HYSOLATE

The CISO's Guide to

Extending Zero Trust to the Endpoint

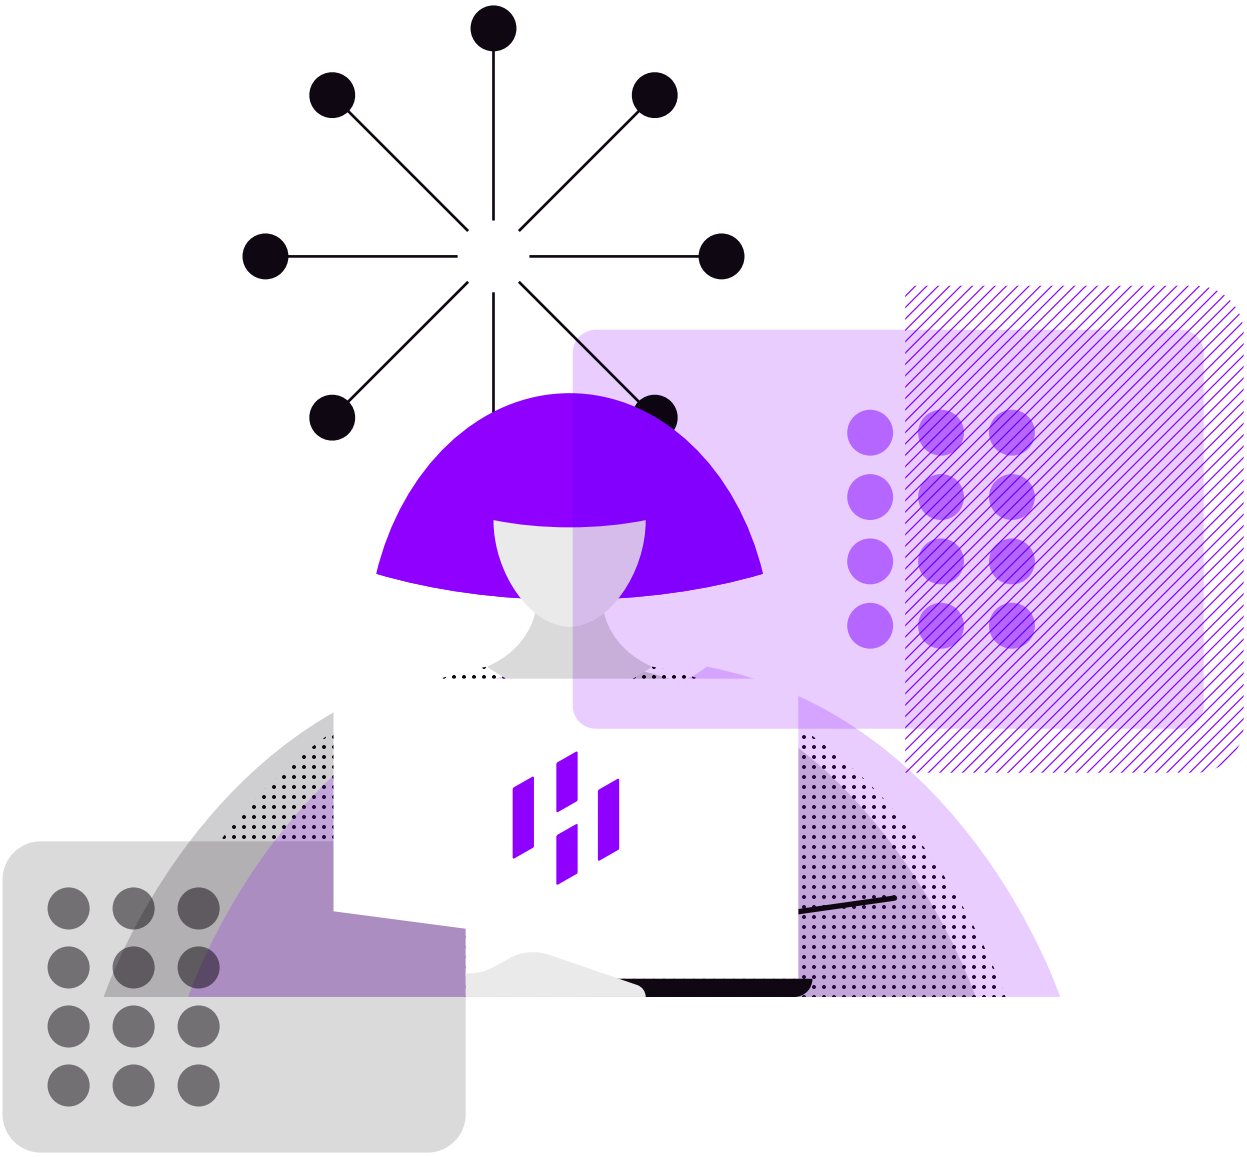

What Do We Mean By Zero Trust?

Zero Trust is a holistic approach to network security, which requires the verification of each person and device whenever it attempts to access resources on a private network. This remains true whether or not that device or person is already inside the network perimeter.

This resolves many issues in the traditional network security model, which relied on the concept of a ‘security perimeter’. Access to a network was tightly controlled, but once inside, connections were trusted by default and an attacker could cause significant damage. In today’s distributed environment, with data and applications running on remote cloud services, employees working from home or from personal devices, and the growing use of mobile and IoT, the security perimeter approach is no longer valid and is being replaced by the Zero Trust model.

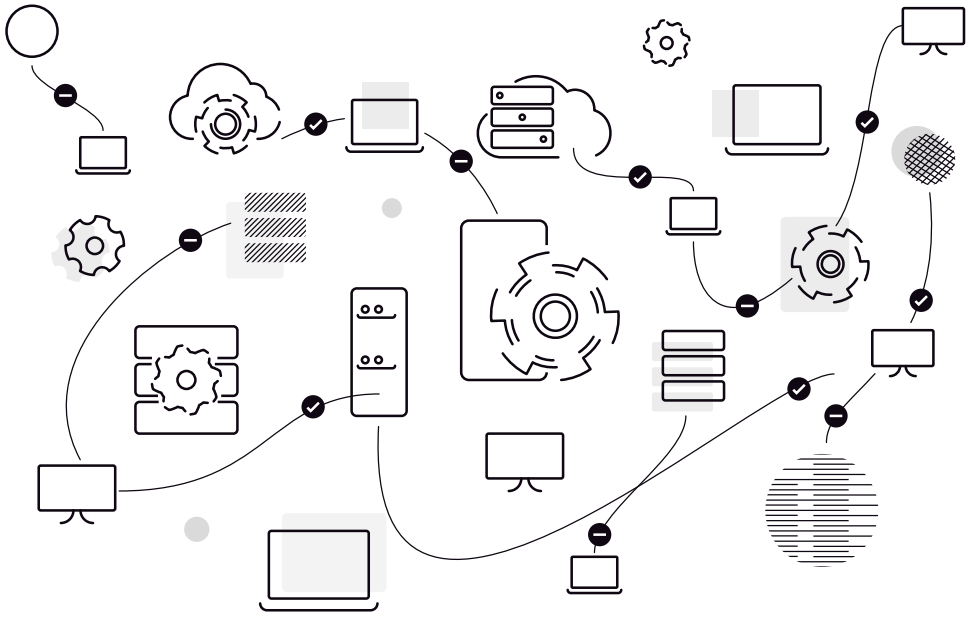

The Zero Trust model comprises a set of principles, and recommends the use of technologies and techniques in line with those principles. There are many technical and operational approaches to implementing Zero Trust.

Why is Zero Trust Important?

In recent years, it has become clear that data breaches are not only, or even primarily, caused by breaches of the network perimeter. Increasingly, breaches are caused by malicious or careless insiders, accounts compromised by social engineering or other techniques, or focus on weaker links of the IT environment, such as unsecured personal endpoints or cloud systems.

Before the advent of Zero Trust, companies used solutions like firewalls and VPNs to control access to networks and applications. The inherent flaw of these solutions is that once the user is successfully authenticated, they are “trusted” and granted unconditional access to corporate resources. Users were exposed to unnecessary data and systems, including mission-critical resources.

To resolve this situation, organizations implemented complex, expensive layers of security to stop attackers, such as intrusion detection, behavioral analytics and endpoint protection, with no real guarantee that any of these layers will prevent a breach.

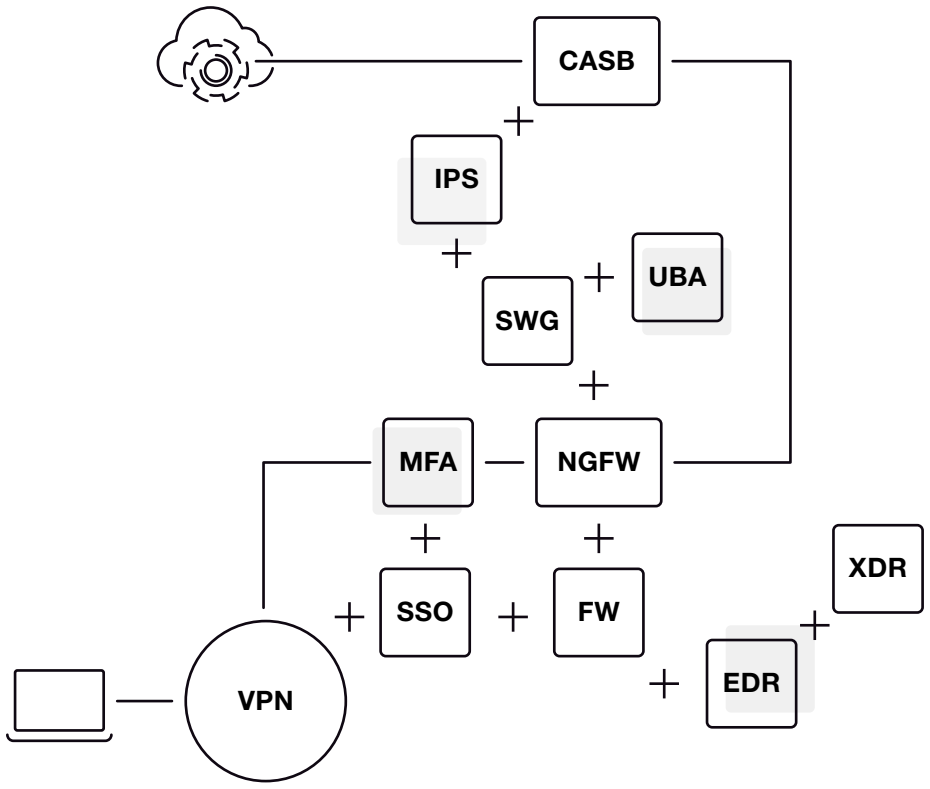

Zero Trust is a more holistic solution that assumes attackers have already breached the network, but prevents them from escalating privileges and moving laterally within the network. It reduces the need for complex security measures to detect and mitigate threats, because it creates an inherently secure network environment.

Another benefit of Zero Trust is that it centralizes and standardizes the problem of access control. Instead of requiring every application on the network to be inherently secure and implement strong authentication measures, the network manages access and authentication centrally. Applications do not handle authentication on their own, relying on a Zero Trust “access broker” to check if users are eligible for access, and verify their identity.

Core Principles of a Zero Trust Architecture

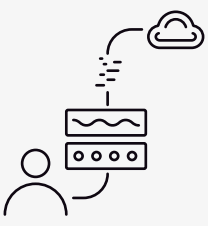

Strict Evaluation of Access Controls

The Zero Trust model assumes that potential attackers may exist inside and outside a network, therefore all users or devices attempting to access network resources must be authenticated, and each access request must be authorized and encrypted.

Preventative Techniques

To prevent breaches and minimize their damage, a variety of preventive techniques are available.

1/ Multi-factor Authentication

The most common method of confirming user identity. It requires the user to provide at least two forms of evidence to confirm credibility. These may include security questions, SMS or email confirmation, and/or logic-based exercises. The more means required for access, the better the network is secured.

2/ Limiting Access for Authenticated Users

Another layer used to gain trust. Each user or device only gains access to the minimal amount of resources required, thus minimizing the potential attack surface of the network at any time. All else remains blocked, thereby denying lateral movement for trusted entities.

3/ Micro-segmentation

A network security technique that involves separating networks into zones, each of which requires separate network access. The damage a hacker can do, even once security is breached, remains limited to the microsegment they have managed to penetrate.

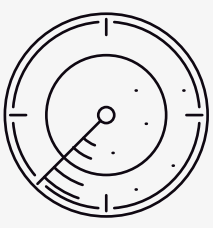

Real-Time Monitoring to Identify Malicious Activity

The Zero Trust model is mainly a preventative one. In addition to preventive measures, real-time monitoring is important, because it can minimize the time between an initial breach and the moment a threat spreads to additional systems on the network. Swift monitoring enables detection, investigation, and remediation, closing the window of opportunity for attackers.

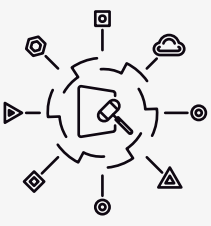

Alignment with the Broader Security Strategy

The Zero Trust model is insufficient in itself as a comprehensive security strategy. Keeping endpoints healthy, as well as monitoring, detection, and incident response capabilities are critical to ensure network safety. Technology solutions, though important, cannot replace a holistic security approach that considers the organization's broader security needs and compliance obligations.

The National Cyber Security Center of Excellence recommends four main features of a Zero Trust model:

Identify

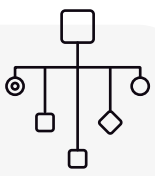

Creates an inventory of systems, software, and other resources, classifies them, and sets baselines to allow for detecting anomalies.

Protect

Authentication and authorization processing. Zero Trust protection includes policy-based resource authentication and configuration, as well as software, firmware, and hardware integrity checks.

Detect

Identifies anomalies and suspicious events, by continuously monitoring network activity to proactively detect potential threats.

Respond

once a threat is detected, handles threat containment and mitigation.

Source: "Implementing a Zero Trust Architecture"

NCCoE
NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

These capabilities are typically implemented by several IT and security solutions, which work together to create a Zero Trust environment.

According to the National Institute of Standards and Technology:

All applications, infrastructure entities and data sources

are defined as resources that need to be protected

All communication

whether inside the corporate network or involving external networks, must be secured

Users and services

- Must be authenticated and authorized before they access resources
- Their activity is monitored and recorded

Users

authorized to use services only for specific purposes, and access should be revoked when no longer needed.

Source: "NIST Special Publication 800-207 Zero Trust Architecture"

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

The Technologies Behind Zero Trust Architecture

The main technologies used to implement a Zero Trust architecture

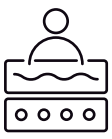

Strong User Verification

Achieved through measures like role-based access control (RBAC).

Identity and Access Management (IAM)

Helps to define and manage user permissions. The IAM system decides whether to grant or deny access requests.

Multi-Factor Authentication (MFA)

Helps protect the network against weak or reused passwords.

Endpoint Protection

Attackers use compromised endpoints to exploit authorized user sessions and gain unauthorized access to company resources. Endpoint security can help protect against compromised accounts.

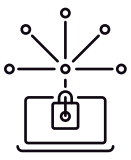

Zero-Trust Network Access (ZTNA)

Remote connections often use telework. To ensure secure remote access, ZTNA technologies provide continuous monitoring for remote connections.

Microsegmentation

Enables the enforcement of Zero Trust policies inside the network.

Challenges with Implementing a Zero Trust Architecture

Zero Trust is a paradigm shift for most organizations, and implementing it in large scale networks can be challenging. Here are some of the key challenges faced by organizations as they adopt a Zero Trust model.

Legacy Applications and Protocols

Mainframes, old HR systems, shell scripting languages like Powershell, and legacy protocols like POP, SMTP, and IMAP are typically incompatible with the Zero Trust approach. Unpatched or decades-old legacy systems that simply can't be patched anymore are exactly where gaps in security and flaws may occur, making it far easier for attackers to make that first step into your data center. There are two approaches for dealing with this:

1/ Excluding legacy systems from the Zero Trust implementation

which can defeat the point of Zero Trust, because those legacy systems become a weak link for attackers to target.

2/ Shutting down or restricting access to legacy systems

which can seriously impair employee productivity, because these systems are part of critical business processes in many organizations.

The Endpoint Zero Trust Gap

While being a great step in the right direction, common Zero Trust approaches have a fundamental design flaw that is the result of a wrong assumption. The wrong underlying assumption is that the Zero Trust broker can check the health of user endpoints and then trust them with access to enterprise resources. This might be true for some extremely locked-down endpoints. However, most enterprise user endpoints run operating systems like Windows and have a very large and potentially vulnerable code base, a wide variety of legacy applications/middleware, and access to risky malicious networks or internet resources. These endpoints can easily be compromised by determined attackers. Once a device is compromised, the operating system can no longer be trusted as malware resides in the same operating system kernel and can tamper with operating system health checks.

This means that many enterprises that adopt Zero Trust may still mistakenly trust user endpoints. This is a critical flaw as it allows attackers to breach a user's device and then ride the user's authenticated session to do harm.

Without this missing link of strong device identity, Zero Trust creates a false sense of security as it encourages enterprises to allow access to corporate resources from personal/unmanaged/BYOD endpoints, relying on basic (and easily forgeable) health checks to prevent malware from getting in. To make matters worse, personal/unmanaged endpoints already have a higher probability of getting infected than corporate devices.

Some enterprises try to close this gap by deploying endpoint detection/protection agents on the user's device, but this ends up being a cat-and-mouse game with endpoint malware. Furthermore, such agents and restrictions often limit what users can do and can lead to issues with user privacy.

Things get worse as more and more employees work remotely and mix both legacy corporate apps and brand new collaboration tools on the same endpoint, opening it up to new types of threats (not to mention the increasing personal usage of endpoints in unmanaged home network environments).

Designing Zero Trust Endpoints

To close this gap in ZTA — and make ZTA a dramatically more secure architecture — enterprises must ensure employees use trusted devices. By re-establishing trust in user devices, it is possible to let users access corporate resources anywhere.

However, this is a challenging task, as enterprises still rely heavily on Windows (or other monolithic operating systems) and legacy applications that are vulnerable and untrusted. Making devices trusted again must also support the migration of existing devices, as solutions that require a fresh start with a new operating system or new devices would fail in any realistic enterprise environment.

One way to increase trust in user access is by requiring that enterprise access is always done through a secure trusted OS, e.g. through a VDI desktop or through a jump box. Even if the OS on the physical device is compromised, having enterprise access run on a separate trusted and isolated OS eliminates a lot of common threats. However, setting up, configuring, and maintaining a whole separate OS image on VDI/jump box is expensive, complicated, and error prone. Furthermore, it degrades the user experience, as it requires users to go through another hop in the network for every interaction they do with every app on that remote OS. In the era of flexible work this has become a serious issue, as connectivity may vary for people in different locations, in different home networks, etc.

To make Zero Trust a true end-to-end security solution, organizations must design their endpoints to be trusted. However, IT cannot just lock down their endpoints to achieve that level of trust – IT must make sure end-users get a great user experience and can get their jobs done in an efficient way. Solutions that limit which apps users can use or remove local admin rights will not fly for today's knowledge workers. Users need a way to use the latest tools and apps, without having IT whitelist every app, website, and service.

Extending Zero Trust for the Endpoint with Hysolate

Hysolate provides an innovative Zero Trust solution to the endpoint. It essentially splits a user's endpoint Windows 10 and Windows 11 device into two segregated zones, each running in its own OS, leveraging the latest hypervisor and virtualization-based security technologies. One OS is the user's unmanaged/less-trusted/personal OS and another is a trusted corporate OS.

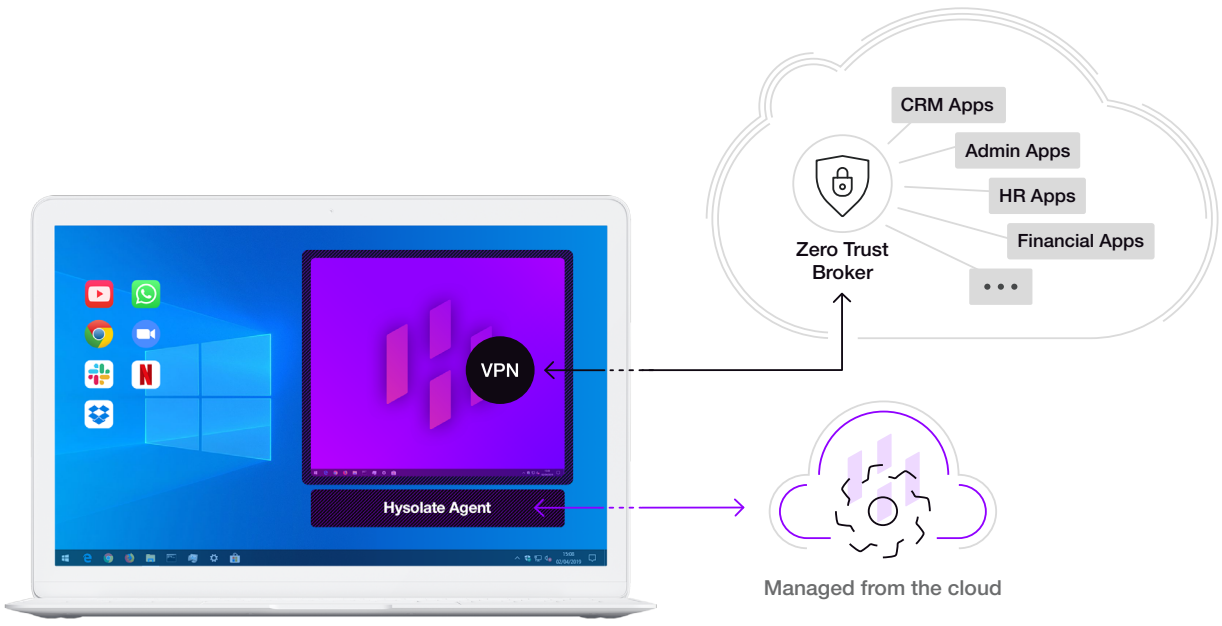

The corporate environment is a fully locked-down operating system that can contain an inaccessible client certificate that vouches for the integrity of the OS. The ZTA broker would only allow that corporate OS to have access to sensitive enterprise applications. The end-user would be unable to access these applications from any other untrusted environment/device.

With Hysolate, IT can isolate the corporate OS from the user's "riskier productivity zone" OS, including detailed controls over clipboard, USB, network, applications and more.

Hysolate is Enterprise-ready

Hysolate is fully enterprise ready, with features including SSO compatibility and SIEM integration, and can be deployed and scaled out to entire teams in minutes. Admins manage Hysolate from a cloud based management console, via customized granular policies, and can easily package and deploy apps during and after an initial deployment. With Hysolate's Zero Trust endpoint solution in place, enterprises can really move to a secure-by-design architecture, without affecting user productivity.

Learn more about Hysolate's Zero Trust access solution, or request a demo.

Hysolate's Zero Trust Access Solution

Request a Demo