



# 7 STEPS TO HARDEN CLOUD SECURITY POSTURE AND PREVENT BREACHES





What you'll learn in this paper:

- ✔ A pragmatic approach to industrialize cloud security and prevent breaches
- ✔ Insights into high-profile breaches and how they could have been prevented
- ✔ Making sense of security-tool acronym soup, and what to adopt when and why
- ✔ Key indicators and considerations to measure the success of your security program

## The Formula for Hardening Cloud Security:

- 1 Assess the security posture of cloud runtimes
- 2 Prioritize remediation based on measurable risk
- 3 Track drift to prevent new breach paths
- 4 Stop risky deployments before they happen
- 5 Extend visibility across hybrid cloud environments
- 6 Remediate high-risk hybrid attack paths
- 7 Continuously assess cloud and cyber risk

# CLOUD BREACHES ARE COMMONPLACE

Despite massive investments in security, such as threat detection and response tools, we continue to see disproportionate growth in the volume of cloud breaches. In fact, IBM's latest Cost of Data Breach 2022 report, found that nearly half of all breaches were cloud based. And most cloud breaches are the result of poor cloud hygiene— a simple misconfiguration, vulnerability, or excess privilege— exposures that could easily have been addressed, but that went undetected or unremediated.

Common examples of each type of exposure:

- ⚠️ **Misconfigurations:** Open ports, unencrypted data, expired certificates, clear text passwords
- ⚠️ **Vulnerabilities:** Unpatched systems and applications, undetected malware
- ⚠️ **Privileges:** Excessive admin or root-level permissions, no multi-factor authentication

## This begs the question, "Why?"

Despite all its benefits, cloud is a perfect storm of three factors that are the natural enemy of security: speed, scale and skills shortages.



Let's explore each:

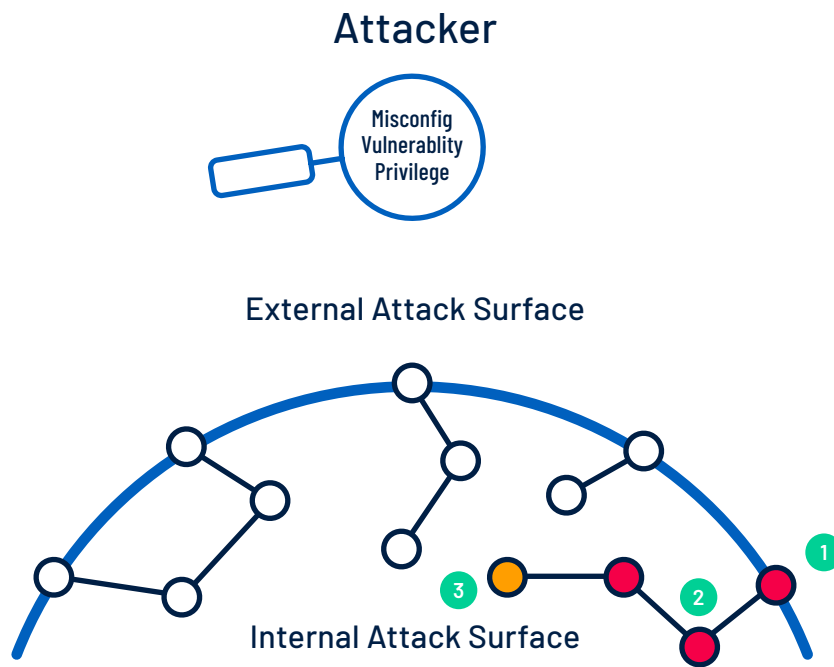
- ✅ **Speed:** The pace of cloud adoption has only accelerated in recent years, with more workloads moving to the cloud, and more cloud developers building and deploying on cloud on a continuous basis. In fact, in cloud first organizations, the ratio of developers to security teams can reach as high as 100:1.
- ✅ **Scale:** Traditional architectures have given way to complex cloud native architectures, with orders of magnitude more moving and reusable parts – microservices, containers, Kubernetes, infrastructure as code (IaC). One misconfiguration, vulnerability or excess of privilege can quickly be replicated at scale.
- ✅ **Skills:** To mitigate risk, most organizations have adopted a multi-cloud strategy, adopting a mix of cloud vendors and platforms. Each cloud platform brings its own unique services, security tools, configuration best practices and skills requirements.

The result is overburdened security teams that are struggling to keep pace with growing alert volumes and an inability to cost-effectively scale people, skills and processes – even as cloud adoption continues to accelerate.

This presents a perfect opportunity for attackers.

# Attackers focus on what we don't see






The reality is, while security teams chase growing volumes of alerts in cloud runtimes, attackers spend most of their time looking for the one weakness in your cloud security posture.



Consider the approach to virtually every attack:

- 1 Identify misconfigs, vulnerabilities and excess privilege in the external attack surface to gain entry
- 2 Exploit additional exposures to move laterally, and identify high value targets
- 3 Execute against endgame objective: exploit data, impact availability, hold for ransom

## High Profile Cloud Breaches:

|        | <a href="#">Capital One</a>  | <a href="#">MGM Grand Hotels</a>  | <a href="#">Accenture</a>   |
|--------|--|---|---|
| Breach |  106 Million Client Records                                 |  142 Million Client Records                    |  137 Gigabytes of Client Data                |
| Cause  |  Misconfigured Firewall + <a href="#">Excess Privileges</a> |  <a href="#">Cloud Server Misconfiguration</a> |  <a href="#">Open AWS S3 Storage Buckets</a> |
| Impact |  \$270 Million in Fines and Lawsuits                        |  Client data for sale on the Dark Web          |  40,000 Client Passwords Compromised         |

What all of these have in common is they could have been prevented using cloud security posture management (CSPM) to maintain secure configuration and enforce robust digital hygiene.

# A PRESCRIPTIVE APPROACH TO HARDENING CLOUD SECURITY

The following seven steps to harden cloud security posture are based on real-world adoption strategies from leading enterprises and the best practice recommendations of security and compliance frameworks.

92% of companies have a multi-cloud strategy, according to a recent [“State of the Cloud Report”](#) by Flexera.

## 1. ASSESS THE SECURITY POSTURE OF CLOUD RUNTIMES

### Adopt a best-practice framework

While most security operations teams have well documented best practices for on-prem security, many lack an equivalent set of best practices for cloud. Further, with increased multi-cloud adoption, it can be difficult for security teams to define consistent best practices across individual cloud providers. This is because each provider has its own unique infrastructure, services and security tools.

Benchmarks like the [Center for Internet Security \(CIS\) Foundations Benchmarks](#) can be a valuable asset, providing prescriptive guidance and best practices on how to securely configure specific cloud vendor environments, including: [Amazon Web Services](#), [Microsoft Azure](#), [Google Cloud Platform](#), and other cloud platforms. CIS Benchmarks are developed and approved by leading governments, businesses, industries, and academia, and are available at no cost.

Other cloud benchmarks to explore include: Cloud Security Alliance’s (CSA) Cloud Controls Matrix (CCM), and ISO/IEC 27017:2015.



## Establish an effective foundation for cloud security

As we noted earlier, attackers can gain entry through the use of misconfigurations, including unrestricted ports, vulnerabilities such as unpatched operating systems, and escalated privileges, such as those gained through phishing attacks. This is why many risk management frameworks, such as the widely recognized [National Institute of Standards and Technology \(NIST\) Cybersecurity Framework](#) recommend organizations begin with an inventory of assets, which includes all hardware and software.

Having an accurate inventory of your attack surface, both external and internal, is critical, as breaches can begin anywhere, and attackers can quickly move laterally. The goal should be to achieve a single source of truth for asset inventory, that is as accurate, and up to the minute as possible. This should include all running cloud instances across your cloud runtimes: including production, staging, QA/test and development environments.

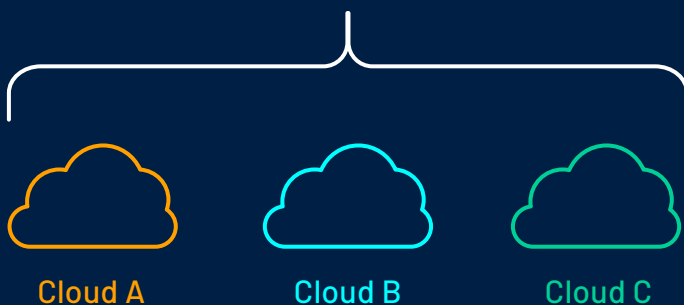
While an inventory of assets can be achieved through manual audits, or by aggregating data from multiple cloud provider tools, CSPM solutions can greatly simplify the time and manpower needed. They do this by continuously discovering new cloud accounts and assets across cloud providers as they are added.

Once an accurate asset inventory is achieved, it is time to assess cloud runtimes for potential exposure. CSPM solutions provide out-of-the-box policies which directly align to popular control frameworks and benchmarks, such as those from NIST and CIS, greatly simplifying the assessment of runtime security posture. For CIS benchmarks, for example, the best-practice configurations are automatically provided allowing for faster remediation across any cloud, without the need for deep expertise in that cloud.

Organizations can apply these policies as they are or customize their own policies to the needs of their organization. Another value of CSPM tools is the ability to apply consistent policy enforcement across cloud providers, augmenting or eliminating the need for individual cloud service provider (CSP) security tools.

### Cloud Security Posture Management

#### Consistent Policy Enforcement



#### Selection Tips:

If you require custom policies, CSPM tools with low code editors can greatly simplify policy creation.

Explore vendors that offer robust benchmark certifications, such as CIS, for leading cloud vendors.

## Assess the full spectrum of exposure

Out of the box, all CSPM tools detect high-risk misconfigurations in runtime cloud environments such as public access to cloud data stores, unencrypted data and expired certificates.

But to fully assess security posture, organizations must have visibility into Common Vulnerabilities and Exposures (CVEs) in your runtime environment, as well as excess privileges not required by roles and systems to adequately perform their job or function. For example, root-level privileges are not required for daily administration tasks, and such accounts should be restricted, with their access monitored, API keys deleted and MFA strictly enforced to avoid critical risks to your cloud infrastructure through account-takeover attacks.

### Selection Tips:

When evaluating cloud security tools, prioritize solutions that can provide a full view of asset risk, including detection of misconfigurations, CVEs and excess privileges.



Misconfigs



CVEs



Privileges

## ONUS Breach

2 million customer credentials and personal information were exposed, when [ONUS](#), one of the biggest cryptocurrency platforms in Vietnam, was [hacked](#) via the [Log4j vulnerability](#) and escalated due to AWS [S3 bucket misconfigurations and permissions](#).



2 Million Client Records



Log4j Vulnerability  
S3 Bucket Permissions



\$5 Million Ransom



## Cloud-native architectures change the game

Cloud native architectures using microservices, containers and Kubernetes bring more moving parts than traditional 'lift and shift' applications running on virtual machines. The use of a reusable services model and autoscaling means a single vulnerable container image can be replicated at scale in a production environment, sometimes as much as 1:100, equally multiplying possible exposure. Further, images can come and go in days or hours, before a scheduled scan can detect the vulnerability, but long enough to be exploited by an attacker.

This ephemeral nature of images, combined with the labor overhead required to deploy agents to each new container is not only unscalable for staff, but it quickly becomes cost prohibitive due to escalating license fees for agents.

Agentless vulnerability scanning takes a different approach, most often using APIs or read-only access to assess cloud provider accounts. In doing so, it bypasses the need for agent deployment, reducing the labor overhead and licensing cost, while providing near real time visibility.

### **Selection Tips:**

Agentless and agent-based approaches can be highly complementary when dealing with a mix of long-lived applications on virtual machines and containerized applications. Consider cloud security solutions that offer a choice of both.



## 2. PRIORITIZE REMEDIATION BASED ON MEASURABLE RISK

With limited staff on hand to support remediation efforts, and a growing shortage of cloud security experts globally, the ability to prioritize remediation based on measurable risk is critical to breach prevention and assessment of overall cyber and cloud exposure. It is also a key way in which organizations will reduce the time spent chasing security alerts and incidents in production.

### Using control frameworks and benchmarks

Frameworks such as [NIST 800-53B](#) provide a mapping of what specific controls should be applied to secure assets based on their strategic importance or impact to an organization's mission. The specific standards for classifying those assets based on low, medium, and high impact criteria are provided within NIST publication [FIPS 199](#).

### SP 800-53 Rev. 5.1 and SP 800-53B Latest Version

Controls Contain:  
ACCESS CONTROL Family  
Showing 25 controls

| No.                   | Control Name   | Low-Impact | Moderate-Impact               | High-Impact                             |
|-----------------------|--|------------|-------------------------------|---|
| <a href="#">AC-1</a>  | POLICY AND PROCEDURES                                      | AC-1       | AC-1                          | AC-1                                    |
| <a href="#">AC-2</a>  | ACCOUNT MANAGEMENT   | AC-2       | AC-2 (1) (2) (3) (4) (5) (13) | AC-2 (1) (2) (3) (4) (5) (11) (12) (13) |
| <a href="#">AC-3</a>  | ACCESS ENFORCEMENT   | AC-3       | AC-3                          | AC-3                                    |
| <a href="#">AC-4</a>  | INFORMATION FLOW ENFORCEMENT                               |            | AC-4                          | AC-4 (4)                                |
| <a href="#">AC-5</a>  | SEPARATION OF DUTIES                                       |            | AC-5                          | AC-5                                    |
| <a href="#">AC-6</a>  | LEAST PRIVILEGE  |            | AC-6 (1) (2) (5) (7) (9) (10) | AC-6 (1) (2) (3) (5) (7) (9) (10)       |
| <a href="#">AC-7</a>  | UNSUCCESSFUL LOGON ATTEMPTS                                | AC-7       | AC-7                          | AC-7                                    |
| <a href="#">AC-8</a>  | SYSTEM USE NOTIFICATION                                    | AC-8       | AC-8                          | AC-8                                    |
| <a href="#">AC-9</a>  | PREVIOUS LOGON NOTIFICATION                                |            |                               |   |
| <a href="#">AC-10</a> | CONCURRENT SESSION CONTROL                                 |            |                               | AC-10                                   |
| <a href="#">AC-11</a> | DEVICE LOCK  |            | AC-11 (1)                     | AC-11 (1)                               |
| <a href="#">AC-12</a> | SESSION TERMINATION  |            | AC-12                         | AC-12                                   |
| <a href="#">AC-13</a> | SUPERVISION AND REVIEW — ACCESS CONTROL                    |            |                               |   |
| <a href="#">AC-14</a> | PERMITTED ACTIONS WITHOUT IDENTIFICATION OR AUTHENTICATION | AC-14      | AC-14                         | AC-14                                   |



This mapping from NIST 800-53B shows which specific controls should be applied to each asset classification.

CIS Controls and CIS Benchmarks provide alternative approaches to NIST. CIS Controls are classified in three levels, or Implementation Groups (IG): IG 1 is essential cyber hygiene, and the minimum standard for information security which should be

implemented across the board, IG 2 builds on IG 1, while IG 3 represents all possible controls. CIS Benchmarks for specific clouds provide similar levels of classification in the form of Profiles: Level 1 being a base level of security for all systems, Level 2 offering in depth security, and STIG (formerly Level 3) being configuration standards used by the U.S. Department of Defense.

## Risk scoring

The methods used to measure risk can vary greatly based on the approach of the independent body or vendor technology used, but generally align to two main considerations: first, the probability of a cyber event occurring, such as the exploitation of a vulnerability, and second, the business impact of a cyber event factoring in variables such as asset criticality or other indicators of potential business impact. We explore examples of these below.

### CVSS score

Perhaps one of the most well known models for calculating the probability of a cyber event occurring is the Common Vulnerability Scoring System (CVSS). CVSS is a free and open industry standard for assessing the severity of computer system security vulnerabilities. CVSS attempts to assign a base severity score from 1-10 (10 being most severe). While many utilize only the CVSS base score for determining severity, temporal and environmental scores also exist to factor in availability of mitigations and how widespread vulnerable systems are within an organization.

Many regulatory bodies have adopted a maximum CVSS score for compliance as part of their requirements for regulations. For example, a score of 4.0 or higher, for Payment Card Industry Data Security Standard (PCI-DSS), qualifies as non-compliant, and requires remediation. According to NIST, CVSS is not an actual measure of risk. It is rather a measure of threat severity. It is also specific to CVEs, not to all flavors of potential exposure, such as cloud misconfigurations.

### EPSS Score

Another example of a model for calculating the probability of a cyber event occurring is the Exploit Prediction Scoring System (EPSS). According to the Forum of Incident Response and Security Teams (FIRST), EPSS is “an open, data-driven effort for estimating the likelihood (probability) that a software vulnerability will be exploited in the wild.” The goal of EPSS is to enable defenders to better prioritize vulnerability remediation by leveraging current threat information from real-world exploit data. The model produces a probability score between 0 and 1 that represents the likelihood that a vulnerability will be exploited.

EPSS is primarily focused on overcoming the fact that CVSS does not provide any measurement on exploitation risk. Unfortunately, EPSS does not provide any guidance on the impact of exploitation or the importance of the asset that is vulnerable, so on its own it is not sufficient for making well-informed decisions around risk-informed prioritization.

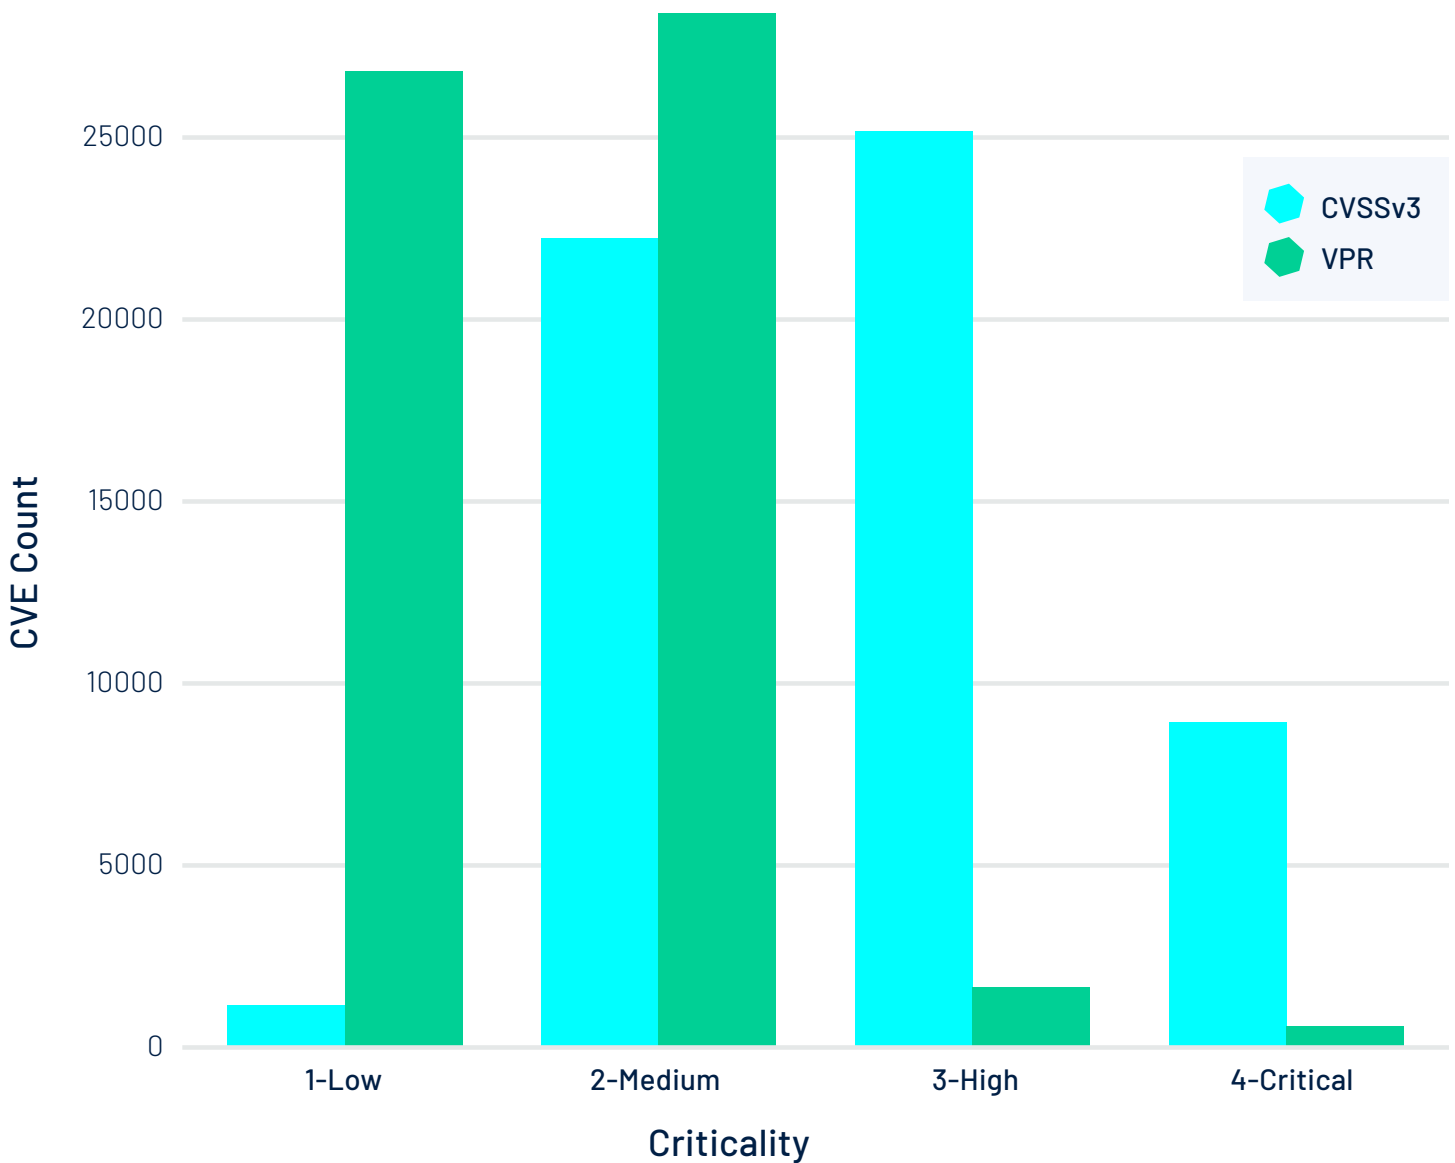
### Asset Exposure Score

Our third example, Asset Exposure Score (AES) is a model that uses both the probability of a cyber event occurring, and the business impact of a cyber event to calculate a risk score. AES leverages a multidimensional understanding of attack paths, including identity, privileges, and topology to assess both exploitability using a Vulnerability Priority Rating (VPR) and the potential impact, using an Asset Criticality Rating (ACR). VPR assesses the age of a vulnerability, static CVSS score, exploit code maturity, threat severity, and threat intensity based on recent activity by attackers. ACR assesses the location of the asset on your network and proximity to the internet, asset type and asset capabilities.

The intent is to look at risk holistically, and calculate a combined risk score that assesses all known risk factors associated with a given asset and vulnerability. One immediate value of this approach is that it does not silo risk into separate buckets for ad-hoc remediation. Instead, a measurable AES ranks which specific assets pose the greatest total threat of exposure for the business. This allows for prioritized remediation, and makes it possible to track risk reduction trends in a measurable way.

Organizations have seen a [97% reduction](#) in Critical and High Severity vulnerabilities using multidimensional VPR as compared to CVSSv3 alone.

Compare Distributions of CVSSv3 CVEs by VPR and CVSSv3 Criticality



This graph provides a visual comparison of CVSSv3 CVEs by severity as compared to the same CVEs scored using VPR. To learn more, check out this Tenable [blog](#).

**Selection Tips:**

Look for cloud security solutions that can assess AES based on an understanding of attack paths and dynamic VPRs which reduce the number of vulnerabilities requiring immediate remediation by 97%.

### 3. TRACK DRIFT TO PREVENT NEW BREACH PATHS

Once runtime risk has been remediated, a secure baseline is established. Optimally, any configuration changes to infrastructure in runtime should be disallowed and should instead be made prior to deployment.

#### Drift happens

However, it is possible that configuration changes to cloud infrastructure can happen in runtime for a variety of reasons. A recent study found that over 90% of organizations allow users to make changes to cloud infrastructure in runtime. Any change to a cloud configuration has the potential to be a security policy violation which introduces new breach paths. It is also possible that a change is made to the runtime environment to address a security concern. As such, it is critical to continuously monitor runtime infrastructure for drift. Common drifts include changes to security group configurations and IAM policies.

#### Drift comes in multiple forms


There are two primary varieties of drift. Cloud-to-cloud drift represents any changes to runtime configuration from the baseline state. Code-to-cloud drift represents changes between Infrastructure-as-code (IaC) and cloud configurations in runtime. The latter is used when an IaC baseline of compliant code has been established, allowing for comparison of that IaC baseline with the current cloud configuration to assess if there is a policy violation. For this reason, it is important to ensure uniform governance policies are used to assess both varieties of drift.

#### Microsoft breach

250 million [Microsoft](#) customer service and support records spanning a 14-year period were exposed when a change was made to the network security group that controls inbound traffic to servers.



**250 Million Client Records**



**[Server Misconfiguration](#)**



**Client Satisfaction  
& Brand Health**


## Drift remediation

Another important consideration with respect to drift is remediation. Next generation CSPM tools can provide detailed visibility into configuration and change history for both varieties of drift. More importantly, they can greatly streamline remediation by providing detailed step-by-step instructions for compliant configuration. Code-to-cloud drift specifically allows for further optimization of remediation, providing a line-by-line comparison of IaC vs current runtime configuration and the exact lines of code or runtime configuration that have changed. This is important not only to catch drift in runtime that introduces new risk, but also to inform developers if a security issue was addressed in runtime, so it can be addressed in IaC to ensure future releases remain compliant. Pull requests can be automatically generated and sent to developers with the specific code fixes, or even sent to auto remediate, which can significantly reduce mean-time-to-remediation. Given developers can often move on to new projects, remediation can take days from the time drift is detected to the time a developer implements a fix. Auto generated code and auto-remediation therefore significantly reduce developer overhead, and more importantly, can shrink the window of exposure from non-compliant configurations that can lead to breaches.

### Selection Tips:


If you use IaC or plan to, look for CSPM solutions that can detect both cloud-to-cloud and code-to-cloud drift to ensure alignment across DevOps teams.

To reduce mean-time-to-remediation, consider CSPM tools that offer prescriptive remediation steps to fix out of compliance configurations, as well as automated code fixes and pull requests that can greatly reduce the time and effort of remediation by developers.



## 4. STOP RISKY DEPLOYMENTS BEFORE THEY HAPPEN

Earlier, we discussed how cloud-native architectures using microservices, containers and Kubernetes, due to the reusable nature of components and auto scaling, can greatly increase the volume of moving parts that must be managed in runtime environments. As such, a single misconfiguration in a base or golden image can be replicated at scale, putting greater burden on teams that must remediate those issues. This same principle holds true for IaC, which is often replicated at scale and shared by developers. The best way to minimize this risk is to prevent risky deployments from reaching cloud runtimes.



**Shift left gradually.** There are several points in the DevOps lifecycle where this can be accomplished.

- ✔ **CI/CD pipeline scanning:** Policies can be run against IaC files in CI/CD tools (such as Terraform Cloud, Jenkins, CircleCI, GitHub Action and Azure DevOps Pipeline) to detect violations at build or release stages, notifying developers of failed status, or entirely stopping releases for more severe problems.
- ✔ **Source Code Management (SCM)/repository scanning:** Policies can be run against source code management repositories to assess IaC projects. Pull requests can be generated with remediation steps, as well as auto-remediation that corrects errors in code, preventing risky deployments.
- ✔ **Container registry scanning:** Container images stored in registries (such as Amazon Elastic Container Registry (ECR), Azure Registry, Google Container Registry (GCR) can be scanned to identify vulnerabilities prior to deployment.
- ✔ **Code scanning by developers:** Policies can be run against code files on developer desktops to identify violations. As the earliest compliance check in the DevOps lifecycle, self-scanning of code by developers can massively reduce the issues that make it to production.

CSPM, IaC security scanners, and container security tools offer varying degrees of pre-release scanning, guardrails and remediation.

### Selection Tips:

Consider CSPM tools that can scan CI/CD, Container Registries, SCM, and Developer Code using a consistent set of policies. This maximizes visibility into potential risky configurations across the DevOps lifecycle, before release for improved breach prevention, and greater economies of scale.

## 5. EXTEND VISIBILITY ACROSS HYBRID CLOUD ENVIRONMENTS

At this point, we have established a secure baseline for our cloud runtimes, which are hardened against misconfigurations, vulnerabilities and excess privileges that lead to the majority of breaches. We have put additional checks in place to identify drift from that hardened state and correct it, plus guardrails to prevent new risky configurations and image vulnerabilities from being deployed into our secure runtime.

As with any variety of risk, it is what you don't see that can leave you exposed. And for many organizations, the reality is hybrid environments that span traditional IT on premises and multiple clouds. Further, the number of internet-connected assets can range from tens of thousands to well over 70,000 assets. Savvy attackers can use any gap in visibility to enter and move laterally. With rapid adoption of hybrid-cloud applications which span data centers and public clouds, bring-your-own-device (BYOD), work from anywhere, Internet-of-things (IoT), and shadow cloud accounts, it is vitally important to ensure that we can see the full external attack surface - both known and unknown assets - in the data center and in the cloud.


The problem for most enterprises is that the process of maintaining an inventory of externally facing assets, or any assets for that matter, is often very manual, prone to human error, and time consuming. This is further compounded by the fact that the resulting asset inventory is quickly out of date. The reality is enterprise infrastructures are often highly dynamic, living environments. An employee working from home, a partner remotely connecting to the network from an unsecured or compromised device, or a misconfigured web application can provide a means of access to a resourceful attacker.

### Equifax breach

One of the most impactful examples of this is the [Equifax breach](#) in which 155 million client records were compromised due to a [web application vulnerability](#) that was missed due to poor asset management. The impact to Equifax was combined fines and lawsuits equaling \$1 billion.



**155 Million Client Records**



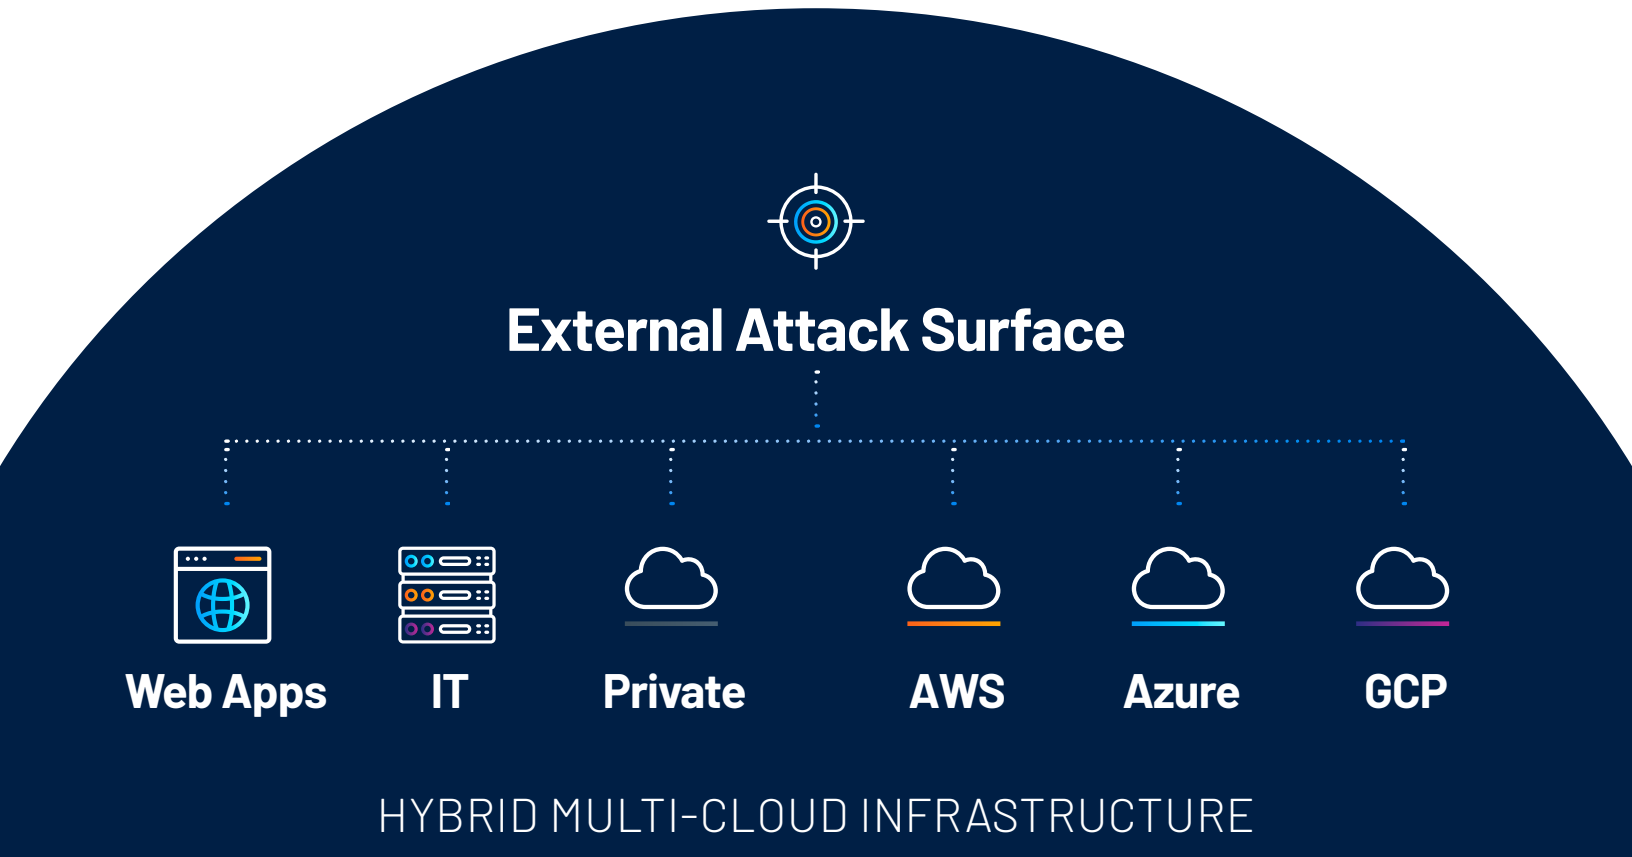
**[Web App Vulnerability](#)**



**\$1Billion in  
Fines and Lawsuits**

[External attack surface management \(EASM\)](#) extends visibility across hybrid environments by discovering internet-facing assets, applications and resources which might provide an entry point for attackers. It also provides context and attribution details, such as geo-IP physical location, cloud or Content Delivery Network (CDN) provider, Content Management System (CMS) type, certificate status and many other

metadata fields which can help cybersecurity teams determine asset criticality, identify asset ownership and help make more informed decisions based on potential exposure. [Web application scanning \(WAS\)](#) augments the visibility provided by EASM with Dynamic Application Security Testing (DAST) that assesses running web applications to identify application flaws, vulnerable web app components and API weaknesses.



**Selection Tips:**

When exploring EASM and WAS solutions that can extend visibility from public cloud environments to on-premises private clouds, IT environments, and hybrid web applications, those that can bring together data and context into a single pane of glass can help maximize asset visibility and workflow efficiency across teams.



## 6. REMEDIATE HIGH RISK HYBRID ATTACK PATHS

One of the highest visibility areas of cloud security is attack path analysis (APA), and for good reason. The purpose of this ebook is to address the problem of rising exposure from growing cyber attacks. In attacks that originate with a point of entry, the attackers then move laterally and escalate privileges until they can achieve a desired outcome. This journey is by definition an attack path – the steps an attacker takes to reach their goal.

So far we have taken a linear path to harden cloud security posture, starting with foundational security to protect assets from common risks. And as we discussed in this ebook's step 2, an understanding of attack path dependencies is needed to effectively calculate the Asset Exposure Score for prioritization. APA is recommended at this phase because APA tools require a detailed understanding of assets, relationships, vulnerabilities and identity to build an accurate picture of an attack path and to enable additional use cases. And many of the core security tools we discussed lay that foundation for APA.

### APA for cloud vs hybrid cloud

Today, some APA tools provide mapping for specific domains, such as cloud. While valuable for specific teams and use cases, they lack the underlying data needed to identify attack paths that can cross from on prem to cloud and vice versa. This can ultimately lead to unseen exposure and breaches. In contrast, APA solutions built on a common data framework offer the added benefit of aggregating data from multiple points of telemetry into a single data model. For example, deep privilege context, attack surface exposures spanning identity systems, IT assets, and web apps, and passive monitoring sensors. The advantage of this approach is the ability to build a more comprehensive, and detailed view of end-to-end attack path, which helps prioritize and provide context across multiple teams and use cases.

### SolarWinds Breach


Perhaps the largest breach in terms of the number of organizations impacted is the SolarWinds breach, which affected 18,000 of their clients. Malicious code in the SolarWinds platform provided a backdoor to enter their clients' on-premises network where they used vulnerabilities to move laterally and escalate privileges. From there they used a [SAML attack](#) to gain permissions that allowed them to move from on-premises to cloud environments unchecked. This approach demonstrates how hybrid attack paths are being actively exploited today.



**18,000 Client Exposed**




**[SAML Attack](#)**



**Billions in  
Financial Losses**

## Common APA use cases

- ✔ APA can be used alongside Asset Exposure Score to help prioritize exposure remediation, focusing on the 'choke points' that could prevent the highest number of critical attacks.
- ✔ APA can help security teams demonstrate to different parts of the business why security changes are required - changes that the business is not always motivated to make. A visualization of the exact path an attacker can take and its blast radius clearly demonstrates the risk to the business.
- ✔ APA is highly valuable for incident investigation, providing context that normally has to be manually correlated by pulling data from many different tools.
- ✔ APA can be used to help organizations identify how vulnerable they are to specific threats or vulnerabilities, such as ransomware, providing a view of the attack paths a given exposure enables, and the potential impact on the business.
- ✔ With access to EASM data, APA can be used to show rogue assets, shadow IT and unauthorized cloud accounts, and visualize related paths that would normally go unseen.
- ✔ APA can also provide valuable context to assess the potential impact of breaches that bypass the external perimeter or provide escalated privileges to see the potential impact of phishing, supply chain and ransomware attacks.



### Selection Tips:

To get the most complete and accurate view of your landscape and risk, look for APA solutions that leverage data from a broad range of tools and domains, such as: vulnerability management for on premises, CSPM for cloud, Active Directory or Cloud Infrastructure Entitlement Management for identities, and EASM and WAS for the perimeter. These allow for more robust visibility into attack paths and blast radius.

## 7. CONTINUOUSLY ASSESS CLOUD EXPOSURE

In steps 1-6, the focus has been on hardening our cloud security posture in order to prevent breaches. We also discussed the nature of hybrid-cloud environments as highly dynamic with new cloud assets, such as instances and containers being introduced at scale on a continuous basis, as well as new vulnerabilities and misconfigurations. To effectively manage cyber risk, exposure scoring must be recalculated on a continuous basis to reflect the current environment. Continuous risk assessment is also critical to optimize investments where they will have the greatest impact, and to demonstrate to constituents, such as C-suite and board of directors, overall risk and therefore that the potential for a breach has been reduced in a measurable way.

As we discussed earlier, an Asset Exposure Score can significantly improve prioritization, reducing the number of Critical and High Severity vulnerabilities that need to be remediated by as much as 97%. AES can also be used to understand where risk is concentrated. For example, cloud assets vs IT devices. By looking at assets that have both a 'High' Asset Criticality Rating, and 'Critical' VPR we can quickly assess what remediations will have the biggest impact. Further, looking at the average length of time these high-severity vulnerabilities remain open can help us effectively measure progress in reducing our overall exposure.

### Selection Tips:

The ability to easily aggregate data on assets, related risk, and their overall impact on cyber exposure can greatly improve the effectiveness of your [exposure management](#) program.

### Cyber Exposure Score

Another key metric is [Cyber Exposure Score \(CES\)](#). As with attack path analysis, the more we are able to aggregate data from different tools into a single data set, the easier it is to put that data to work. CES averages AES from multiple domains to calculate an overall CES, thereby quantifying total cyber risk exposure for the organization. For example, AES associated with on prem vulnerabilities, cloud misconfigurations, web application flaws and identity system exposures, are aggregated into an overarching score. The higher the score the higher the risk.

CES allows organizations to understand their cyber risk at any given time, understand how their cyber risk is trending over time and provide event details and explanations as to why the scores have changed. Security teams can visualize CES according to business applications or services, business processes or functions, geographic location, operating environment and more. This helps security leaders better communicate to non-technical business leaders using clear and concise cyber risk metrics to demonstrate effectiveness and overall process improvement.

CES for peers in the same industry, and CES that crosses all industries can then be compared to the CES for your organization. These three scores also serve as a benchmark from which to measure continued progress as well as providing an effective means of aligning perceived business risk appetite with measurable cyber risk exposure. CES can also be calculated to see other trends: For example, CES by region or by cloud provider, giving additional context on where investments can have the biggest impact.

## ABOUT TENABLE

Tenable® is the Exposure Management company. Approximately 43,000 organizations around the globe rely on Tenable to understand and reduce cyber risk. As the creator of Nessus®, Tenable extended its expertise in vulnerabilities to deliver the world's first platform to see and secure any digital asset on any computing platform. Tenable customers include approximately 60 percent of the Fortune 500, approximately 40 percent of the Global 2000, and large government agencies.

Our goal is to arm every organization, no matter how large or small, with the visibility and insight needed to answer four critical questions at all times: Where are we exposed? Where should we prioritize based on risk? Are we reducing our exposure over time? How do we compare to our peers?

Tenable can help you gain comprehensive visibility into your attack surface so you can stay one step ahead of attackers. To learn more about how Tenable can help: [View Infographic](#)