WHITE PAPER

AppOmni

SaaS Security Primer

Comparison: SaaS Data Security and Open S3 Buckets

Introduction

Technology professionals have come to appreciate both the promise and the peril of Amazon S3. This simple web services interface empowers users to store and retrieve any amount of data, at any time, from anywhere on the web. That's a powerful capability. But such power comes with a price, as an open S3 bucket can leave valuable data exposed.

Extensive media coverage and security research has shown how misconfiguration can inadvertently lead to this data being stolen and exploited by bad actors. While much has been made of the vulnerability of misconfigured S3 buckets, IT experts have often overlooked an analogous threat that is even more potentially damaging – the threat of misconfigured SaaS applications.

S3 draws its strength from the sheer speed, fluidity, and volume with which it handles data. The same can be said of the SaaS applications that today support key enterprise business functions. And just as with S3 buckets, misconfigurations here can give bad actors ready access to the data in these applications.

If anything, the stakes may be higher in SaaS, where key information on personnel, financials, and other vital business functions is clearly labeled and readily identified.

The Risks in S3

In order to fully appreciate the configuration risks in SaaS, it's helpful to work by analogy. The perilous terrain around S3 buckets is already familiar to many, and serves to illustrate the challenges that arise when key data-driven systems are not managed appropriately.

As a public cloud storage resource available in Amazon Web Services' (AWS) Simple Storage Service (S3), an Amazon S3 bucket stores objects consisting of data and its descriptive metadata. As has been well documented, there are any number of online tools available on GitHub that enable a user to find a web site's S3 bucket: They effectively offer a Google-type search for exposed data.

At the same time, AWS access control settings are known to be complex, and are susceptible to a wide range of attacks. These known vulnerabilities have led security researchers to pursue S3 bug bounties with some enthusiasm – often with startling results. "Recently during my bug hunting, I came across a misconfigured AWS S3 bucket of an Indian E-commerce Company which gave me full access to their S3 bucket, allowing [me] to download, upload and overwrite files," wrote one bug hunter.[1]

Recent statistics suggest that as many as 7% of all S3 servers are completely publicly accessible without any authentication, and 35% are unencrypted.[2] Moreover, vulnerable S3 buckets can be all too easy to identify.

An S3 bucket can be accessed through its URL, which typically comes in one of two readily identifiable forms:

http://s3.amazonaws.com/[bucket_name]/
http://[bucket_name].s3.amazonaws.com/

Since bucket names are not secret, access can sometimes be achieved merely through brute-force guessing. There are also multiple ways to make an S3-bucket actually reveal itself independent of proxies in front of it. The domain history may expose the bucket name, response headers of objects in the bucket may have metadata that reveals the bucket name or content may be tagged with the bucket name.

In order to test the openness of a bucket, one need only access the bucket's URL from a web browser. A private bucket will report "Access Denied," and will not show any bucket contents. A public bucket, on the other hand, will list the first 1,000 files contained in that bucket – simply by clicking the URL.

[1] #BugBounty — AWS S3 added to my "Bucket" list!
[2] 7% of All Amazon S3 Servers Are Exposed, Explaining Recent Surge of Data Leaks

Once a bucket has been located, bad actors have considerable room to operate. By using the AWS Command Line to talk to Amazon's API, the attacker can list and read files in S3 bucket, write/upload files to the S3 bucket, and change access rights to content contained in the files. A few examples help to describe the potential pitfalls:

- The widely reported "Million Dollar Instagram Bug"[3] allowed a security researcher to access every image and account on Instagram via the S3 bucket, where the company stored everything from source code to images.
- One bug bounty operation found misconfigurations in its own S3 bucket, which allowed any authenticated AWS user to write to it.
- Another white hat hacker has reported using S3 misconfigurations to control assets on high profile websites -- overwriting files, uploading vulnerable files, and downloading Intellectual property.
- One major telecommunications provider's poorly secured S3 systems allowed high-profile data exposure, with 6 million customer records exposed when an employee of a partner firm placed log information from customer service calls on a publicly accessible S3 server.

The security bloggers have spelled out in excruciating detail the many ways in which S3 buckets can be misconfigured. Permissions that should be powerful and granular may instead be very broad. Users may fail to implement appropriate encryption for sensitive data. Private and public data may reside in the same buckets, leading to overly broad access.

These and other configuration mistakes can be damaging in S3, where much of the data is unstructured – and they can be lethal in SaaS applications, where data is classified, labeled and ripe for the picking.

The SaaS Landscape

For all that IT has done to harden its S3 profile, far less attention has been paid to the corollary risks that exist in the SaaS environment. That's especially problematic given the vast tracts of critical business intelligence that reside in these systems.

While there is no tool that allows a Google-type search for open SaaS connections, such as exists with S3, it's likely that such tools will arise as more people learn about the potential for misconfiguration in this space. Even in the absence of such a mechanism, SaaS vulnerabilities are easy enough to spot for those who know where to look.

Improperly defined security groups may fail to restrict ingress and egress. Misconfiguration and open ports can allow access from IPs not associated with your organization. Secrets may be stored in unencrypted application code. Default settings may allow unlimited ability to share data outside the organization. Upgrades may introduce new security risks. Misconfigured permissions of high-powered accounts create additional vulnerabilities.

> *"Security and risk management leaders should invest in cloud security posture management processes and tools to proactively and reactively identify and remediate these risks"* – Gartner

As SaaS leaks and hacks become increasingly common, Cloud Security Posture Management (CSPM) is more vital than ever. The CSPM methodology enables an organization to discover, assess, and resolve cloud misconfigurations. As defined by Gartner, this class of security tools can be applied to support compliance monitoring, DevOps integration, incident response, risk assessment, and risk visualization.

"Security and risk management leaders should invest in cloud security posture management processes and tools to proactively and reactively identify and remediate these risks," according to Gartner.

[3]"Instagram's Million Dollar Bug": Case study for defense

The CSPM approach has been most closely identified with Infrastructure as a Service (IaaS), as a means to continuously manage cloud security risk. It typically will deliver detection, logging, reports, and other automated means to remediate cloud service configurations and other security issues. Used proactively, it will prevent leaks and hacks from happening in the first place.

In IaaS deployments, CSPM has been shown to reduce cloud security misconfiguration incidents. It can help to reduce fatigue among security operators, manage cloud security at scale, and deliver a higher level of compliance. All these same benefits apply when CSPM is implemented in support of SaaS applications.

Those who take this proactive approach to SaaS, on a par with their efforts to secure S3, have an opportunity now to get ahead of these emerging security issues.

Specific Risks

Why are the risks potentially higher in SaaS? It has to do with the nature and quality of the data under fire. Where S3 data may be unstructured and arbitrary, SaaS applications are rich with structured, labeled, and often highly valuable business intelligence.

These applications may support key business processes and may contain sensitive financial information. They may hold data that must be protected for legal or compliance reasons, such as healthcare data or personally identifiable information (PII). Such data may also be embedded in documents and files within these systems.

> *"Your company might have thousands of users with access to Salesforce, and you can't necessarily trust everyone with the same level of access."* – Max Feldman, AppOmni Security Engineering Manager

Access control around these sensitive systems is critical. "You can have various different types of users, and some of those people might be on the public internet. They might not be employees of your company," said Max Feldman, AppOmni security engineering manager. "Your company might have thousands of users with access to Salesforce, and you can't necessarily trust everyone with the same level of access."

The brand-name nature of SaaS may also induce a false sense of security. It's easy to assume that a Salesforce or Box frontend ensures that all permissions and configurations are pre-configured to inherently safeguard the goods. Due to unique needs of individual deployments, it's not realistic to expect the SaaS providers to provide pre-configured security settings that are suitable for all deployments, Much of the power and responsibility surrounding the appropriate configuration of SaaS rests with the customer.

In addition, familiarity factors in. Many of the steps needed to secure an S3 bucket are well known and well documented. SaaS configurations are newer and less familiar to administrators.

This risk is multiplied for any organization that incorporates customers or user data within a SaaS deployment. If the application supports a customer- or external-user portal, for example, it may be a repository for this kind of information. If those customers aren't provisioned correctly, with appropriate access restrictions monitored and enforced, it is possible to inadvertently expose private information to third parties.

You don't even have to be a hacker to take advantage of this. Any outsider who becomes a party in the portal can potentially gain access to other users' information, either deliberately or simply by typing in the wrong URL, depending on the extent of the misconfiguration.

These systems may expose data that is obviously critical – financial information for example. Or the data may be seemingly innocuous. An executive calendar that is accidentally shared with the public may not look like much of a threat, for instance, unless it contains an entry like "hold for IPO discussion."

Overall, the sheer ease of use within the SaaS environment can lead to lax security practices. The very promise of a SaaS application with its ease of adoption and use – can put the enterprise in an inherently vulnerable position.

Going Forward

Given the potential vulnerabilities that may exist in the SaaS environment, organizations may need to adopt a new and more robust stance toward SaaS security. They will need to seek out solutions that implement CSPM organically in the as-a-service environment, taking the best practices from their experiences in the IaaS world and applying them to their SaaS deployments.

To address possible exposure of user data, they need solutions that leverage automation in order to continuously monitor for risky configurations, and to automatically alert the watchers so that timely corrections can be made.

In order to more fully safeguard SaaS, IT leaders need to begin by taking inventory, assessing who exactly has access to what data within these complicated and sprawling systems. This effort to review permissions and policies is a critical first step in ensuring that appropriate controls are in place, and it helps to highlight those areas where default configurations may not have supplied adequate protections.

Along with a review of existing permissions, it's also helpful up front to inventory the many and varied applications that may already be in use, and to review the relative sensitivity of the data they contain. This initial risk assessment will help to target resources as an enterprise moves toward more rigorous management of SaaS configurations.

ABOUT APPOMNI AppOmni's platform implements guardrails for internal and external users of mission critical SaaS applications, enabling employees to work productively and securely. As the world's leading provider of CSPM for SaaS, AppOmni enables organizations to quickly identify and remediate risky configurations, improper access controls, and data exposures before a breach can occur. The platform provides customers an always-on dashboard that delivers visibility into the current state of their cloud/SaaS applications, deviations from their business intent, and a simple way to immediately verify the state of any application for both functional and security requirements. The platform can be deployed within minutes and immediately integrates into existing workflows to provide actionable insight and visibility on Day 1.