function	Ports Used
RPC, TCP	135,139,445,593
SNMP, UDP	161,162
Optional:	
WINS, TCP	42
UDP	42, 137
PrintSpooler, TCP	139, 445
TCP/IP PrintServer, TCP	515

These are relevant only if you are using any Windows-based server device driver.

B

Cruz NFV Permissions

This section describes the permissions for NFV records, descriptors, and operations. Each list includes a description of its type or function, valid actions, and portlet from which each action is performed.

NFV Records

Record types(Table B-1) areobject-scoped, where READ, WRITE, ADD, DELETE, and EXECUTE are against objects of the specified type.

NOTE:

By default, all users have permission to create and maintain these NFV records.

Table B-1. Record Types

Type	Description	Actions
Software Image	A catalog of VDU/VM images. Maintain and deploy/undeploy these records from the Software Images Portlet .	Read, Write, Add, Delete, Execute
Vim Instance	Virtual Information Manager OpenStack deployments under management. Maintain these records from the Virtualized Infrastructure Managers Portlet .	R,W,A,D,E
Network Service	Provisioned virtual network functions and physical functions used to produce connectivity. Maintain these records from the Network Service Records Portlet .	R,W,A,D,E
Vnf	At least one VDU/VM, such as a router. Maintain these records from the Virtual Network Function Records Portlet .	R,W,A,D,E
Pnf	Physical device to which you connect a virtual device. Maintain these records from the Physical Network Function Records Portlet .	R,W,A,D,E
Vdu Reservation	Virtualization Deployment Unit (VDU) reserved within a staged VNF. View VDU reservation records from the Virtual Reservations Portlet .	R,W,A,D,E

NFV Descriptors

Descriptor types (Table B-2) are object-scoped, where READ, WRITE, ADD, DELETE, and EXECUTE are against objects of the specified type.

NOTE:

By default, only users with administrative permissions (such as system administrators, network service designers or engineers, and so on) can add and maintain the NFV descriptors.

Table B-2. Descriptor Types

Type	Description	Actions
Network Service	The template used to create network service records that provision virtual network functions and physical functions used to produce connectivity. Maintain these descriptors from the Network Service Descriptors Portlet .	Read, Write, Add, Delete, Execute
Vnf	The template used to create VDU/VM records, such as a router. Maintain these descriptors from the Virtual Network Function Descriptors Portlet .	R,W,A,D,E
Pnf	The template used to create records for physical device to which you connect a virtual device. Maintain these descriptors from the Physical Network Function Descriptors Portlet .	R,W,A,D,E
Monitor Attributes	User-defined monitoring attribute descriptors where you define: • Scope (Network Service, Vdu, or Vnf) • Units of measure • Data Type (Boolean, Double, Float, Integer, IP, Long, or String) • Metric Type (Boolean, Counter, Gauge, or Says Up Time) Maintain these descriptors from the NFV Monitoring Attributes Portlet .	R,W,A,D,E

NFV Operations

Operation functions (T able B-3) for which there is only an EXECUTE action defined.

The listed functions are performed from the applicationportlets.

NOTE:

Bydefault, all users have permission to execute all these functions except for the Vnf Discovery function. Only users with administrative permissions (such as system administrators, network service designers or engineers, and so on) can execute the Vnfd Discovery function.

Table B-3. Operation Functions (Sheet 1 of 4)

Function	Description	Actions
Vim Capacity Resync	Refreshes namespace resources and capacity for the selected VIM or OpenStack capacity. Performed from the Virtualized Infrastructure Managers Portlet .	Execute
Image Download	Creates a software image descriptor (such as F5 Firefly) for a snapshot image and saves it to disk in the specified location. Performed from the Software Images Portlet . Right-click and software image and then select Download.	Execute
Image Deploy	Implements the selected software image (such as F5 Firefly) in the specified VIM. Performed from the Software Images Portlet . Right-click an software image and then select Deploy.	Execute
Image Undeploy	Removes the selected software image setup from the VIM. Performed from the Software Images Portlet . Right-click an software image and then select Undeploy.	Execute
Network Service Deploy	Implements a network service in the specified VIM by establishing network connectivity, processing the VNF instantiate event, and resyncing VIM resources. Performed from the Network Service Records Portlet . Right-click a network service record and then select Deploy.	Execute
Network Service Undeploy	Removes the network service from the specified VIM by processing the VNF terminate event and resyncing VIM resources. Performed from the Network Service Records Portlet . Right-click a network service record and then select Undeploy.	Execute
Network Service Discovery	Brings existing Network Services under management by interrogating the virtualization platform and matching the instance properties to a descriptor. Performed from the Network Service Records Portlet . Right-click a network service record and then select Discovery/Resync.	Execute

Table B-3. Operation Functions (Sheet 2 of 4)

Function	Description	Actions
Network Service Scale-In	Runs the Scale-in event to reduce the service scope. To successfully run, this option requires a scale-in event with a lower scale value. Performed from the Network Service Records Portlet . Right-click a network service record and then select Scaling > Scale In.	Execute
Network Service Scale-Out	Runs the Scale-out event to increase the service scope. To successfully run, this option requires a scale-out event with a higher scale value. Performed from the Network Service Records Portlet . Right-click a network service record and then select Scaling > Scale Out.	Execute
Network Service Diagnostics	Runs the Diagnostics event for the selected service. The user defines this event. Performed from the Network Service Records Portlet . Right-click a network service record and then select Maintenance > Diagnostics.	Execute
Network Service Upgrade	Runs the Upgrade event for the selected service to ensure that the latest descriptor is used. Performed from the Network Service Records Portlet . Right-click a network service record and then select Maintenance > Upgrade.	Execute
Network Service Recovery	Runs the Disaster Recovery event for the selected service. The user defines this event. Performed from the Network Service Records Portlet . Right-click a network service record and then select Maintenance > Disaster Recover.	Execute
Network Service Shutdown	Gracefully shuts down the selected network service. Performed from the Network Service Records Portlet . Right-click a network service record and then select Shutdown.	Execute
Network Service Link Deploy	Deploys link configuration actions that define the end points to pass traffic. Performed from the Network Service Records Portlet . Right-click a network service record and then select Deploy Virtual Links.	Execute
Vnfd Discovery	Creates a base descriptor record from an existing VNF deployment. Performed from the Virtual Network Function Descriptors Portlet. Right-click a VNF descriptor and then select Discover Vnfd. **Note:** Only users with administrative permissions (such as system administrators, network service designers or engineers, and so on) can execute this function.	Execute

Table B-3. Operation Functions (Sheet 3 of 4)

Function	Description	Actions
Vnf Discovery	Brings existing VNF instances under management by matching the configuration properties with those of a descriptor in inventory. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Discover/Resync.	Execute
Vnf Deploy	Implements a VNF instance in the specified VIM. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Deploy.	Execute
Vnf Undeploy	Runs the VNF terminate event. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Undeploy.	Execute
Vnf Start	Resumes running the selected VNF instance in a non-frozen state. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Start/Stop > Start. Can also use the OpenStack Unpause command.	Execute
Vnf Stop	Pauses the selected VNF instance, stores the VM state in RAM, and the instance continues to run in a frozen state. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Start/Stop > Stop. Can also use the OpenStack Pause command.	Execute
Vnf Resume	Continues with the selected VNF instance maintenance state. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Resume/Suspend > Resume.	Execute
Vnf Suspend	Pauses the selected VNF instance maintenance state. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Resume/ Suspend > Suspend.	Execute
Vnf Shutdown	Gracefully shuts down the selected VNF instance. Part of the lifecycle event. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Shutdown. Can also use the OpenStack Shut Off Instance action.	Execute
Vnf Migrate	Move the VNF from one host to another. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Maintenance > Migrate.	Execute
Vnf Scale-In	Runs the Scale-in event to reduce the VNF scope. To successfully run, this option requires a scale-in event with a lower scale value. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Scaling > Scale In.	Execute
Vnf Scale-Out	Runs the Scale-out event to increases the VNF scope. To successfully run, this option requires a scale-out event with a higher scale value. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Scaling > Scale Out.	Execute

Table B-3. Operation Functions (Sheet 4 of 4)

Function	Description	Actions
Vnf Diagnostics	Runs the Diagnostics event for the selected VNF. The user defines this event. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Maintenance > Diagnostics.	Execute
Vnf Heal	Runs the Heal event for the selected VNF. Part of the lifecycle event. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Maintenance > Heal.	Execute
Vnf Upgrade	Runs the Upgrade event for the selected VNF to ensure that the latest descriptor is used. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Maintenance > Upgrade.	Execute
Vnf Recovery	Runs the Disaster Recovery event for the selected VNF. The user defines this event. Performed from the Virtual Network Function Records Portlet . Right-click a VNF and then select Maintenance > Disaster Recovery.	Execute
Vnf Configure	Lays the down base configuration necessary for VNF to operate, once the VNF is instantiated.	Execute

C

Notifications

Provides a list of notifications (Table C-1) that the Cruz application publishes to the Cruz event bus. Refer to the *Cruz Event Management User Guide* to understand what you can do with the notification's w/r/t forwarding, suppressing, running actions, correlation, and so on.

Table C-1. Notifications (Sheet 1 of 4)

Notification	OID
Network Service Descriptors	
Network Service Descriptor Created	1.3.6.1.4.1.3477.36.2.2
Network Service Descriptor Branched	1.3.6.1.4.1.3477.36.2.3
Network Service Descriptor Updated	1.3.6.1.4.1.3477.36.2.4
Network Service Descriptor Deleted	1.3.6.1.4.1.3477.36.2.5
Network Service Records	
Network Service Record Staged	1.3.6.1.4.1.3477.36.4.2
Network Service Record Deploy Initiated	1.3.6.1.4.1.3477.36.4.3
Network Service Record Deploy Completed	1.3.6.1.4.1.3477.36.4.4
Network Service Record Deploy Failed	1.3.6.1.4.1.3477.36.4.5
Network Service Record Undeploy Initiated	1.3.6.1.4.1.3477.36.4.6
Network Service Record Undeploy Completed	1.3.6.1.4.1.3477.36.4.7
Network Service Record Undeploy Failed	1.3.6.1.4.1.3477.36.4.8
Network Service Record Scaling Initiated	1.3.6.1.4.1.3477.36.4.9
Network Service Record Scaling Complete	1.3.6.1.4.1.3477.36.4.10
Network Service Record Scaling Failed	1.3.6.1.4.1.3477.36.4.11
Network Service Record Deleted	1.3.6.1.4.1.3477.36.4.12
VNF Descriptors	
Vnf Descriptor Created	1.3.6.1.4.1.3477.36.6.2
Vnf Descriptor Branched	1.3.6.1.4.1.3477.36.6.3
Vnf Descriptor Updated	1.3.6.1.4.1.3477.36.6.4
Vnf Descriptor Deleted	1.3.6.1.4.1.3477.36.6.5
VNF Records	
Vnf Record Staged	1.3.6.1.4.1.3477.36.8.2
Vnf Record Deploy Initiated	1.3.6.1.4.1.3477.36.8.3
Vnf Record Deploy Completed	1.3.6.1.4.1.3477.36.8.4
Vnf Record Deploy Failed	1.3.6.1.4.1.3477.36.8.5
Vnf Record Undeploy Initiated	1.3.6.1.4.1.3477.36.8.6
Vnf Record Undeploy Completed	1.3.6.1.4.1.3477.36.8.7
Vnf Record Undeploy Failed	1.3.6.1.4.1.3477.36.8.8
Vnf Record Scaling Initiated	1.3.6.1.4.1.3477.36.8.9
Vnf Record Scaling Complete	1.3.6.1.4.1.3477.36.8.10
Vnf Record Scaling Failed	1.3.6.1.4.1.3477.36.8.11
Vnf Record Deleted	1.3.6.1.4.1.3477.36.8.12
Vnf Record Offline	1.3.6.1.4.1.3477.36.8.13
Vnf Record Failed	1.3.6.1.4.1.3477.36.8.14
Vnf Record Suspended	1.3.6.1.4.1.3477.36.8.15
Vnf Record Resumed	1.3.6.1.4.1.3477.36.8.16
Vnf Record Stopped	1.3.6.1.4.1.3477.36.8.17
Vnf Record Started	1.3.6.1.4.1.3477.36.8.18

Table C-1. Notifications (Sheet 2 of 4)

Notification	OID
Vnf Record Migrated	1.3.6.1.4.1.3477.36.8.19

Table C-1. Notifications (Sheet 3 of 4)

Notification	OID
VNF Records (continued)	
Vnf Record Normal	1.3.6.1.4.1.3477.36.8.20
Vdu Record Staged	1.3.6.1.4.1.3477.36.10.2
Vdu Record Instantiation Initiated	1.3.6.1.4.1.3477.36.10.3
Vdu Record Instantiation Complete	1.3.6.1.4.1.3477.36.10.4
Vdu Record Instantiation Failed	1.3.6.1.4.1.3477.36.10.5
Vdu Record Virtual Machine Booted	1.3.6.1.4.1.3477.36.10.6
Vdu Record Resource Under Management	1.3.6.1.4.1.3477.36.10.7
Vdu Record Termination Initiated	1.3.6.1.4.1.3477.36.10.8
Vdu Record Termination Complete	1.3.6.1.4.1.3477.36.10.9
Vdu Record Termination Failed	1.3.6.1.4.1.3477.36.10.10
Vdu Record Offline	1.3.6.1.4.1.3477.36.10.11
Vdu Record Failed	1.3.6.1.4.1.3477.36.10.12
Vdu Record Suspended	1.3.6.1.4.1.3477.36.10.13
Vdu Record Resumed	1.3.6.1.4.1.3477.36.10.14
Vdu Record Stopped	1.3.6.1.4.1.3477.36.10.15
Vdu Record Started	1.3.6.1.4.1.3477.36.10.16
Vdu Record Migrated	1.3.6.1.4.1.3477.36.10.17
Vdu Record Rebooted	1.3.6.1.4.1.3477.36.10.18
Vdu Record Normal	1.3.6.1.4.1.3477.36.10.20
Lifecycle Records	
Lifecycle Vdu Record Task Initiated	1.3.6.1.4.1.3477.36.12.2
Lifecycle Vdu Record Task Complete	1.3.6.1.4.1.3477.36.12.3
Lifecycle Vdu Record Task Failed	1.3.6.1.4.1.3477.36.12.4
Lifecycle Vnf Record Task Initiated	1.3.6.1.4.1.3477.36.12.5
Lifecycle Vnf Record Task Complete	1.3.6.1.4.1.3477.36.12.6
Lifecycle Vnf Record Task Failed	1.3.6.1.4.1.3477.36.12.7
Lifecycle Service Record Task Initiated	1.3.6.1.4.1.3477.36.12.8
Lifecycle Service Record Task Complete	1.3.6.1.4.1.3477.36.12.9
Lifecycle Service Record Task Failed	1.3.6.1.4.1.3477.36.12.10
Lifecycle Link Record Task Initiated	1.3.6.1.4.1.3477.36.12.11
Lifecycle Link Record Task Complete	1.3.6.1.4.1.3477.36.12.12
Lifecycle Link Record Task Failed	1.3.6.1.4.1.3477.36.12.13
Virtual Link Records	
Virtual Link Record Deploy Initiated	1.3.6.1.4.1.3477.36.14.2
Virtual Link Record Deploy Complete	1.3.6.1.4.1.3477.36.14.3
Virtual Link Record Deploy Failed	1.3.6.1.4.1.3477.36.14.4
Virtual Link Record Undeploy Initiated	1.3.6.1.4.1.3477.36.14.5
Virtual Link Record Undeploy Complete	1.3.6.1.4.1.3477.36.14.6
Virtual Link Record Undeploy Failed	1.3.6.1.4.1.3477.36.14.7

Table C-1. Notifications (Sheet 4 of 4)

Notification	OID
VIM Instances	
VIM Instance Created	1.3.6.1.4.1.3477.36.16.2
VIM Instance Deleted	1.3.6.1.4.1.3477.36.16.3
Software Images	
Software Image Deployed	1.3.6.1.4.1.3477.36.18.2
Software Image Undeployed	1.3.6.1.4.1.3477.36.18.3
VNF Licenses	
VNF License Threshold Warning	1.3.6.1.4.1.3477.36.20.2
VNF License Threshold Cleared	1.3.6.1.4.1.3477.36.20.3

Glossary

Access Control Refers to mechanisms and policies that restrict access to computer resources. An access control list (ACL), for example, specifies what operations different users can perform on specific files and directories.

Advanced Firewall Manager (AFM)
An F5 Networks BIG-IP Web application firewall that uses both positive and negative security models to identify, isolate, and block sophisticated attacks without impacting legitimate application transactions. AFM is a high-performance, stateful, full-proxy network security solution designed to guard data centers against incoming threats that enter the network on the most widely deployed protocols.

Agent Mapper
The Mediation Agent handles SNMP requests, SNMP traps, and ASCII grammars through the Agent Mapper. You must configure the Agent Mapper with each Mediation Agent location.

Alarm
A signal alerting the user to an error or fault. Alarms are produced by events. Alarms produce a message within the Alarm Window.

API
Application Programing Interface—A set of routines used by the application to direct the performance of procedures by the computer's operating system.

Applet
A small application that performs a single task and runs within a widget engine or larger program, often as a plug-in.

Application Server
A component frameworks container. For example, Enterprise JavaBeans, that includes fundamental services for access and intra-system communication. Cruz runs in the Application Server.

ASCII
American Standard Code for Information Interchange protocol.

ATM
Asynchronous Transfer Mode. ATM is a high bandwidth, low-delay, connection-oriented, packet-like switching and multiplexing technique.

Authentication
The process of determining the identity of a user that is attempting to access a network. Authentication occurs through challenge/response, time-based code sequences or other techniques. See CHAP and PAP.

Bare Metal
A computer system or network in which a virtual machine is installed directly on hardware rather than within the host operating system (OS). Bare metal refers to a hard disk where the OS is installed.

Blade
See Solution Blade.

BOM
Business Object Manager; responsible for managing all defined objects. Includes object database. Handles all scheduling, notification, caching and modification services for objects.

Business Class
A category of objects. A business class contains a name that defines the class, properties, and attributes that define the individual data within the business class. For example, you can have a business class named Employee that has a persisted property and a Name attribute.

Note: A business class is not the same as a Java class. Business class is concept that is unique to Oware.

BusinessObject

One specific instance in a class. For example, the Employee business class wouldinclude the John Smith object.

Ceilometer

A metering function that collects, normalizes, and transforms data produced by OpenStack services. The information collected allows you to track and manage resources.

Center

One of three integrated Oware parts; each part has its own function within Oware. The three parts are:

- Oware Creation Center (OCC) — for creating and rendering solution blades
- Oware Execution Center (OEC) — for powering solution blades
- Oware Management Center (OMC) — for deploying and managing solution blades

Class

See Business Class.

CORBA

Common Object Request Broker Architecture — An architecture that enables pieces of programs, called objects, to communicate with one another regardless of what programming language they were written in or what operating system they're running on.

CoS

Class of Service—Describes the level of service provided to a user. Also provides a way of managing traffic in a network by grouping similar types of traffic.

Creation Center, (OCC)

This Oware center is where you create and render solution blades.

Database

An organized collection of Oware objects.

Deployment

The distribution of solution blades throughout the domain.

Descriptor

A file that describes the VNF attributes and network services specifications.

A VNF descriptor describes the virtual functions behavior. A network services descriptor describes the deployment requirements, operational behavior, and policies required.

These descriptors are templates from which you build network service records and VNF records.

Digital Certificate

A digital certificate is an electronic “credit card” that establishes your credentials when doing business or other transactions on the Web. It is issued by a certification authority (CA). It contains your name, a serial number, expiration dates, a copy of the certificate holder's public key (used for encrypting and decrypting messages and digital signatures), and the digital signature of the certificate-issuing authority so that a recipient can verify that the certificate is real.

Domain

A goal-oriented environment that can include an industry, company, or department. You can use Oware to create solutions within your domain.

EJB

Enterprise JavaBean — An application designed to extend the JavaBean component model to a cross-platform, server-side application.

Element

A type of data that defines and organizes Ow are data for easy references. There are six types of elements: business class, event, form, grammar, graphic, and rule.

Encryption

Scrambling data in such a way that it can only be unscrambled through the application of the correct cryptographic key.

Equipment

A network device managed by the system.

Ethernet Access Point

These represent the access points through which Ethernet frames flow in and out of an Ethernet service.

For an Ethernet Service that uses an Ethernet Trunk Service, an Ethernet Access Port must be associated with either one of the two Ethernet Access Services.

Ethernet Access Service

Since an Ethernet trunk can be shared by multiple Ethernet Services, each Ethernet Service relates to a shared trunk via a unique Ethernet Access component.

Because Ethernet trunk is a point-to-point connection, there are two Ethernet Access Services per trunk per Ethernet service instance.

EthernetService

An Ethernet service represents a virtual layer broadcast domain that transports or transmits Ethernet traffic entering from any one endpoint to all other endpoints.

Often, this is a VLAN service across multiple devices.

An Ethernet service may or may not use Ethernet trunk, depending on the desired connection between two neighboring devices. If the connection is exclusively used for this Ethernet service, no Ethernet trunk is need ed. On the other hand, if the connection is configured as an aggregation which can be shared by multiple Ethernet services, an Ethernet trunk models such a configuration.

Each Ethernet service can have multiple Ethernet Access Ports through which Ethernet traffic flows get access to the service.

EthernetTrunk

An Ethernet Trunk service represents a point-to-point connection between two ports of two devices. Ethernet frames transported by the connection are encapsulated according to IEEE 802.1Q protocol. Each tag ID value in 802.1Q encapsulated Ethernet frames distinguishes an Ethernet traffic flow. Thus, an Ethernet trunk can aggregate multiple Ethernet VLANs through a same connection which is why “trunk” describes these.

Ethernet Trunk Port

An Ethernet trunk port is a port that terminates a point-to-point Ethernet trunk. Since Ethernet trunk is a point-to-point connection, each Ethernet trunk contains two Ethernet trunk ports.

Event

Notification received from the NMS (Network Management Sy stem). Notifications may originate from the traps of network devices or may indicate an occurrence such as the closing of a form. Events have the potential of becoming alarms.

Event Definition

Parameters that define what an event does. For example, you can tell Oware that the event should be to wait for incoming data from a remote database, then have the Oware application perform a certain action after it receives the data.

Event Instance

A notification sent between two Oware components. An event instance is the action the event performs per the event definition.

Event Template

Defines how an event is going to be handled.

Event Threshold

Number of events within a given tomfooleries that must occur before an alarm is raised.

Execution Center (OEC)

This <Oware> center renders solutions by executing solution blades. It does this by completing the application. It loads classes into the classes database, it and reads and creates objects in the database.

Exporting

Saving business objects, packages, or solution blades to a file for others to import.

FCAPS

Fault Configuration Accounting Performance Security

Filter

In network security, a filter is a program or section of code that is designed to examine each input or output request for certain qualifying criteria and then process or forward it accordingly.

Flavors

Cruz deployment configurations that specify the number of network services, virtualized network functions, *etc.* The deployment flavors are useful when there is a need to scale in/out the number of instances to decrease/increase capacity, respectively.

In an OpenStack configuration, flavors define resources, such as CPU, memory, and disk space.

Form

An interface that lets users interact with a solution blade. Users can use forms to enter data, view data and/or events, and trigger rules and/or events in a solution blade.

Grammar

Translates ASCII text communications between the solution blade and the external systems via patterns built into the OCC Grammar Composer.

Hypervisor

Software that partitions the underlying physical resources and creates Virtual Machines (VMs) and isolates the VMs from each other.

A hypervisor runs one or more virtual machines on a host machine, and each virtual machine is called a guest machine. This allows multiple guest VMs to effectively share the system’s physical resources, such as CPU cycles, memory space, network bandwidth, and so on.

IDE

Integrated Development Environment. A programing environment integrated into an application.

IIOP

Internet Inter-ORB Protocol. A protocol used to implement CORBA solutions over the World Wide Web

Image

A blueprint from which new VMs are deployed. An image captures the operating system, applications, and their configuration settings. The image does not contain any sensitive data (such as private keys and passwords).

Importing

Loading business objects, packages, or solution blades into the OCC for you to edit.

Instantiate

The process of creating an object of a specific class.

ISATAP

The Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) is an IPv6 transition mechanism which is defined as a tunneling IPv6 interface and is meant to transmit IPv6 packets between dual-stack nodes on top of an IPv4 network.

Key

In cryptography, a key is a variable value that is applied using an algorithm to a string or block of unencrypted text to produce encrypted text. The length of the key generally determines how difficult it will be to decrypt the text in a given message.

Key Management

The establishment and enforcement of message encryption and authentication procedures, in order to provide privacy-enhanced mail (PEM) services for electronic mail transfer over the Internet.

LCE

A Lifecycle Event (LCE) that identifies a specific behavior as defined within the Cruz Management and Orchestration NFV specification.

Managed Object

A network device managed by the system.

Management Center, Oware (OMC)

This Oware center deploys and manages solution blades.

Mbeans

Management Beans. These objects specify the protocol of the device with which they are communicating.

Mediation

Communication between this application and external systems or devices, forexample, printers. Mediation services let this application treat these devices as objects.

Mediation Agent

Any communication to and from equipment is handled by the Mediation Agent. This communication includes SNMP requests, ASCII requests, and unsolicited ASCII messages. In addition, the Mediation Agent receives and translates emitted SNMP traps and converts them into events.

MEG

Maintenance Entity Group

MEP

Maintenance End Point

Metadata

Attributes that describe Oware database objects; data about data (for example, title, size, author, or subject).

Metadatabase

An organized collection of Oware attributes (also called meta-data). For example, the ZIP Code attribute in the meta-database can describe all database business objects with five digits (such as, 94101).

MIB

Management Information Base. A database (repository) of managed objects containing object characteristics and parameters that can be monitored by the network management system.

Multi-Threading

Concurrent processing of more than one message by an application.

Network Functions Virtualization Infrastructure (NFVI)

All hardware and software components used to create the environment in which you deploy virtualized network functions (VNFs).

Note: The NFV infrastructure can span across several locations (such as different data center operations) and provides the connectivity between these locations.

NFVI and VNF are the are the main entities of network function virtualization.

OAM

Operation, Administration and Maintenance

OCC

See Creation Center, (OCC).

OEC See Execution Center (OEC).

OID Object ID.

OMC See Management Center, Oware (OMC) .

Onboarding A process used to implement network services.

OpenStack An open-source cloud operating system that controls large pools of compute, storage, and networking resources. The OpenStack system is often deployed as an infrastructure-as-a-service (IaaS).

Orchestration The automated arrangement, coordination, and management of computer systems, middleware, and services.

OSPF Open Shortest Path First routing protocol.

Oware Oware is an Internet infrastructure software product that delivers advanced user, content and network aware Internet Commerce Systems. Oware allows users to quickly render Internet Commerce Systems that have inherent knowledge of user profiles, content, and network resources, enabling a class-of-service based Internet.

Package The container for Oware elements.

ParameterEvent An event with rule parameters.

PersistentData Data stored in the database for future use. (As opposed to transient data).

Pessimistic Locking A pessimistic locking scheme denies all access (create, read, write, update) to database items already being read by other users.

Policy A rule made up of conditions and actions and associated with a profile. Policy objects contain business rules for performing configuration changes in the network for controlling Quality of Service and Access to network resources. Policy can be extended to perform other configuration functions, including routing behavior, VLAN membership, and VPN security.

Policy Enforcement Points (PEP) In a policy enforced network, a policy enforcement point represents a security appliance used to protect one or more endpoints. PEPs are also point for monitoring the health and status of a network. PEPs are generally members of a policy group.

Policyrouting Routing scheme that forwards packets to specific interfaces based on user-configured policies. Such policies might specify that traffic sent from a network should be routed through interface, while all other traffic should be routed through another interface.

PolicyRules In a policy enforced network (PEN), policy rules determine how the members and endpoint groups of a policy group communicate.

PPTP (Point-to-Point Tunneling Protocol) Point-to-Point Tunneling Protocol (PPTP) is a network protocol that enables the secure transfer of data from a remote client to a private enterprise server by creating a virtual private network (VPN) across TCP/IP-based data networks. PPTP supports on-demand, multi-protocol, virtual private networking over public networks, such as the Internet.

PrivateKey
In cryptography, a private or secret key is an encryption/decryption key known only to the party or parties that exchange secret messages. In traditional secret key cryptography, a key would be shared by the communicators so that each could encrypt and decrypt messages. The risk in this system is that if either party loses the key or it is stolen, the system is broken. A more recent alternative is to use a combination of public and private keys. In this system, a public key is used together with a private key.

Profile
A profile is an abstract collection of configuration data that is utilized as a template to specify configuration parameters to be applied to a device as a result of a policy condition being true.

Pseudo-optimistic locking
Oware provides pseudo-optimistic locking. The first caller to attempt to get an object for update is granted an “update” lock. Subsequent callers can still read the object, but Oware denies any attempt to obtain an update lock until the initial caller ends the session.

Public Key
A public key is a value provided by some designated authority as a key that, combined with a private key derived from the public key, can be used to effectively encrypt and decrypt messages and digital signatures. The use of combined public and private keys is known as asymmetric encryption. A system for using public keys is called a public key infrastructure (PKI).

QoS
Quality of Service. In digital circuits, it is a measure of specific error conditions as compared with a standard. The establishment of QoS levels means that transmission rates, error rates, and other characteristic s can be measured, improved, and, to some extent, guaranteed in advance. Often related to Class of Service (CoS).

RADIUS
RADIUS (Remote Authentication Dial-In User Service) is a client/server protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service. RADIUS allows a company to maintain user profiles in a central database that all remote servers can share.

Rendering/Render
The process of generating OMG metadata, compiling and saving application components within the context of Oware. Also called solution rendering.

RIP
Routing Information Protocol

RMI
Remote Method Invocation — A set of protocols (Sun’s JavaSoft) that enables Java objects to communicate remotely with other Java objects. It works only with Java objects.

RMO
Resource Management and Operation (RMO) is the Cruz core resource management, configuration, backup, and image management feature.

RMON
Remote **Mon**itoring; a monitoring set of SNMP-bas ed MIBS used to manage networks remotely. Nine diagnostic groups are defined. Of these, Cruz software supports Events and Alarms.

Rule
Defines what a solution blade does in a situation — for example, how the solution blade responds to an event. An event, another rule, or program code can trigger a rule.

Scale In
An event used to decrease an instance’s capacity by selecting the appropriate deployment flavors.

Scale Out
An eventused to increase an instance’s capacity by selecting the appropriate deployment flavors.

Self-signed Certificate

A self-signed certificate uses its own certificate request as a signature rather than the signature of a CA. A self-signed certificate will not provide the same functionality as a CA-signed certificate. A self-signed certificate will not be automatically recognized by users' browsers, and a self-signed certificate does not provide any guarantee concerning the identity of the organization that is providing the website.

Servlet

A small Java program that runs on a server, noted for quickness. (See also: Applet.)

SMTP

Simple Mail Transfer Protocol.

SNMP

Simple Network Management Protocol. Network management protocol used almost exclusively in TCP/IP networks. SNMP provides the means to monitor and control network devices, and to manage configurations, statistics collection, performance, and security.

Solution Blade

The properties specified in the OCC that describe the application to be rendered. Also called a meta application or a blade.

Solution Rendering

The integration of all Oware elements and components into an executable application and a set of instructions (metadata) for creating the application.

Spanning Tree Protocol (STP)

The inactivation of links between networks so that information packets are channeled along one route and will not search endlessly for a destination.

Spanning Tree Protocol (STP)

The inactivation of links between networks so that information packets are channeled along one route and will not search endlessly for a destination.

SSH (Secure Shell)

A protocol which permits secure remote access over a network from one computer to another. SSH negotiates and establishes an encrypted connection between an SSH client and an SSH server.

SSL (Secure Sockets Layer)

A program layer created by Netscape for managing the security of message transmissions in a network. Netscape's idea is that the programming for keeping your messages confidential ought to be contained in a program layer between an application (such as your Web browser or HTTP) and the Internet's TCP/IP layers. The “sockets” part of the term refers to the sockets method of passing data back and forth between a client and a server program in a network or between program layers in the same computer.

Telemetry

An automated communications process used to collect measurements and other data from remote or inaccessible points and transmit it to receiving equipment for monitoring.

Transaction

An atomic unit of work that modifies data. A transaction encloses one or more program statements, all of which either complete or roll back. Transactions enable multiple users to access the same data concurrently.

Transaction Context

The scope of a transaction; defined by a transaction context shared by the participating objects.

Transaction Service

This defines interfaces that allow multiple, distributed objects to cooperate. These interfaces enable the objects to either commit all changes together or to rollback all changes together, even in the presence of failure. The Transaction Service typically places no constraints on the number of objects involved, the topology of the application or the way in which the application is distributed across a network.

Transient Data

Temporary data that exists only while being used by a solution blade.

Trap (SNMP Trap) A notification from a network element or device of its status, such as a server startup. This notification is sent by an SNMP agent to a Network Management System (NMS) where it is translated into an event by the Mediation Agent.

Trap Forwarding The process of re-emitting trap events to remo te hosts. Trap Forwarding is available from the application through Actions an d through the Resource Manager.

TrapForwarding The process of re-emitting trap events as tr aps to another management system, such as the NMS (Network Management System).

UUID (Universally Unique ID) A unique 128-bit, system generated value assigned to each network service, virtual network function, and physical function descriptor.

Virtual Rule Machine (VRM) A rule engine encapsulated by a ubiquitous meta schema that manages logic, processes requests, and associates behaviors required by a solution blade.

Virtual Solution Machine (VSM) A complete, integrated set of technologies used to create, execute, and manage Internet commerce business systems. All components, services and technology in a single product, already integrated and ready to use.

Virtualized Infrastructure Manager (VIM) A functional block that controls and manages Network Functions Virtualization Infrastructure (NFVI) resources (compute, storage, and network) within an operator’s infrastructure domain.

VLAN A virtual local area network (LAN), commonly known as a VLAN, is a group of hosts with a common set of requirements that communicate as if they were attached to the Broadcast domain, regardless of their physical location. A VLAN has the same attributes as a physical LAN, but it allows for end stations to be grouped together even if they are not located on the same network switch. Network reconfiguration can be done through software instead of physically relocating devices.

Index