

The Evolution of Cloud Security

To Real-Time CSPM and Beyond



Table of Contents

03

CHAPTER I **See what others don't**

CSPM was born to tackle
cloud misconfigurations

Agentless scanning is
a snapshot of your risk posture

CSPM can be noisy:
Context provides clarity

Real-time risk visibility
for effective prioritization
and remediation

10

CHAPTER II **Prioritize with context**

Mind the cloud security gap:
The need for real-time visibility

Common CSPM misconceptions

Prioritize, prioritize, prioritize:
The power of the context

Modern-day CSPM:
Focusing on what matters

CHAPTER III **Remediate risks faster**

Coming Soon!

CHAPTER I



See what others don't

Cloud Security Posture Management (CSPM) has had quite a bit of discussion around what it is, what it does — and doesn't — do, why it's an essential part of a cloud native security strategy, and why the time has come for its evolution.

The rapid growth of hybrid and multi-cloud environments has created complex architectures and makes it difficult to see what's going on across platforms. Adding to the challenges, recent surges in cloud attacks and breaches have amplified the destructive impact and put the spotlight on how teams protect and run applications in the cloud.


In this e-book, we'll discuss the role that CSPM has played in modern cloud security. We'll dive into the benefits and challenges of existing CSPM solutions and discuss how only a real-time CSPM solution can provide:

- **Complete risk visibility across multi-cloud environments**
- **Contextual insights and better prioritization of risks**
- **Faster remediation to dramatically reduce critical risk**



“ Through 2025, 99% of cloud security failures will be the customer's fault.

Gartner



CSPM was born to tackle cloud misconfigurations

Misconfigurations top the list of security threats in cloud environments and are one of the most preferred launching pads for cloud-based attacks.

CSPM was developed initially to address misconfigurations in cloud infrastructure, which are often caused by human error. Many organizations began their cloud security journey by adopting CSPM as it first emerged to address basic cloud compliance needs such as:

- **Discovery of cloud configuration issues**
- **Remediation of issues**
- **Reporting and auditing to demonstrate compliance**

CSPM solutions are evolving, and today they're designed to enable organizations to automatically discover, assess, and remediate security issues, providing risk assessment across not only the infrastructure but also running workloads.




Agentless scanning is a snapshot of your risk posture

Most modern CSPM solutions offer an agentless scanning component. Agentless scanning technology works by taking snapshots of running workloads using the cloud providers' APIs and scanning them for issues. It provides fast, automated discovery of the resources in your cloud account. It provides quick visibility into cloud workloads and risk posture management, while detecting some risks, including misconfigurations and vulnerabilities.


Agentless scanning allows you to:

- **Automatically discover and map all your cloud resources**
- **Get basic visibility into workload risks**
- **Rapidly prove compliance**

The benefit of an agentless solution is that it provides workload scanning to quickly assess your basic risk posture, but it's only a point-in-time **snapshot** of the total picture. To understand your real risk exposure, you need to go beyond periodic workload scans and gain real-time visibility into in-workload risks.



Ghost containers are containers spun up for a limited time. Typically, they're spun up after the snapshot is taken and removed before the next snapshot is scheduled. They're ghosts, leaving a mark on your build, but you can't see what they did because they disappear before an agentless solution can capture them. They're the perfect example of what can be missed with a point-in-time snapshot.



CSPM can be noisy: Context provides clarity

Advanced threat actors see the security gaps and are increasingly targeting environments to steal data or disrupt business operations using state-of-the-art, custom-made malware that's often undetectable by agentless solutions.

So, while traditional CSPM solutions provide partial insight into your cloud native environments, there are certain limitations, specifically around:

An overwhelming amount of noise


Cloud environments are maturing, becoming too complex and too large. This results in too many alerts and an overwhelming number of findings that represent both low effective risks (noise) and real risks that need to be remediated (threats). Organizations need a solution that will cut through the noise to focus on real issues.

Lack of visibility


For everything that an agentless CSPM solution can see, there's much that it can't — especially with today's threats, such as in-memory attacks, transitional containers, and evasive behaviors.



Increase in
fileless attacks



Security pros
experience alert
fatigue



cloud native attacks
evade agentless
detection

These limitations contribute to:


- **Inconsistent and conflicting data**
- **Lack of context to filter and prioritize**
- **Difficulty in operationalizing response**

Ultimately, this creates a “Groundhog Day” effect where there’s no improvement in the enterprise’s cloud security.

Adding a lightweight eBPF-based sensor provides in-workload visibility. It’s non-intrusive, using less than 1% of CPU to pinpoint memory-based attacks, unknown malware, and unknown exploit attempts such as zero-days. Its behavioral detection capabilities catch threats that other solutions can’t see.

Context is key

Log4J is a known vulnerability, but depending on how it’s used and where it is in a build, it may not pose a high risk – or even a risk at all. Yet many alerts are sent out to flag it as an issue. Could it be one? You won’t know if it’s a real issue unless you know where it lives and what it’s being used for.



Real-time risk visibility for effective prioritization and remediation

Only Aqua's Real-Time CSPM combines agentless and in-workload scanning to detect risks accurately and catch threats that other solutions miss, such as memory-based attacks, unknown malware, and zero-days.

It delivers a complete and prioritized view of the cloud security risk in real time across cloud workloads and infrastructure, delivering deeper insights and uncovering threats that evade agentless detection.

Once you have full visibility into what's happening in your environment, then what? Knowing something is there does little to help you remediate, especially if the alerts you're receiving don't provide you with:

- **The context of where to find the issue**
- **What the issue is**
- **The urgency with which it should be resolved**



Aqua Nautilus honeypot data across all attacks over the last six months of 2022 discovered more than 700,000 attacks, many undetectable by agentless and traditional antivirus technologies.

In some of these attacks, files were written to disk and then deleted, and their history wiped out.

63% of the attacks were known malware. Solutions that look only for identified malware would have failed to detect 37% of the attacks.

50% of the attacks included a masquerading technique, such as a file executed from /tmp, and obfuscated files or information, such as dynamic loading of code.



CHAPTER II


Prioritize with context

Mind the cloud security gap: The need for real-time visibility

Today, developers and DevOps teams can spin up cloud resources with a click of a button and deploy applications as often as a few times a day, making it hard for security teams to know what's happening in their cloud environment.

On the other hand, the multitude of siloed scanners and other tools generates so many alerts that make it impossible to sift through all of them manually to separate the wheat from the chaff and determine what's critical and what's not. This profound gap desperately calls for the next evolution in cloud security.

Traditional CSPM is not enough for today's rapidly evolving cloud environments. Modern security teams demand better cloud security. They need complete, real-time visibility and deep cloud context, which are the foundation for rigorous risk prioritization across multiple clouds and environments. Teams need to be able to cut through the noise, discern the most critical and significant issues, and remediate them rapidly.



Common CSPM misconceptions

Over the past few years, many security practitioners were led to believe that agentless snapshot scanning is true cloud security. Agentless security solutions deliver quick visibility, basic compliance, and posture management and are great for getting started with cloud security. However, relying solely on agentless CSPM solutions inevitably leads to gaps in your security coverage. There are several reasons why this is the case:

Point-in-time visibility

Agentless scans usually run once per 24 hours, so they show the lay of the land for that specific point in time when the scan was done. The rest of the time, you're running blind and have little to no insight into what's happening in your environment. Since cloud attacks today happen in seconds and target ephemeral workloads that quickly come and go, 24 hours is too big of a gap to afford.

No real enforcement

Agentless solutions work by taking a copy, or snapshot, of a disk image, so they're not looking at the actual running code. In fact, once the snapshot is taken, they have no connection of any kind to the running workload. If they can identify an attack from a disk image snapshot, they can only alert about the problem. There's no mechanism for interdicting the attack.

Sophisticated fileless techniques

The cloud native threat landscape is constantly evolving. Attacks are getting increasingly sophisticated and often use fileless malware to evade detection and leave no footprint. Agentless solutions miss such sophisticated threats because they can't see the process running in memory from a static disk image. Again, another blind spot.

All these visibility gaps make the lives of security practitioners incredibly hard because they're unable to see the full picture and identify all the risks in the environments. Hence, they can't prioritize and remediate those risks efficiently to reduce the attack surface.

In February 2023, Aqua's security research team, Aqua Nautilus, detected a global campaign attacking Redis servers with custom-made fileless malware called HeadCrab

[HeadCrab: A Novel State-of-the-Art Redis Malware in a Global Campaign >](#)

Prioritize, prioritize, prioritize: The power of the context

In many organizations, security teams are historically understaffed and overstretched, and their time and resources are limited. What's more, in the current economy, they're asked to achieve better security outcomes with tighter budgets. At the same time, they're being bombarded with thousands or millions of alerts every day, leading to severe alert fatigue and burnout.

In this environment, trying to chase down and fix every single issue and wasting time on low-fidelity alerts is not feasible. Moreover, the inability to distinguish what's most important makes it more likely that a real, critical detection will be missed.

The only way to overcome this is to gain deep and detailed context-based insights into your specific environment. What do we mean by context here? It's many different factors or conditions in your environment that affect the issue in question and make it more or less critical in terms of your overall security posture. In short, it helps you understand how large the impact of a given issue is.



con·text

the circumstances that form the setting for an event, statement, or idea, and in terms of which it can be fully understood and assessed.

Reference: Oxford Languages




Let's take an example. You got multiple alerts about instances of the Spring4Shell vulnerability in your environment. Because the exploitation of this vulnerability — like many others — relies on several dependencies and contextual factors that should be present in your system, how do you identify which of the instances could actually be exploited? And even if you could identify them, you can't take care of all of them at once, so how do you know which ones to tackle first?

Not all alerts are created equal. That's why you need to understand which of the issues pose a greater risk and are more important. To determine this, you need to consider various factors:

- Which of the CVEs have exploitable system configurations
- Which of them have remote exploits in the wild
- Which of the packages are located in production workloads and are actively in use
- Which of these running workloads are business-critical
- Which of them are internet-facing
- Which of them have excessive privileges
- Which of them have access to sensitive data or secrets

Spring4Shell is a critical vulnerability in an open source spring frameworks Java-based core module that can be used to trigger remote code execution (RCE) under non-default circumstances

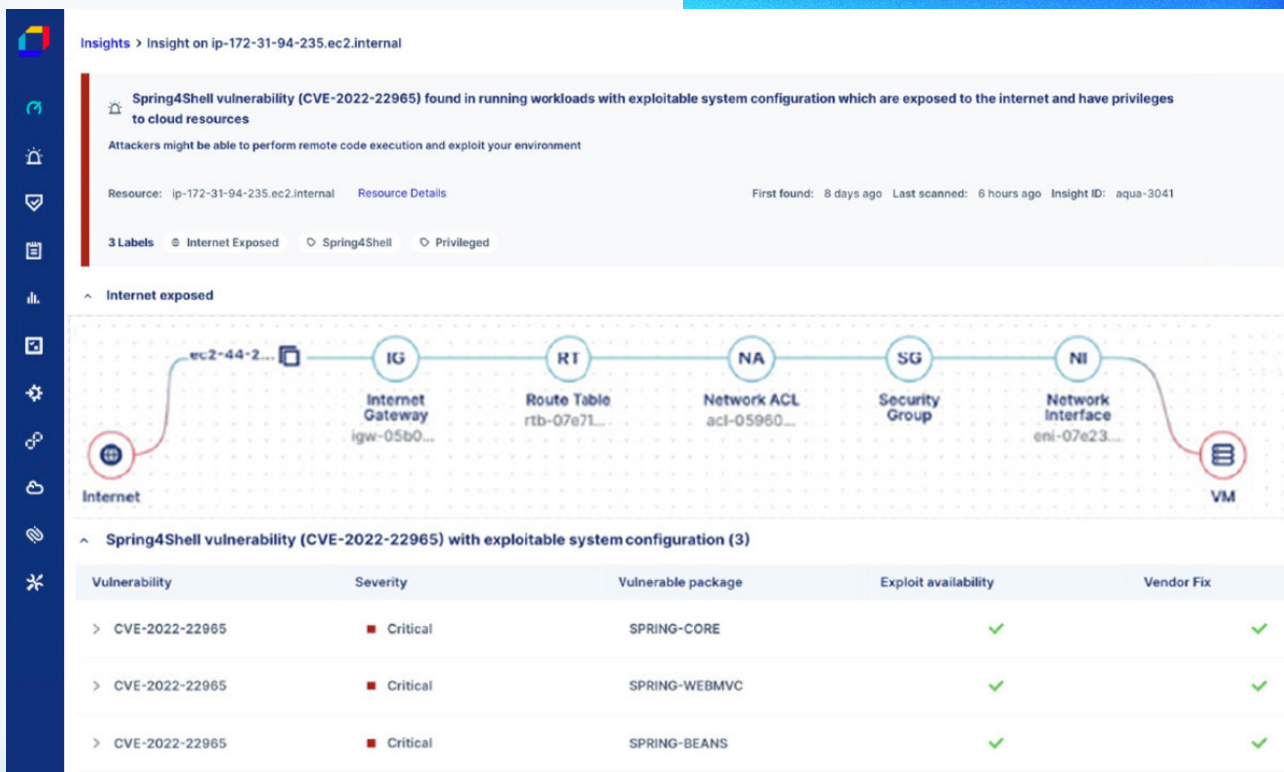
New Zero-day RCE Vulnerability
Spring4Shell: What You Should Know >



This additional context helps you significantly reduce risk and narrow down the list to only a few critical cases that could be exploited effectively and that have the biggest impact on your environment.

To sum up, a context-based approach brings a lot of value to cloud security teams:

- Allows you to understand the impact of an issue
- Determines dangerous combinations of risks
- Helps you identify top issues to make prioritization easier
- Reduces noise and helps eliminate false positives



Insights > Insight on ip-172-31-94-235.ec2.internal

Spring4Shell vulnerability (CVE-2022-22965) found in running workloads with exploitable system configuration which are exposed to the internet and have privileges to cloud resources

Attackers might be able to perform remote code execution and exploit your environment

Resource: ip-172-31-94-235.ec2.internal [Resource Details](#) First found: 8 days ago Last scanned: 6 hours ago Insight ID: aqua-3041

3 Labels Internet Exposed Spring4Shell Privileged

Internet exposed

Diagram illustrating the network path from Internet to VM:

```

graph LR
    Internet((Internet)) --- IG((IG Internet Gateway igw-05b0...))
    IG --- RT((RT Route Table rtb-07e71...))
    RT --- NA((NA Network ACL acl-05960...))
    NA --- SG((SG Security Group))
    SG --- NI((NI Network Interface eni-07e23...))
    NI --- VM((VM))
  
```

Spring4Shell vulnerability (CVE-2022-22965) with exploitable system configuration (3)

Vulnerability	Severity	Vulnerable package	Exploit availability	Vendor Fix
> CVE-2022-22965	Critical	SPRING-CORE	✓	✓
> CVE-2022-22965	Critical	SPRING-WEBMVC	✓	✓
> CVE-2022-22965	Critical	SPRING-BEANS	✓	✓

Modern-day CSPM: Focusing on what matters

With a context based CSPM, you get a prioritized list of insights that enables you to:

Determine your real risk exposure


The reality of vulnerabilities and other risks in the cloud is more complex and nuanced than simple black and white — vulnerable or not vulnerable. That's why it's crucial to take the context into account when you assess the severity of the issue, including the specific configurations or conditions that make it actually exploitable. Contextualized risk scoring brings this to life by calculating and assigning the level of severity (low, medium, high) to an issue based on a bunch of contextual risk factors. This means that you won't get critical alerts about the issues that don't pose a risk, helping you understand your real security posture.

Unify risk for a complete view

As organizations scale, most of them end up with complex multi-cloud architectures spread across different clouds, environments, accounts, services, and workloads. When complexity is growing exponentially, it's essential to not only identify risks in each specific environment or cloud but also to be able to correlate and match them across all your multi-cloud environments. This allows you to detect issues with much higher fidelity and get a single unified view of your risk profile.

Prioritize top risks and reduce noise

The only efficient way to reduce the attack surface in the cloud is to focus on the most critical issues. Only context-based CSPM gains a complete, correlated view of your risk posture, it then prioritizes detected risks so you can see what's most important. You don't need to dive in and analyze anything yourself, you just get a prioritized, ready-to-use list of the biggest issues that you need to focus your remediation efforts on. Ruthless prioritization paves the way for reducing noise and saving time since it removes the issues that aren't that significant or critical.



In modern CSPM solutions, real-time visibility provides full context to make risk-based prioritization faster and easier. Using real-time context such as internet exposure, compensating controls, and workload visibility, you can focus only on the issues that are critical to your organization and significantly improve your security posture.

Aqua Real-Time CSPM

Is the evolution of CSPM, which improves upon traditional CSPM by bringing deep, real-time context and prioritization to discovered issues. This allows you to rapidly identify, prioritize, and remediate the biggest security risks.

Chapter III - Remediate risks faster

In the next and final chapter, we'll delve into remediation, the last component of Real-Time CSPM, and how you can efficiently remediate security risks across your multi-cloud environments.

The Aqua Platform

The Aqua Cloud Security Platform protects the entire development life cycle from dev to cloud and back and is the industry's most integrated cloud native application protection platform (CNAPP). After expanding its platform with fully integrated Software Supply Chain Security, Aqua is the only solution with the end-to-end context to accurately identify and stop threats in any phase of the application life cycle.



[Schedule demo >](#)

