Latio

Cloud Security Market Report

2025

Table of Contents

Report Overview

As early as five years ago, many organizations' cloud infrastructure was a sea full of vulnerabilities and posture misconfigurations. Since then, organizations have matured in their ability to address cloud issues at scale and are looking for ways to optimize their programs. This report covers the past, present, and future of the cloud security market, as vendors have rapidly adapted their ability to help customers address emerging security threats.

Key Insights:

1. In the past cloud security tooling was defined by easy visibility, while the present is focused on using runtime insights to prioritize reducing real risks.
2. Vulnerabilities and misconfigurations continue to converge within asset contexts, allowing teams to fully understand their threat vectors.
3. Mapping cloud alerts back to the code that generates them allows teams to create better fix workflows.
4. Cloud security tooling is moving beyond the cloud, as vendors race to build a single tool for vulnerability management and reporting to support hybrid environments.
5. Security operations teams are being empowered with application layer insights to their running applications.
6. Teams are increasingly able to send their runtime alerts to security operations, and their vulnerability alerts directly to developer workflows.

This report offers insights into what leading vendors are building and who's leading the way to help security teams prepare for the future. Recognizing that cloud security choices depend heavily on each organization's architecture, team size, and security goals, we've designed a buyer's guide that provides clear, actionable recommendations, structured through a decision tree, to help teams design the right cloud security program for their specific needs.

Prepared by:
James Berthoty

2025 CLOUD SECURITY SURVEY RESULTS

Main Survey Insights

- Teams are focusing on reducing exploitable vulnerabilities
- Operationalizing large scale patching programs remains a priority
- Organizations need visibility into AI and applications
- CISOs are looking to get more results with smaller budgets

Voice of the Practitioner

We surveyed cloud security practitioners across organizations ranging in size from ten to tens of thousands employees. The results of the survey had several major themes that emerged across industries and company sizes:

1. Most teams are satisfied with their current tool's cloud vulnerability and posture management capabilities
2. Most practitioners are dissatisfied with their current tool's application security and workload security capabilities
3. Cloud security teams own and are responsible for their tools, but the SOC team responds to runtime alerts
4. Cloud security budgets have stagnated
5. Most teams are satisfied with the ROI of CNAPP solutions, but are willing to switch providers based on product maturity and usability
6. The majority of respondents didn't prioritize or require their CNAPP solution to have application security capabilities
7. A small number of respondents in cloud native organizations opt out of purchasing a CNAPP solution altogether to focus their budget on application security
8. The most requested features from a solution are AI visibility, ADR capabilities, and Access Management.

Recommendations for Teams That Don't Need a CNAPP

Before discussing Cloud Native Application Protection Platforms (CNAPP), we start by acknowledging that not all teams need a CNAPP. Teams with certain architectures are able to achieve cloud security outcomes without needing a dedicated tool.

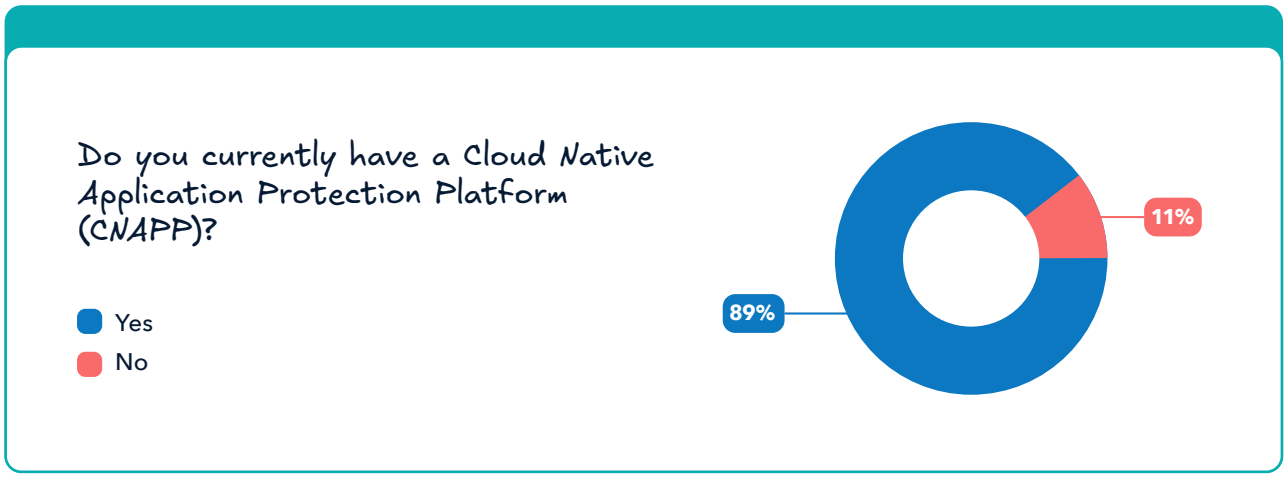

10% of survey respondents from organizations with fewer than 1,000 developers chose not to purchase a CNAPP solution. The respondents indicated that their decision was driven by the following reasons:

- The team decided to prioritize application security
- When looking at capabilities provided by current tools, CNAPPs didn't provide much additional value

If an organization has a highly standardized architecture, such as requiring infrastructure as code for all deployments, or using platforms like Vercel or Railway, CNAPPs are simply not needed.

In the above cases, we'd highly suggest using open source solutions like Prowler or Cloudquery to get simple visibility into your cloud environment. For similar reasons, vendors like Aikido, OX, Cycode, and JIT, provide basic cloud security posture management (CSPM) and code to cloud mapping, despite being application security focused tools. In these situations, survey respondents also expressed a desire for container protection, making pairing with a strong runtime offering a great solution.

The Average Cloud Security Program

Among respondents that shared which CNAPP they use, Wiz and Microsoft Defender were the most common choices, followed by Upwind, Lacework, and Orca. Despite pressures from the "platformatization" trend, it was surprisingly common for teams to combine solutions, such as pairing Tenable with Aqua, or Qualys with Sysdig. On average, a cloud security team is managing three tools as part of their cloud security program.

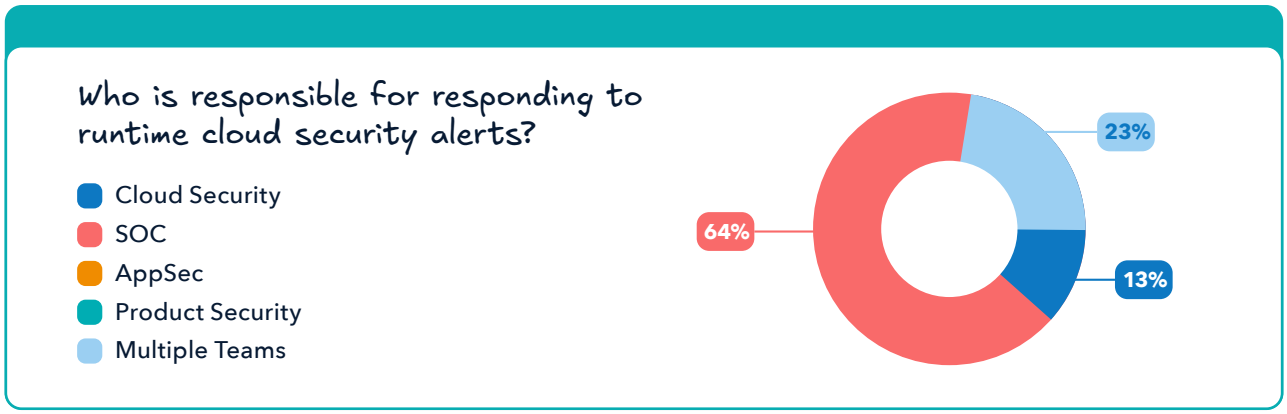

Over 70% of organizations' CNAPP tools are owned by their cloud security team, with the remainder dispersed between product security or vulnerability management teams. As we will contend later in the report, the evolution of cloud security into larger vulnerability management programs will be a continuing trend.

The Stagnant Budget Challenge

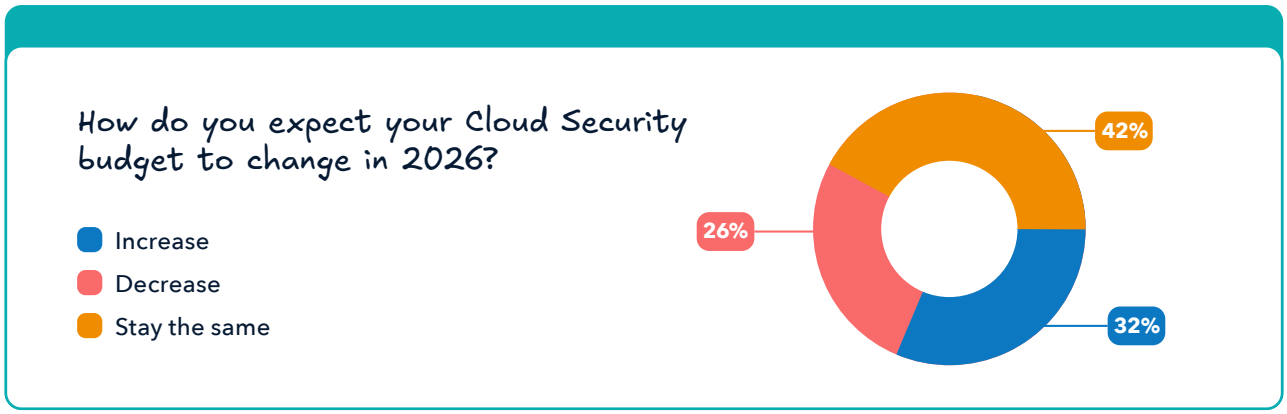

Over 65% of respondents expect their cloud security budget to stay the same or decrease in 2026, with only 30% indicating an increasing cloud security budget. As teams continue to operationalize their existing cloud security solution, there's less need to increase hiring and tool budget in this area. Furthermore, as cloud security programs have matured alongside their tooling, the budget is shifting to AI, security operations, and more traditional security concerns.

Does Cloud Security Care About Application Security?

Do you prefer to have Application Security capabilities integrated within your CNAPP solution?

- Yes
- No
- Maybe
- Will use whatever has the best testing
- Will use whatever has the best developer experience

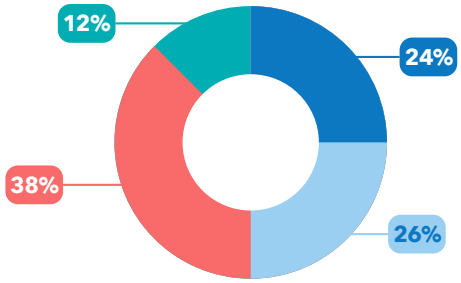

While most CNAPP's offer some kind of basic application security testing solutions, most organizations are both dissatisfied with the offering, and don't prefer having it a part of the solution anyways. 38% of respondents actively prefer having separate tools for application security, while another third of respondents indicated that they will use whatever offers the best experience to developers regardless of where the testing comes from.

Despite the broader push to product security, and vulnerability management uniting application security and cloud security from a data perspective, these teams remain distinctly focused on their own types of testing and workflows.

The Future of CNAPP Features

When asked what emerging CNAPP features teams were most excited about, four choices stood out far above the others:

- 65% - AI Posture Management. Visibility into what AI models are being used, how they're being used, and their general security posture.
- 53% - Application Detection and Response. Detecting and responding to application layer threats at runtime.
- 47% - Access Management. Helping not just detect identity misconfigurations, but also providing ways to enable secure access.
- 35% - Remediation Assistants. Helping teams patch their existing vulnerabilities.

Each of these feature requests reflect clear trends in the broader security landscape, demanding more visibility into AI and applications, alongside an emphasis on remediating vulnerabilities more than ways to discover them.

Introduction:

The Future of Cloud Security Goes Beyond CNAPP

Cloud Native Application Protection Platforms (CNAPP) have had a hold over cloud security practitioners, and it's finally time to break free from the bloated paradigm. In this report, we'll argue that the CNAPP offering is evolving into three unique categories, each with their own nuances and challenges: Application Security Testing (AST), Continuous Threat Exposure Management (CTEM), and Cloud Application Detection and Response (CADR). Each of these categories are pushing the boundaries of what's possible to maximize the effectiveness of security programs.

The definition of CNAPP has expanded year over year, as platforms continue to grow into unwieldy behemoths: striving to cover code, cloud, runtime, SaaS, and now AI security in a single solution. While having all of these security features in a single place may seem like a dream (and in some solutions it can be) the practitioner reality is frequently a nightmarish backlog of noisy "misconfigurations" with years-long backlogs to fix.

The future of cloud security tooling is moving beyond CNAPP as an "everything security" platform. Instead, teams are focusing on application security testing, universal vulnerability management, and advanced workload protection as individual focus areas, each mapped to specific practitioners. Even if an organization buys all of these tools from a single vendor, different teams and workflows are used to meet their evolving goals. To better understand this shift we'll start this report by covering the evolution of CNAPP across three generations:

- **CNAPP Gen 1,** defined by agentless vulnerability scanning and misconfiguration discovery.
- **CNAPP Gen 2,** defined by runtime insights and toxic combinations.
- **The future of Cloud Security,** defined by hybrid cloud vulnerability management and advanced workload protection.

By the end, teams should understand how the market has evolved and where it's going and what tools on the market best fit their needs. Cloud security buying decisions are guided by careful considerations related to a company's architecture in relation to their overall stack, so we end the report with a flow chart of recommendations based on a company's security posture and common sizes.

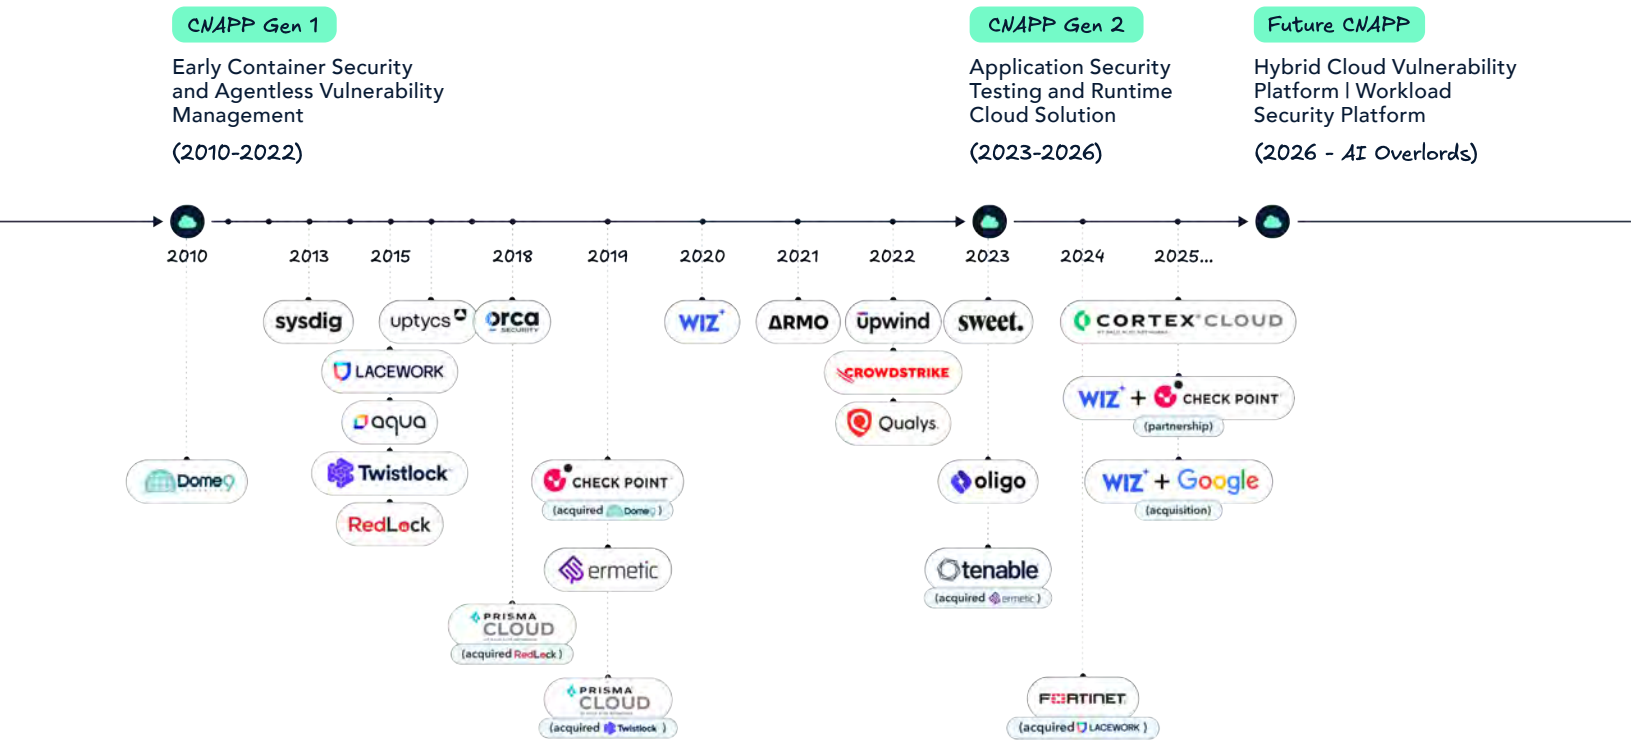

CNAPP Gen 1

Container Security and
Vulnerability Management

CNAPP Gen 1

Early Container Security and Agentless Vulnerability Management (2010-2022)

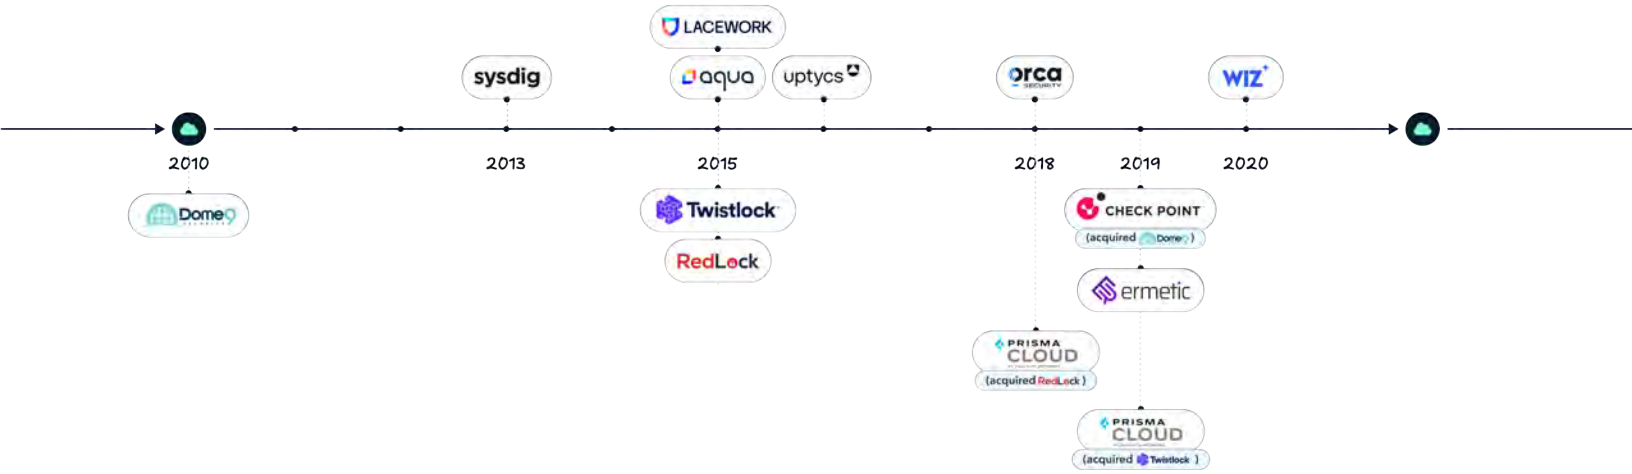

In the late 2010's, many cloud security journeys started with a decree from the CISO, "figure out what the developers have been doing with the CTO's credit card and the AWS account." Security teams have two core responsibilities: preventing security incidents, and responding to ongoing attacks. Because most security journeys started with discovery, preventing security incidents became the primary goal as they sought to discover these new environments.

These newfound cloud environments largely consisted of two types of infrastructure:

- Lift-and-shifted infrastructure from on-premise to cloud promising cost savings in data center management
- Net-new containerized infrastructure that could be widely deployed and monitored at scale with Kubernetes

Early cloud security investors and innovators saw the container revolution that was unfolding and invested in security tooling to protect this new frontier of computing. Tools like **Sysdig and Aqua** were first to market with powerful runtime insights into containerized environments. However, this would prove to be far ahead of where the market, and most security priorities were.

Security teams naturally gravitated towards securing lift-and-shifted infrastructure that more neatly fit into their established domains, predominantly static instances and storage. While established players had a powerful first-movers advantage, with the acquisition of Twistlock and Redlock, followed by the launch of Prisma Cloud, **Palo Alto brought cloud security into the mainstream.** The total addressable market for cloud security solutions greatly increased as awareness of the need for unique security solutions exploded.

As the market rapidly expanded to teams that were being asked to secure their new cloud environments, those who had invested in **agentless scanning technologies, namely Orca and Wiz, found themselves at a massive advantage.** Security teams needed a way to understand the total scope of their infrastructure, and how a list of assets related to one another.

In short, teams needed a way to turn this AWS dashboard:

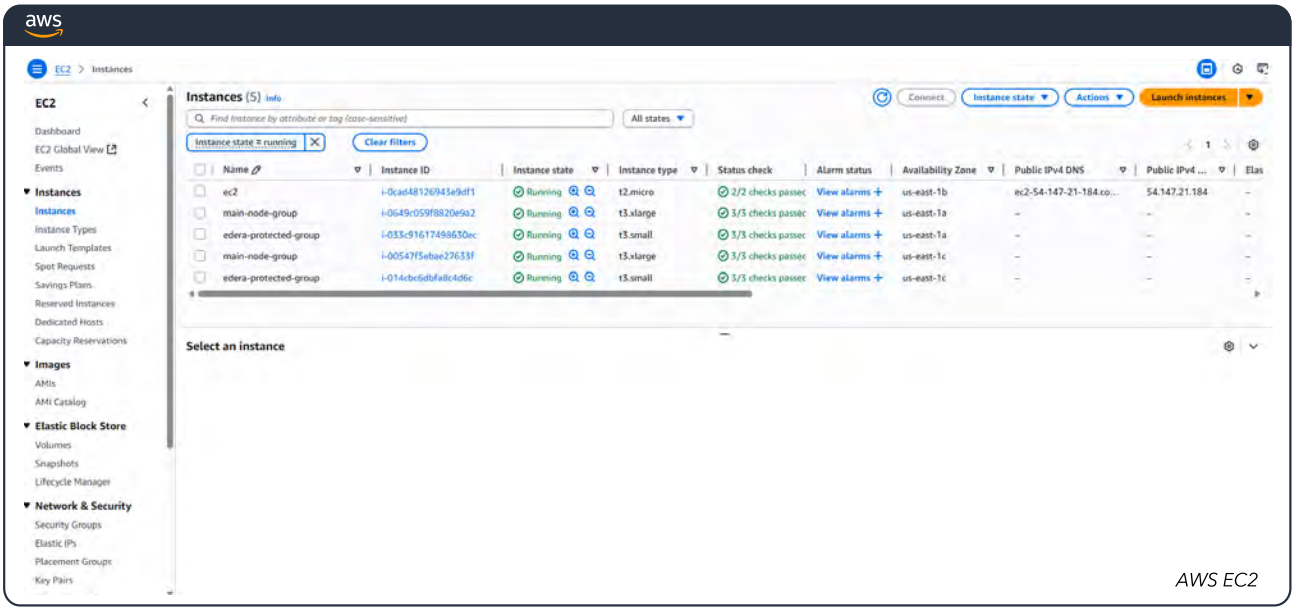

AWS EC2

Into this graphical one:

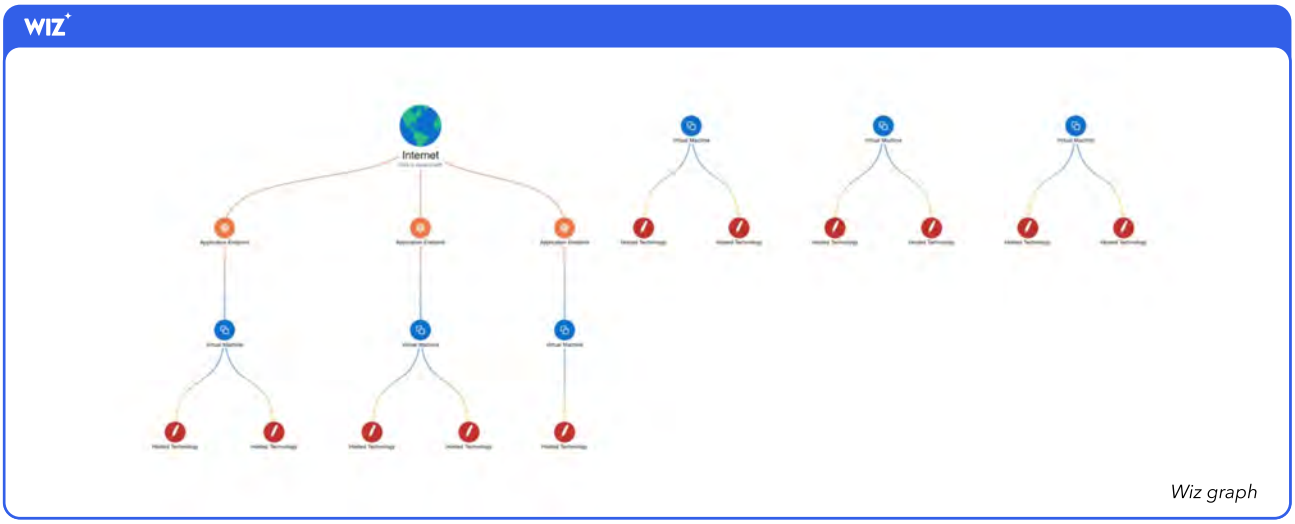

Wiz graph

While tools like Sysdig and Twistlock offered greater runtime protection capabilities than many alternatives, security teams still had to manually map asset relationships and plan endpoint agent deployments before they could begin securing even the most basic infrastructure. This was especially challenging because systems teams were learning Kubernetes and trying to standardize their infrastructure as code deployments at the same time.

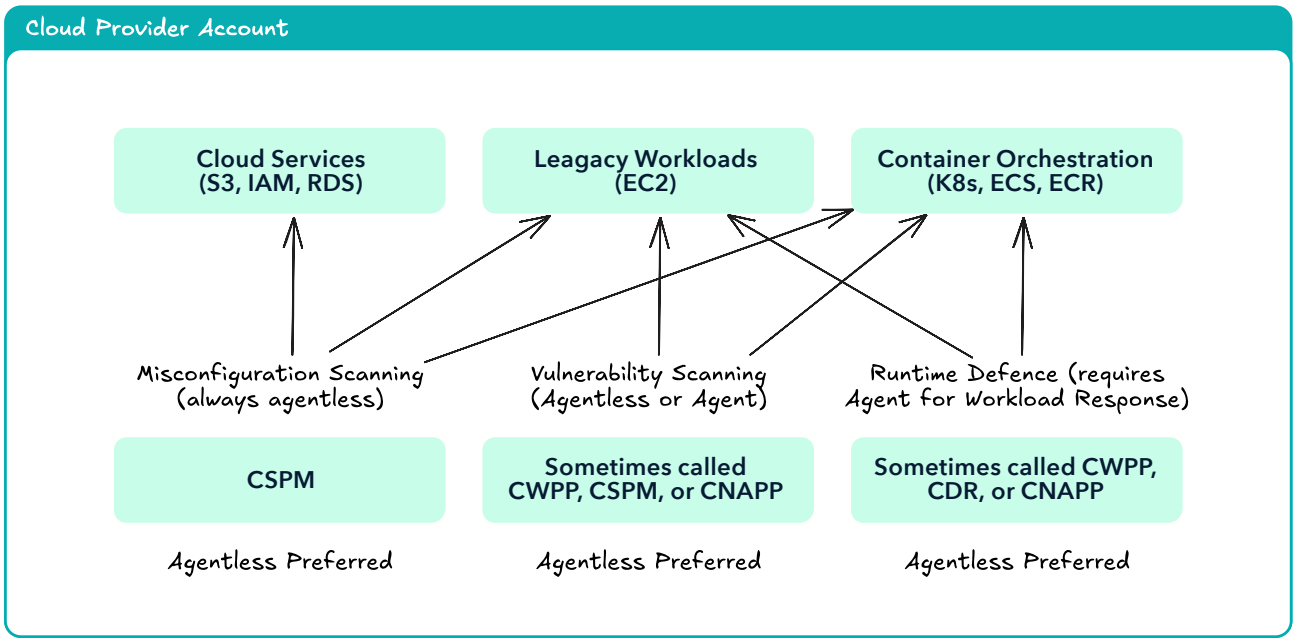

For vulnerability discovery, agentless scanning is always the preferred approach, for reasons that are self explanatory - there's no agent eating up compute, interfering with deployment pipelines, and needing endless maintenance. However, security teams would soon discover that while they could create amazing graphs with this technology, they couldn't do anything to stop an active attack in their environment.

Alongside agentless scanning, Wiz proved that building a graph of asset relationships was essential to driving security outcomes. While this was a pivotal sales tool, offering instant value, their introduction of "toxic combinations" offers immense value to teams looking to reduce their alerts from the thousands down to manageable numbers.

Agentless scanning platforms were able to grow with their security team buyers to meet market demand by deploying kubernetes connectors to map out resources, expanding into container registry vulnerability scanning, and finally introducing runtime agents as customers looked for the functionality.

The first generation of CNAPP defined the following features as necessary to providing a robust user experience:

- An asset relationship graph
- The ability to prioritize findings across discovery types (combining posture and vulnerability findings into a single discovery)
- Agentless vulnerability scanning for instant value

These three features would come to define the first generation of CNAPP, and lead to Wiz handily outperforming the competition.

CNAPP Gen 2

Code to Cloud and Runtime Protection

CNAPP Gen 2

Code to Cloud and Runtime Cloud Solution (2023-2026)

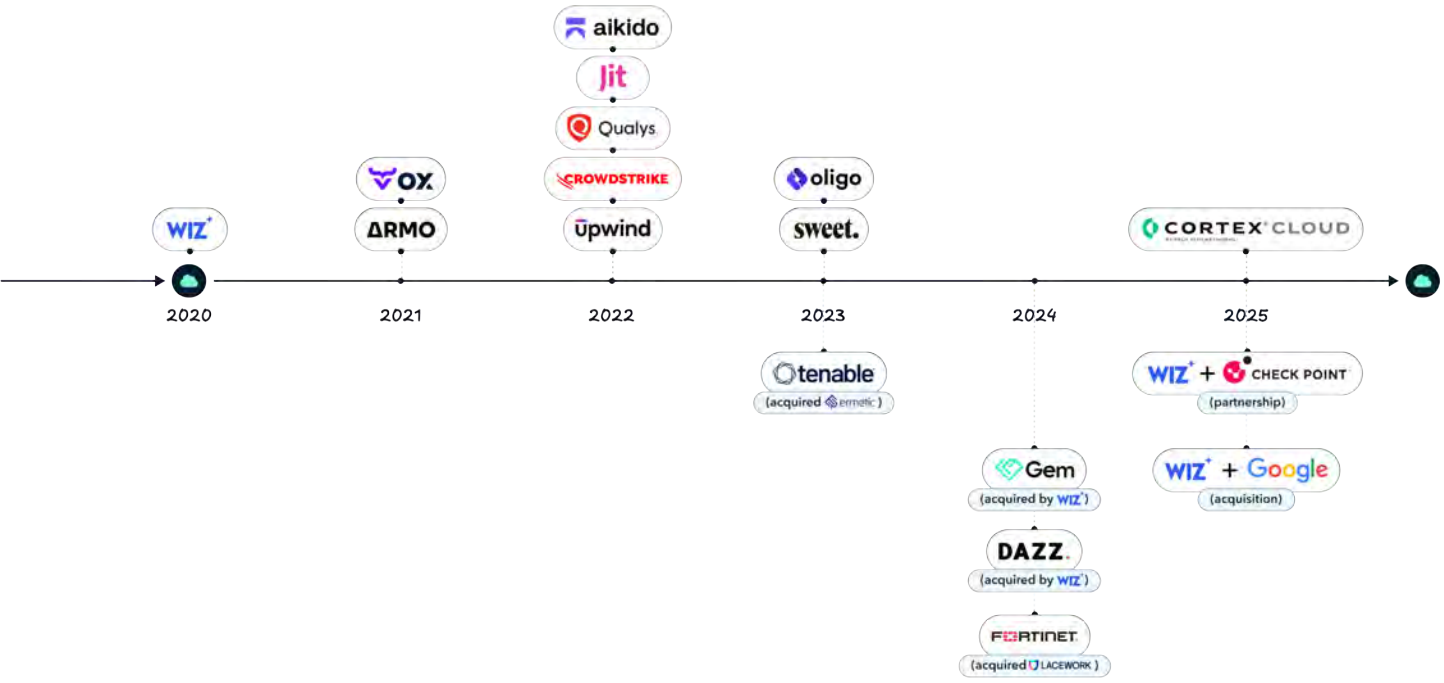

Moving to the present day, the posture side of cloud security is widely agreed to be commoditized. Amazing open source projects like Prowler exist to provide instant posture scanning value to security teams, and most cloud providers offer their own concepts of the graph, agentless scanning, and toxic combinations.

The conversation that came to define CNAPP, agentless versus agent based scanning, would unceremoniously end in 2023 when Wiz released their runtime sensor. While most of their competitors spent 2020-2023 trying to add their core capabilities, namely the graph and agentless scanning, Wiz was busy preparing for the next era of CNAPP with their Gem and Dazz acquisitions. Orca would follow shortly after with their own endpoint agent and acquisition of Opus, and end the era of the agentless CNAPP while preparing for what was to come.

No longer the domain of Orca and Wiz alone, most providers now offer agentless scanning as part of their offering. Alongside the feature becoming more common, many larger customers have realized the hidden cloud costs associated with the EBS volume cloning necessary to make agentless scanning work.

Cloud native organizations have a new set of concerns moving beyond basic visibility:

- Creating sensible remediation workflows for vulnerabilities by finding developer asset owners
- Operationalizing robust runtime protection in containerized environments
- Flexibility in vulnerability scanning methods for diverse use cases
- Building scalable IAM and access patterns for their teams

Each of these needs have led CNAPPs to expand further into new domains like application security and cloud detection and response. In fact, in 2025 it's now almost impossible to find a cloud security tool that doesn't provide an agent, or a GitHub/GitLab integration. This rush to provide features across the board has led to a market that looks like this, defined by code, cloud, and runtime capabilities:

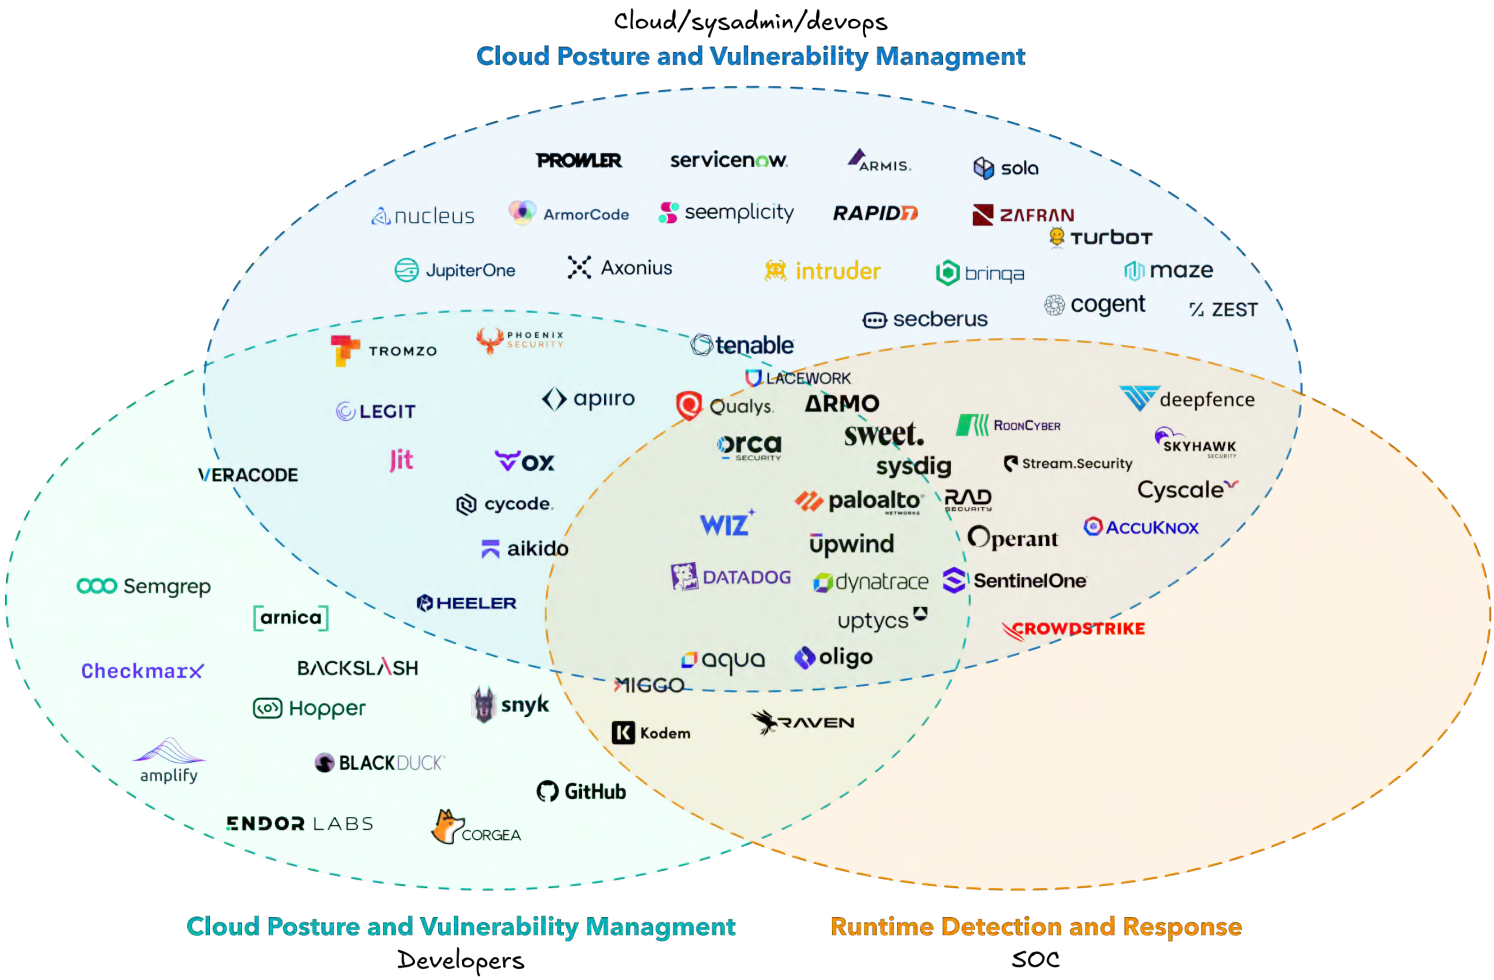

This present view of CNAPP does a few things to the industry:

1. Encourages the "best tool" being the lowest common denominator across a wide variety of use cases
2. Forces powerful emergent security capabilities to be overlooked or ignored altogether
3. Creates an inferior user experience depending on your role
4. Makes it nearly impossible to create a realistic evaluation criteria due to the number of features

Despite the confusion from a feature standpoint, two capabilities have come to define the CNAPP offering in 2025:

1. **Code to Cloud visibility,** to enable teams to accurately prioritize and fix vulnerabilities across tools
2. **Advanced Runtime protection** capabilities to respond to active attacks by expanding into the API and application layers

Code to Cloud Visibility

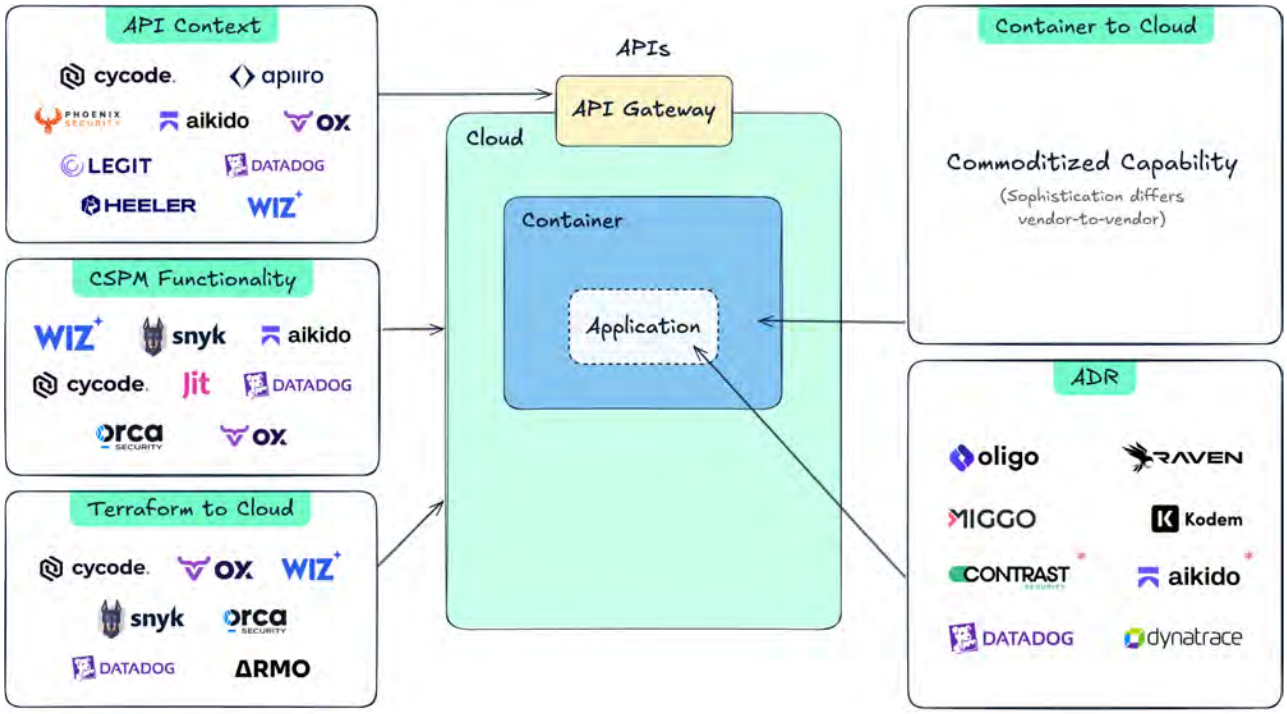

In order to properly fix vulnerabilities in production, teams need complete "code-to-cloud" visibility of their assets. The widespread adoption of infrastructure as code has forced security teams to tie vulnerabilities found in production back to the code that generated them. While teams continue to vulnerability scan in production, the developers that make the fixing changes need to understand what code needs to be changed in order to remediate a finding. In order to make these workflows more obtainable, vendors have begun helping teams automatically associate production assets back to the code that deployed them.

The first generation of CNAPP tools saw teams building complex code to cloud mappings themselves in order to get findings remediated. A combination of asset tagging, python scripts, and business analytics tools were widely used to orchestrate remediation efforts by getting relevant findings to asset owners, who would self assign their remediation tickets.

A new generation of application security tools were the first to address these concerns. One approach for midmarket companies was integrating CSPM features into application security platforms. For example, Aikido, OX and Cycode all even support agentless vulnerability scanning of cloud instances, even if it's not a point of emphasis. Over 10% of survey respondents didn't have a CNAPP, primarily because they chose to prioritize application security and didn't see the value of a standalone solution. These respondents also indicated no desire to enter the CNAPP world.

While "code to cloud" is often listed as a single feature, there's a lot of complexity to the specifics of both what's provided, what makes it work, and how it scales. I cover all of this in more detail in this article, but below is a summary of the latest capabilities.

First, there are three methods of mapping repositories to their corresponding containers at runtime:

- Mapping based on manual or dynamic searches
- Implementing a CLI tool to sign or otherwise verify where the image was built and where it is deployed
- Creating a data model using a variety of factors to create a best guess of mapping based on SBOMs, build times, tags, and other factors.

Mapping the container has become only part of the journey now, as additional mappings and capabilities are possible. The various types of code to cloud mapping are:

The Types of Code to Cloud Mapping

API context mapping

Being able to relate a piece of code to the API endpoint it's surfaced as in order to collect the kind of data being processed by it, as well as map dynamic testing issues to the code owner.

Terraform to cloud mapping

Determining what assets are actually deployed, and where the terraform that built them came from

Container to cloud mapping

Mapping container vulnerabilities back to the repository that builds the image

CSPM functionality

Highlighting application security vendors who also provide agentless vulnerability scanning capabilities

ADR mapping

Mapping runtime function level executions back to the repo that built the container

Many CNAPPs have added application security capabilities to their platforms because they're accessible to create given the integrations they already need; however, the practitioner priority split between application and cloud security practitioners remains distinct. For the most part, application security teams want their own tools to manage their own program separate from infrastructure teams. However, both teams want the outcomes from unifying the data.

Advanced Runtime Protection: CADR

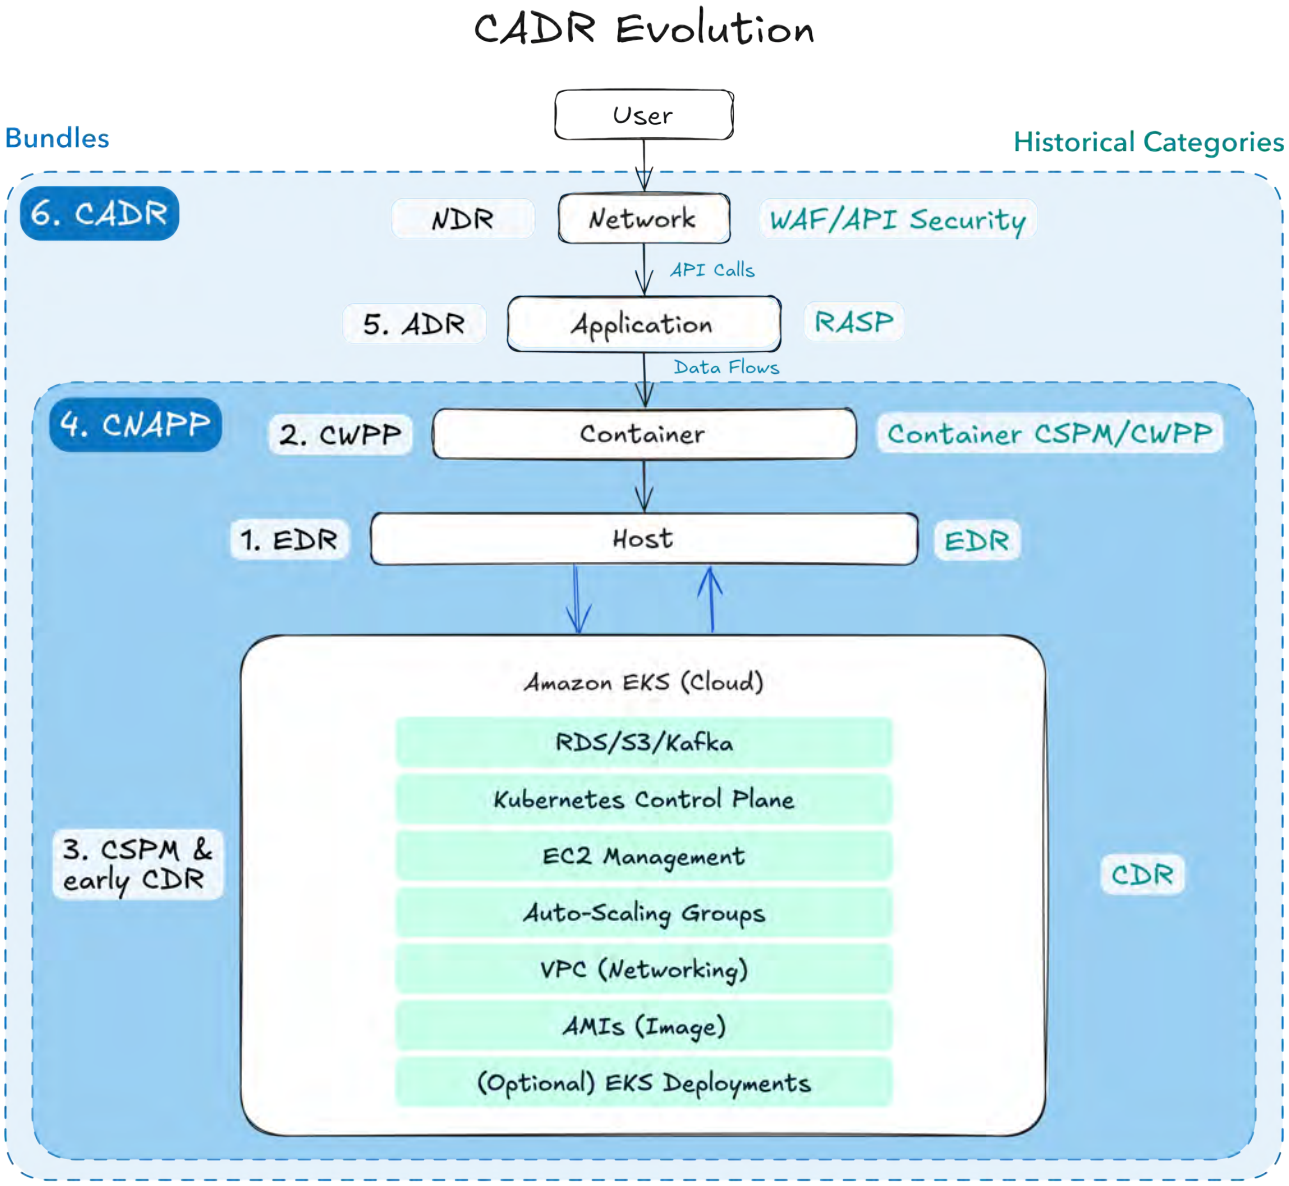

What is Cloud Application Detection and Response, and why are you pushing a new acronym? In short, CADR is the best cloud workload security capabilities you can get. For enterprises that want the best in runtime protection, expanding Cloud Workload Protection Platform capabilities across cloud and application layers is essential for operationalizing them. For mid-market customers, this likely includes traditional CNAPP capabilities, such as cloud vulnerability scanning and posture management.

Instead of the disadvantage they used to be considered, agents are now essential for providing:

- Relevant prioritization and context for containerized environments
- Threat detection across network, container, OS, and application layers
- The ability to take response actions to active attacks

Another surprising aspect of the development of CADR capabilities is the expansion of API security capabilities into CNAPP tools. For example, Wiz and Orca both now offer API discovery capabilities, while Upwind and Sweet extend those into runtime protection.

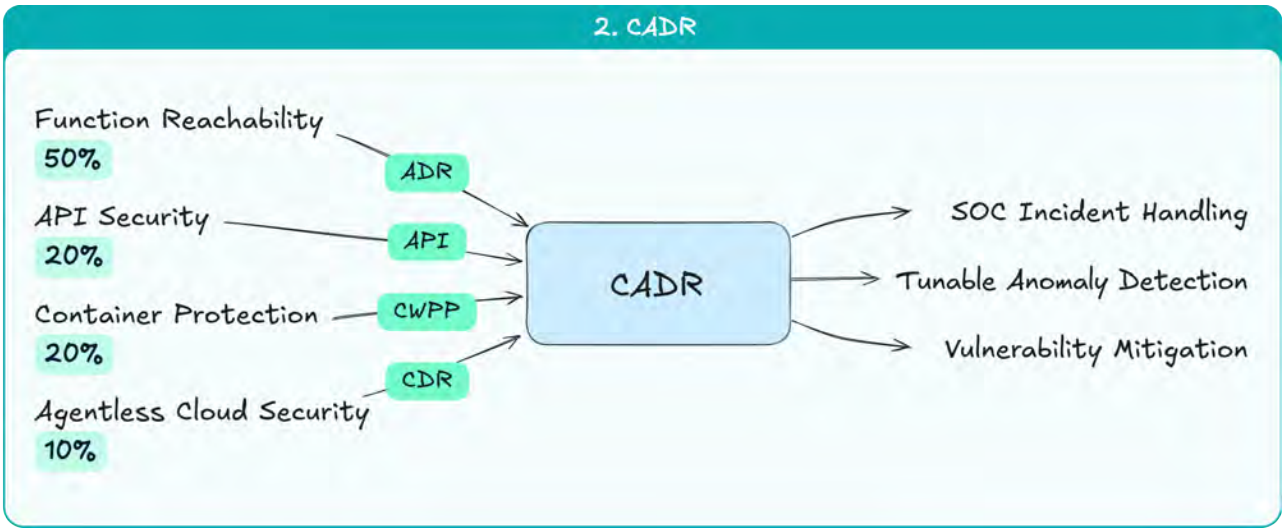

In the above graph, I rank the relative importance of each runtime feature for contextualizing cloud security runtime alerts. Additionally, this graph shows how CADR consolidates several other categories. Report after report has indicated that most cloud breaches are now beginning with web application attacks. In addition to the rise in supply chain malware, it's more important than ever that cloud runtime solutions have visibility into the application layer. Cloud Native Application Protection Platforms simply didn't pay enough attention to the Application part.

Runtime Security Map

Visibility into what application functions are executing via only eBPF capabilities remains an incredibly rare functionality. Oligo, Raven, Kodem, and Miggo remain differentiated by maturity in the capability, while other providers such as Upwind, Datadog, Dynatrace, ARMO and Sweet have all expanded early functionality here. The difference these capabilities make within a tool are massive, and their relative importance will continue into 2026.

When asked which emerging features of cloud security solutions practitioners were most excited about, over 60% of respondents chose application detection and response capabilities, among which I would eagerly count myself. Additionally, Anton Chuvakin highlighted the importance of the application layer by updating the SOC visibility triad into a quad to include the application layer. These tools have been the missing layer to enabling meaningful runtime protection in the cloud.

The Future of Cloud Security

The Future of Cloud Security

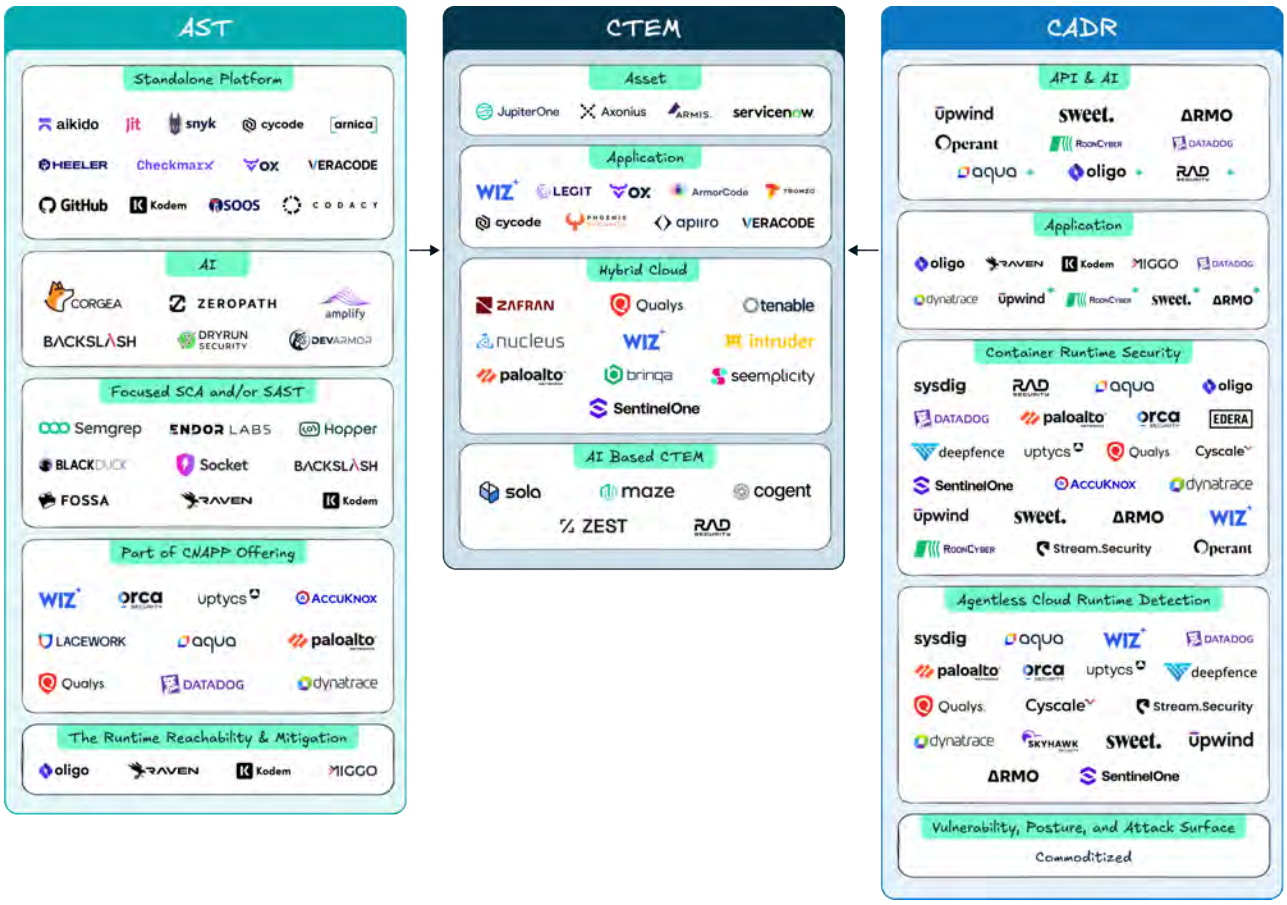

While CNAPPs do a lot, only a few features are actually widely used and deployed by organizations. A few data points make clear what people love about CNAPP, and what they don't:

1. Over 90% of respondents did not want to consolidate their CNAPP with Application Security tooling unless the developer experience was better
2. Over 50% of respondents had cloud security tooling split between posture and runtime capabilities, such as matching Orca with Aqua, or Lacework with Tenable.
3. Over 70% of respondents found posture findings to be the most valuable part of their CNAPP platform

Instead of lumping everything into CNAPP, I instead propose three distinct categories made for different personas: Application Security Testing (AST), Continuous Threat Exposure Management (CTEM), and Cloud Application Detection & Response (CADR). This separation serves two goals:

1. Accurately reflects investments into third party vulnerability management, application testing and securing AI generated code, and meaningful cloud runtime protection.
2. Allow the accurate reflection of where companies are emphasizing their work, to allow for greater consumer freedom.

The cloud is a complex mess, and distilling the buying question into "Which CNAPP should I buy?" is too simplistic. The answer depends on your infrastructure, security goals, and overall risk tolerance. In order to allow consumers to emphasize the things they care most about, greater emphasis is needed for things like "runtime flavored" and "posture flavored" CNAPP.

With those goals in mind, here are the three categories that reflect the future free from CNAPP:

AST	CTEM	CADR
Application Security Testing (AST)	**Continuous Threat Exposure Management (CTEM)**	**Cloud Application Detection and Response (CADR)**
the best developer testing solution suite of tools, focusing on developer experience, true positive rate, and code to cloud prioritizations.	an overarching vulnerability management platform, focused on routing the relevant vulnerability findings to the correct owner. Another acronym that I'd personally prefer is Universal Vulnerability Management (UVM), but CTEM is currently more widespread.	the best runtime solution to protect cloud environments, and surface runtime telemetry for prioritization and threat detection.

Continuous Threat Exposure Management

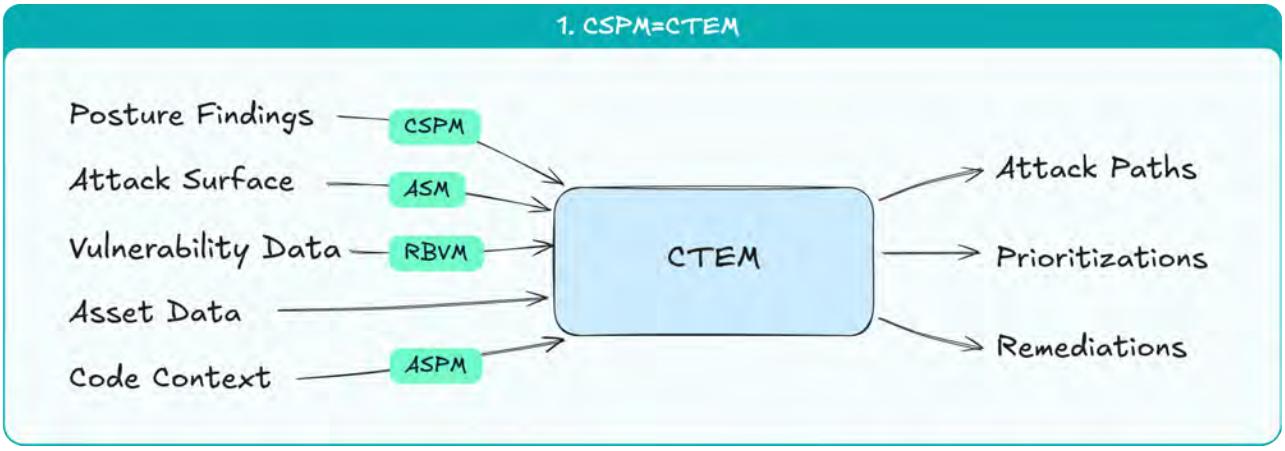

Wiz dominated the early cloud security market via their combination of the graph and agentless scanning. By combining agentless vulnerability scanning with posture findings, and creating toxic combinations for prioritization, Wiz made vulnerability scanning much more approachable. The success of this posture-first approach to vulnerability management is undeniable; however, it's time for these advancements to come to on-premise and hybrid environments.

Over the last two years there was a rapid rise and then acquisitions of a category of tools that was never clearly defined - Gartner's definition of ASPM, RBVM, Exposure Management, all mean nearly the same thing, namely a single place to ingest your vulnerability data and normalize it into a single platform. These investments have led to the future of CSPM: bringing it to your multi-cloud and hybrid environments, and normalizing data across multiple scanners.

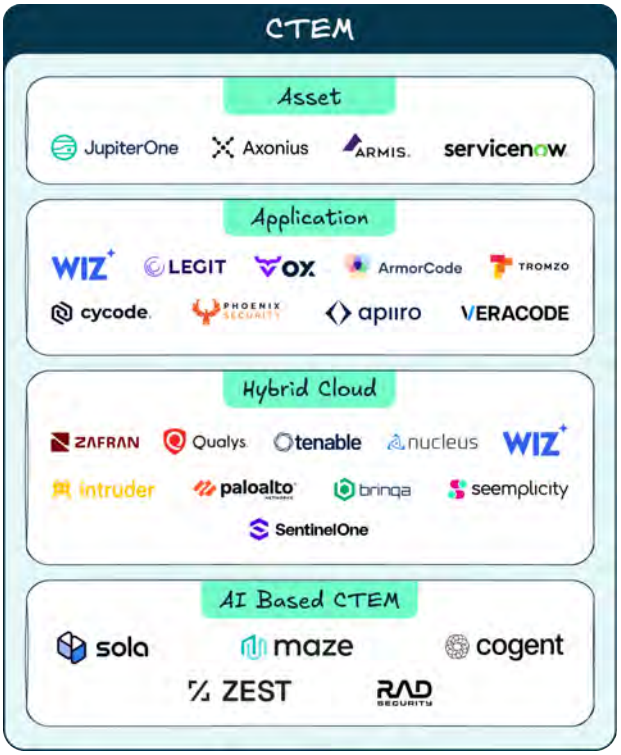

The future of the CTEM platform is a merging of posture findings, attack surface discovery, vulnerability data, asset data, and code contexts. This view of CTEM is forward thinking. No one currently does a perfect job normalizing data across posture, attack surfaces, vulnerabilities, assets, and code. However, it's clear that a group of companies aim at being the single source of misconfiguration management. Here are the five trends driving this frontier:

1. CNAPP is rapidly evolving beyond cloud, to include support for on-premise assets
2. A misconfiguration is just another kind of vulnerability
3. The vulnerability management UX is best as asset management with vulnerabilities layered in
4. Attack surface management is just the asset's posture in relationship to the internet
5. Finding the owner of cloud vulnerabilities requires knowing its code origin
6. A threat is just a buzzword for a high likelihood, exploitable vulnerability

Hybrid Is the New Cloud

While Qualys and Tenable's expansion into cloud occurred without much fanfare due to their later entrance, Wiz recently moved in the other direction with support for agent based on premise vulnerability scanning, bringing the asset graph into on premise environments. Just as the graph was revolutionary for managing cloud environments, it's poised to do the same for on-premise ones.

Increasingly, enterprise security teams need to build complicated workflows for getting vulnerabilities fixed. These tools need to span across all of an organization's assets in order to create consistent and scalable vulnerability programs. Cloud security no longer happens in isolation.

Misconfigurations = Vulnerabilities

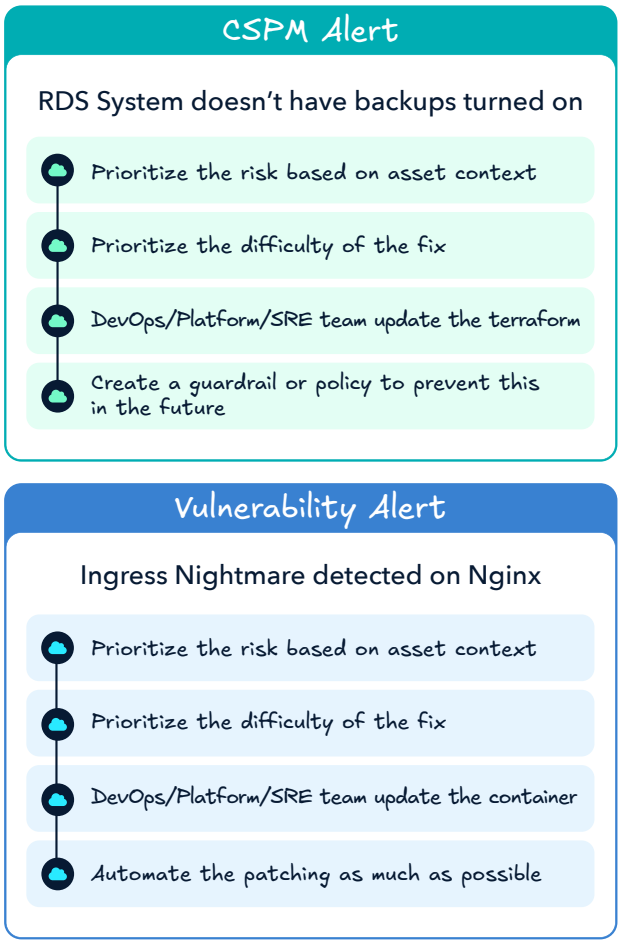

A lot of vendors in the cloud security space suffered from prioritizing misconfigurations over vulnerabilities, thinking the success of the CSPM category was due to posture alone. However, the core of the success was unifying misconfigurations and vulnerabilities together. Ultimately, fixing a misconfiguration and fixing a vulnerability involve the same teams and the same workflows, and should be treated as the same thing - a potentially exploitable risk.

The underlying data, namely the asset context and prioritization metrics, are essential for both alert types, and most teams treat these as the same even if they're technically segmented. On an ongoing basis, posture and vulnerability findings are remediated based on impact and likelihood, as the same team has to prioritize both the issues.

Ultimately, posture findings and vulnerability findings require the same workflows to remediate, and open the same risk of exploitability to the organization.

Attack Surface Management = Asset Exposure

For a long time now, some amount of external asset testing has been a feature of CNAPP. Wiz was the first to take a screenshot of the asset's login page, and in the time since, most CNAPP providers will attempt to validate open ports and share probing responses with you. The line between this basic external testing and DAST continues to evolve over time as well, especially as more vendors add API discovery, testing, and runtime protection capabilities into their platforms.

Basic Attack Surface Management (ASM) features in CNAPPs have always had two weaknesses despite demoing well: first, they tend to do a poor job with API gateways and modern application architectures, requiring an agent based approach to understand microservice communication. Second, CNAPPs usually stop short of doing full DNS enumeration to discover all of the attack surface - typically due to not guaranteeing access to the DNS provider.

Over time, these capabilities will combine to create end-to-end discovery and testing of all of an organization's endpoints.

AI and Cloud Security

In the current market, cloud security practitioners are being asked what AI models their applications are using. As time goes on, security leaders are looking to move into stronger AI governance for their organizations. To account for these use cases, most CNAPP platforms have deployed “AI-SPM” capabilities to detect what models are running on various instances, and where your traffic is going. These features have quickly been adopted into various platforms.

Looking to the future, AI capabilities will undoubtedly change the way that teams do prioritization and remediation, but both of these capabilities will be in the framework of CTEM. AI CTEM startups like Maze and Cogent are leading the way with agentic investigations of vulnerability exploitability, while Zest represents the focus on remediation workflows.

Conversely, there’s the ability to protect AI applications in the cloud, for which CADR tooling is especially well equipped. By combining deep insights into application data through network traffic and internal function monitoring, these tools easily extend to protect AI use cases. Vendors such as Operant, Miggo and Oligo have shown the ability to push their runtime CADR offering into the AI use case. Even still, others like Aqua and Sweet can even do runtime AI security detections as part of their overall CNAPP offering.

EMERGING CAPABILITIES

Emerging Capabilities that Belong in Cloud Security

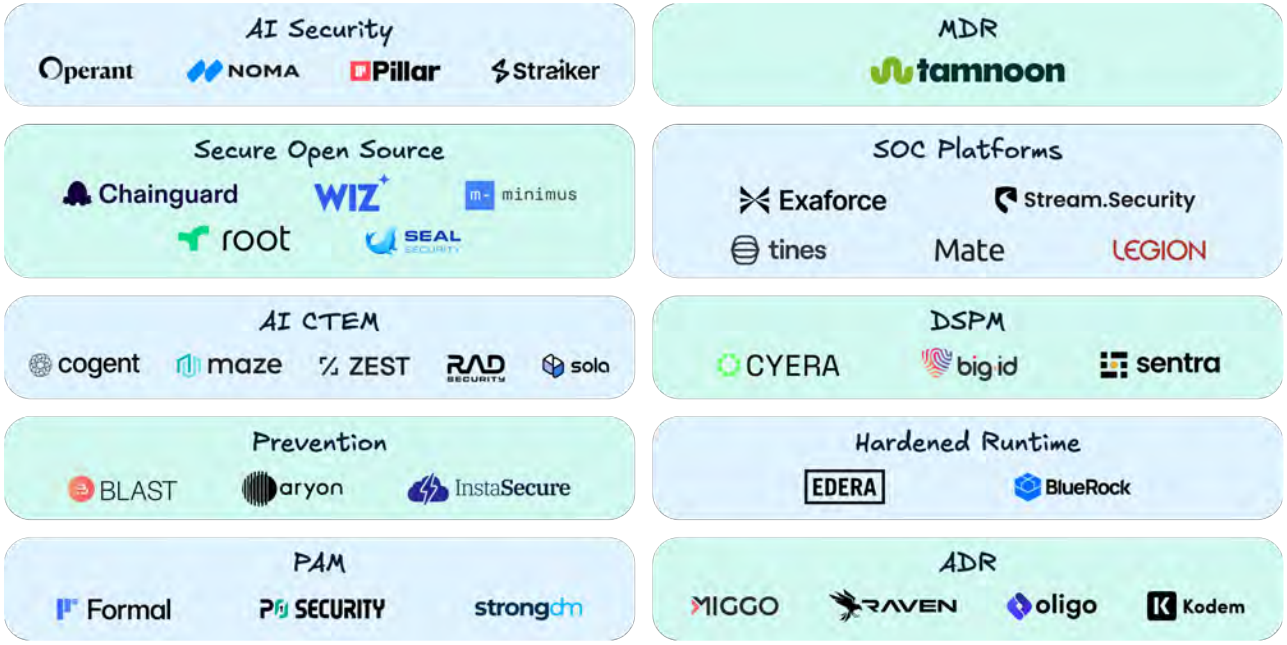

Representative vendors

Alongside the major categories we've discussed, there are a lot of innovative emerging capabilities for cloud security teams to be aware of. This section of the report elucidates those developments, and how they relate to cloud security professionals.

AI Security

Most CNAPPs have evolved to include some level of "AI-SPM" features. In reality, these amount to little more than basic visibility into what models are being used in your environment, usually in terms of what is being run on your workloads. More runtime centric providers will include actual egress network traffic in order to get some sense of where your traffic is going, and the kinds of prompts that are running.

For more in depth solutions, specific AI security vendors are usually needed. Most of the acquisitions that have taken place this year have focused on the enterprise side of the house, monitoring user traffic with browser plugins. For this report, we wanted to highlight some vendors most applicable to cloud security teams, those that get more into securing the workloads themselves. Pillar, Straiker, and Operant all provide robust runtime application security solutions for AI. Noma does this as well, but offers a posture management suite of features that will resonate most closely with cloud security posture capabilities. Finally, of the CNAPP solutions, Aqua is leading runtime AI detection for prompt injection style attacks.

For an in depth look at the AI Security Market, read Latio's 2025 AI Security Report.

Secure Open Source

Chainguard's success providing minimal container base images and operating systems has led to a rise in competitive offerings, each with their own differentiators. Minimus for example allows for CI/CD workflow creations, Root and Seal backport open source patches as part of their offering, Wiz offers migration insights via their container insights, and Rapidfort focuses on reducing the attack surface of existing images. Meanwhile, Chainguard continues to evolve their supply chain scope to include secure supply chain package registries, as well as backporting patches as some of their competitors do.

These solutions should be viewed as a great security maturity step, but not as fundamentally solving the container vulnerability problem, as regular redeployment of these base images is still the root solution to patching containers. No matter what solution you use, regular redeployment of images is the only way to get upstream patches; however, using these solutions is a great way to offload a lot of the engineering work building and maintaining minimal images, or mirroring open source artifactories in house.

AI CTEM

A few recent startups have taken an AI first approach to the CTEM challenge by focusing on vulnerability exploitability using agentic testing. Cogent has taken an asset first approach to the problem, Maze has focused on proving vulnerability exploitability, and Zest has focused primarily on remediation. Each of these companies offer some very cool early capabilities at drastically reducing vulnerability counts by allowing agents to fully understand vulnerability contexts, and offering relevant mitigations or remediations.

Sola has taken an innovative approach to the posture management paradigm as a whole, by offering on the fly AI application creation for monitoring the posture of cloud and SaaS environments. Rad security has expanded to offering an AI first approach to gaining insights to your environment by focusing on accentuating their runtime context with agentic workflows.

Prevention

Another wave of startups has focused on a preventative approach to cloud security. This is a great framing of a set of capabilities involving IaC guardrails, service control policies, and resource based access control. These tools help teams adopt best practices to prevent new misconfigured resources from coming up. In the long run, it's our opinion these tools will shift into the remediation realm, as ultimately the lack of these guardrails is another kind of posture finding and solution. In the short term though, these tools are great for teams struggling to implement guardrails in their cloud environments.

Privileged Access Management (PAM)

Okta's recent acquisition of Axiom is a great step forward for the promise of better PAM for cloud. Additionally, a surprising 50% of respondents stated access management was a high interest feature for the future of their CNAPP tooling. Developer PAM has long been an overlooked category, as even larger providers like Teleport and StrongDM haven't reached critical adoption. These new waves of PAM tools focused on developer experience offer a lot of promise for the future of cloud access by providing easier just-in-time access, access reviews, session management, and threat detection.

Managed Detection Response (MDR)

Managed Detection Response providers have long struggled with cloud environments. On the one hand, runtime alerts require a specialized application layer skillset that most don't have. On the other, posture and vulnerability remediation requires being tightly integrated with a company's DevOps workflows. While some providers like Upwind and Wiz have begun offering their own managed services to compensate, others like Tamnoon are leading the charge when it comes to managed cloud offerings.

SOC Platforms

The transition to cloud security left a lot of security operations teams behind. Increasingly however, the SOC has adapted to incorporating cloud security practices, with over 50% of respondent's alerts going to the SOC instead of dedicated cloud security teams. A future Latio report will go into security operations solutions in more detail, but several advancements are worth cloud security practitioners being aware of.

First, there's platforms like Stream Security that have focused on providing cloud context to amplify the value of existing workload security solutions. By integrating with endpoint agents like Crowdstrike or SentinelOne, Stream provides real-time cloud visibility that's often missing from these other solutions.

Second, there's the massive developments in AI. We highlighted Tines to show the value of AI amplified SOAR technologies, of which most "AI SOC" startups would cleanly fit. Two other approaches however are larger data platforms amplified by AI, and AI browser techniques to automate repetitive tasks, and serve as a copilot for security operations users. Exaforce and AI Strike are two examples in the data platform space, while Mate and Legion are examples of the browser automation approach.

DSPM

Basic DSPM functionality has become widely available in CNAPP platforms, namely through broad detection of sensitive data in underlying systems. These capabilities provide basic visibility into types of information stored on a device - like highlighting if an EBS volume or S3 bucket contains PII or PCI data. A few platforms take this a step further by providing insights into proper databases, showing which tables contain sensitive data. Some others also provide file integrity monitoring as part of their endpoint agents.

Dedicated DSPM solutions however, remain distinct for their ability to discover and determine data context, and govern access to sensitive data across not just cloud, but also on-premises environments as well. They have more robust data tagging capabilities, and are evolving into wider Data Security platforms that include DLP workflows, identity & access, and data AI governance services. These capabilities also tend to expand outside of base volumes and into SaaS providers such as Snowflake and Databricks, where most enterprise data engineering takes place.

The buying decision here on investing into a dedicated DSPM depends largely on the businesses' initiatives around data visibility, data sprawl reduction, access governance, loss prevention, and securing AI adoption - across their entire data environment. CNAPP tools are

clearly fitted for cloud security teams, with the data detection being primarily used for vulnerability prioritization purposes. Conversely, DSPM solutions tend to be built to enable data security, compliance, privacy, data engineering, and IAM teams to improve data security, inform access governance and ensure compliance standards are in place.

Hardened Runtime

An exciting emerging category is "hardened runtime." These solutions provide the deepest workload security capabilities possible by baking into the underlying operating system, offering granular security controls into the environment. While in older style environments this would be too much to ask, Platform teams usually don't care what operating system runs their underlying nodes - creating a massive security opportunity

A great example of this is highlighting Edera, who built a true container virtualization system to enable actual container isolation. This ends a massive category of vulnerabilities involving escaping container contexts in order to expand the attack surface. They're also expanding these capabilities into runtime GPU protection, enabling GPU workloads to also be isolated from one another, a major issue for securing AI usage. Another company in this category is BlueRock, who bake security capabilities into the underlying operating system.

Application Detection and Response (ADR)

In order to highlight the vendors with the most advanced function level reachability, it's worth separating out the dedicated vendors for mitigating application vulnerabilities at runtime. Oligo and Raven provide deep function level insights via eBPF which can be used for the maximum reduction of false positive vulnerabilities possible, and critically stop application exploits without breaking the application. Miggo provides a more macro view of your application, layering in those runtime insights and enabling the creation of specific mitigating WAF rules to buy time for patching. Finally, Kodem provides function level reachability alongside SCA and SAST scanning, enabling false positive reduction.

BUYER'S GUIDE

Choosing a cloud security solution is a series of tradeoffs that needs to be carefully mapped to an organization's architecture. The first and most provocative question you can ask yourself is if you really need a cloud posture solution at all. If your infrastructure is entirely containerized and defined as code, I'd recommend instead purchasing a holistic application security platform alongside robust runtime protection in order to get more security value while paying less than you would for a CNAPP.

Because of how CNAPP platforms developed, many buyers have multiple overlapping tools deployed in order to fulfill specific needs such as file integrity monitoring or agentless vulnerability scanning which have become more commonplace.

A cloud security buying decision should first be guided by your architecture. Some solutions offer robust protections for specific architectures, while others have focused on building broad solutions for a wide range of infrastructures. While we won't be able to fully cover every infrastructure approach in our flow chart below, you can always schedule a Latio consulting session for your infrastructure.

Decision Flow Chart

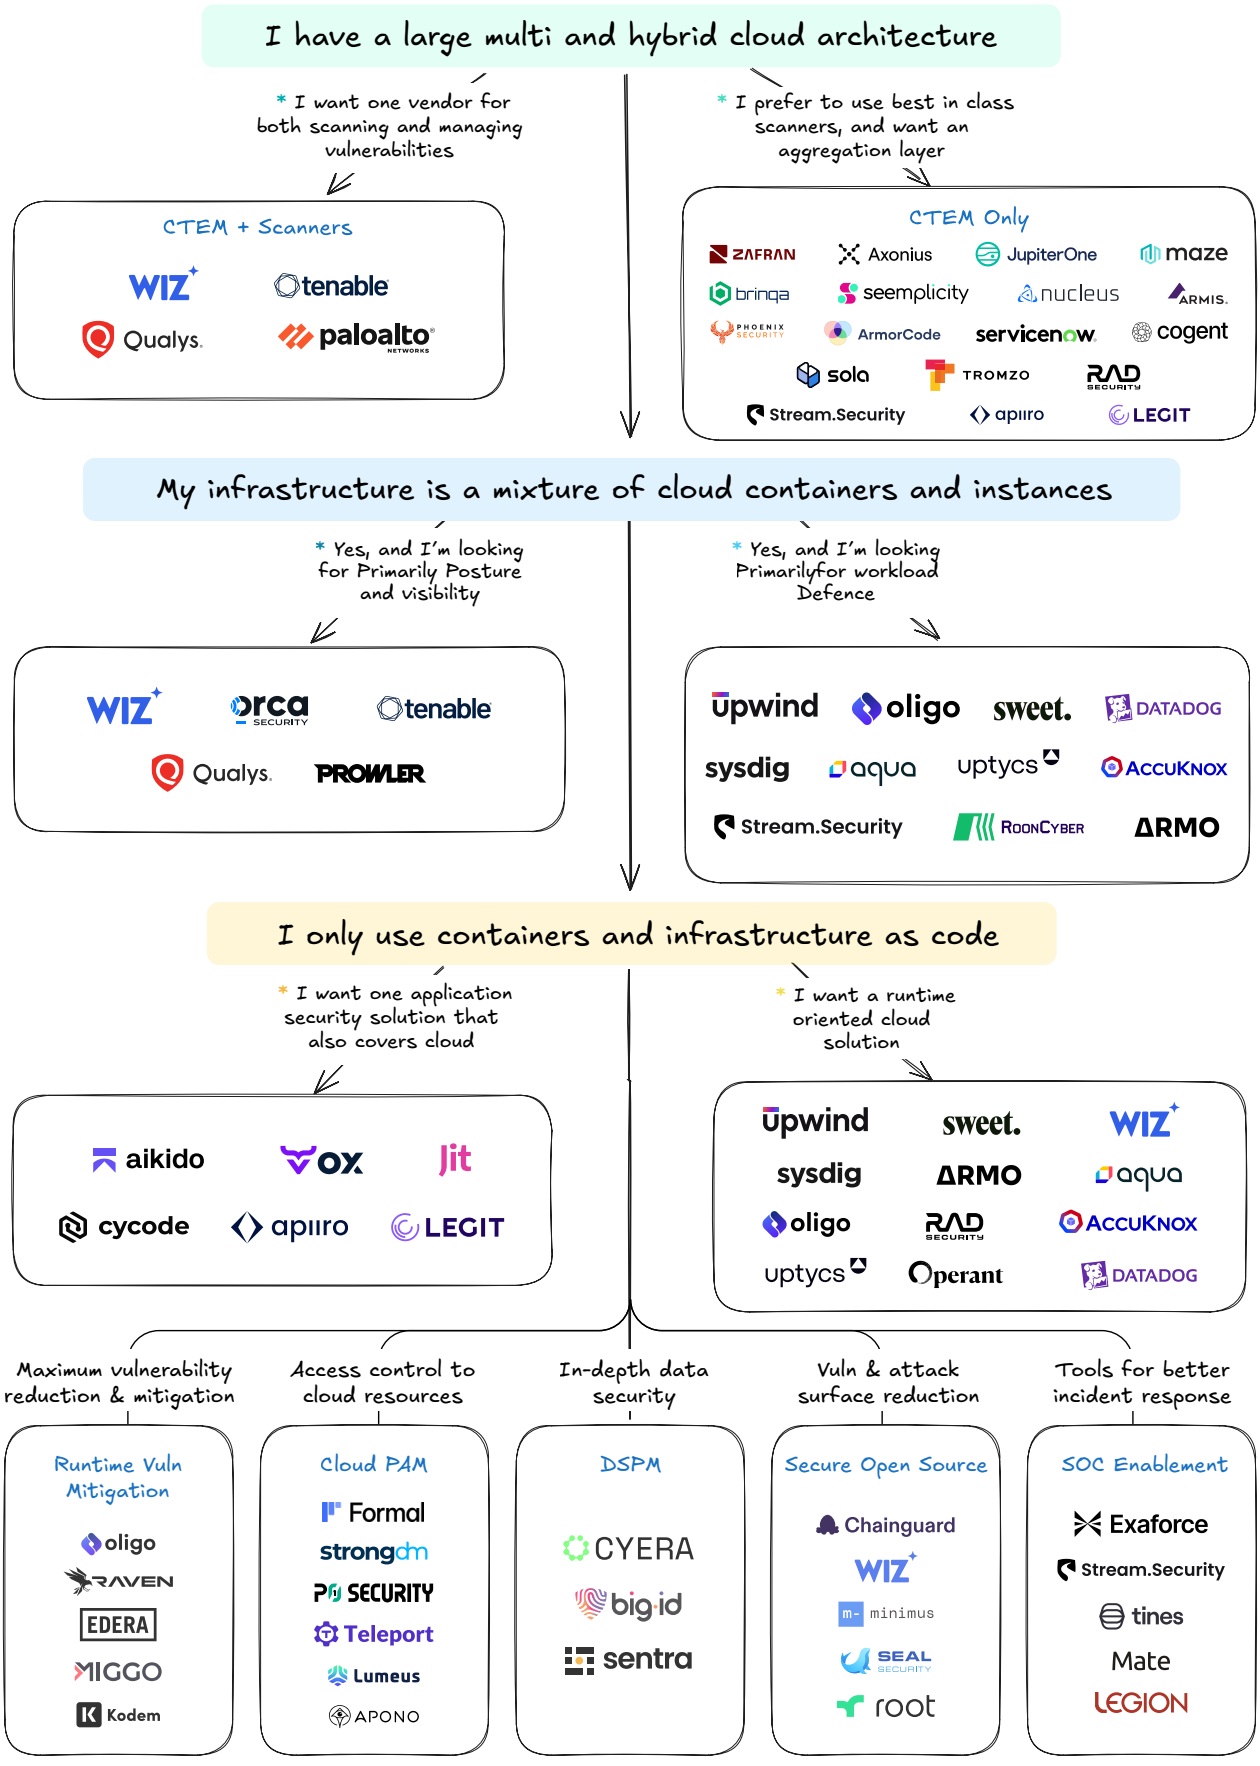

Walking through the flow chart, the first question is if you have a large hybrid cloud footprint. If you do, the key decision is which platform you want to use as your asset and vulnerability source of truth. If you're looking for a single solution to scan and manage all of the vulnerabilities in your hybrid cloud infrastructure, Tenable, Qualys, and Wiz (having a newly released on-premise scanner) remain the most viable options.

Most organizations for the time being however will have multiple scanners running across their infrastructure. In these cases, a dedicated CTEM offering may make the most sense. When selecting a CTEM solution, while most vendors are building towards the same ultimate product, it's important to select based on their point of emphasis. Each of these have various advantages and disadvantages.

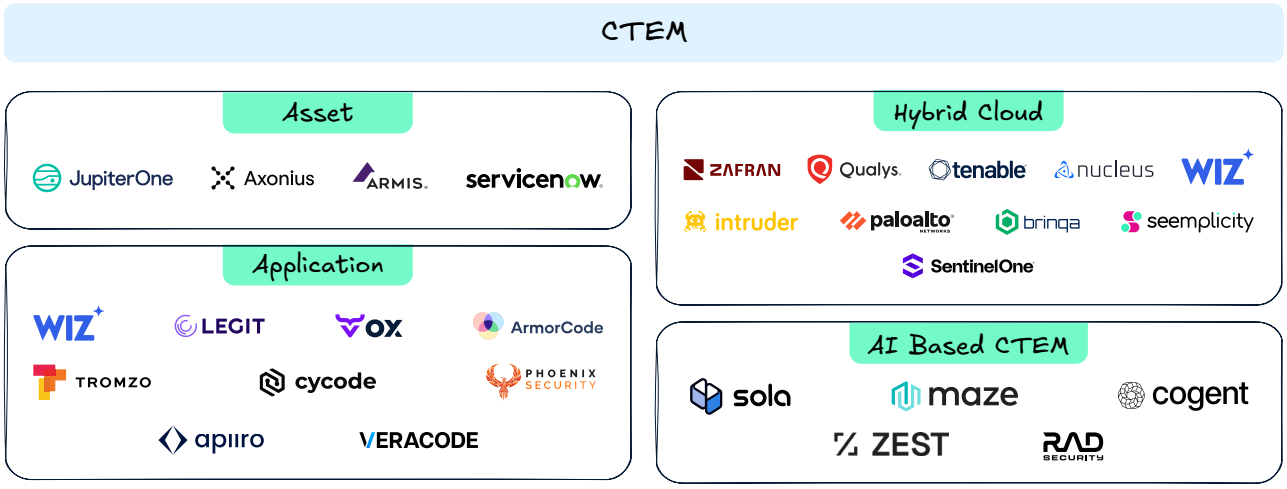

- **Asset** - for teams with large distributed architectures wanting a single source of truth for asset based systems with flexible querying capabilities across endpoints and servers.
- **Application** - for teams looking to consolidate their cloud and code vulnerabilities into a single place to drive effective remediation efforts
- **Hybrid cloud** - for teams looking to standardize their infrastructure patch management across on-prem and cloud systems
- **AI** - for teams looking to use the latest AI innovations to drive remediation, dashboarding, or investigative efforts

Moving then to cloud focused solutions, organizations have to make a prioritization decision in terms of if they want to prioritize a posture or runtime oriented approach. On paper, almost every CNAPP has the same features - agent and agentless scanning, runtime defense, vulnerability management, and some concept of toxic combinations. This enables them to cover a wide range of options, from static instances to containers.

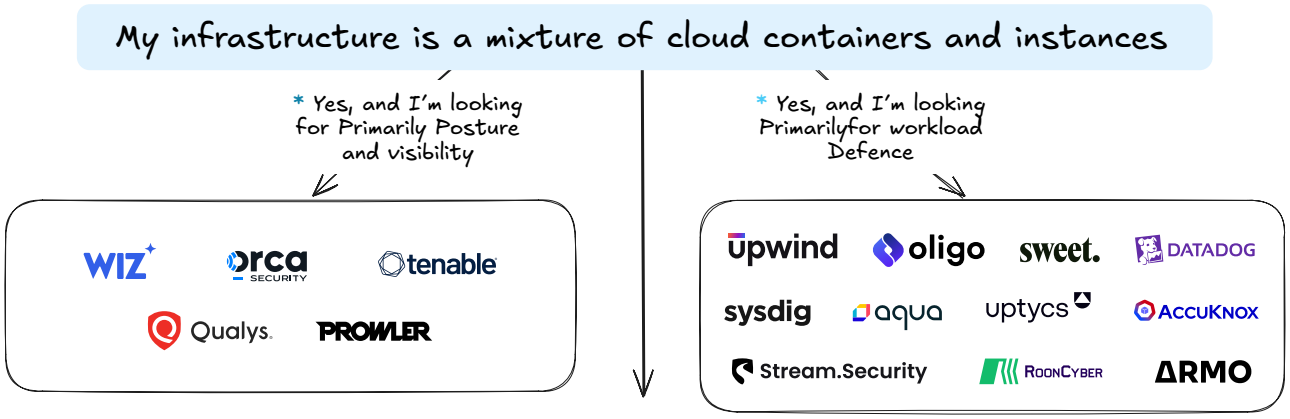

Of the solutions we highlighted in each box, posture and runtime oriented, various pros and cons exist between the competitive options regarding pricing, supported architectures, and general platform approaches.

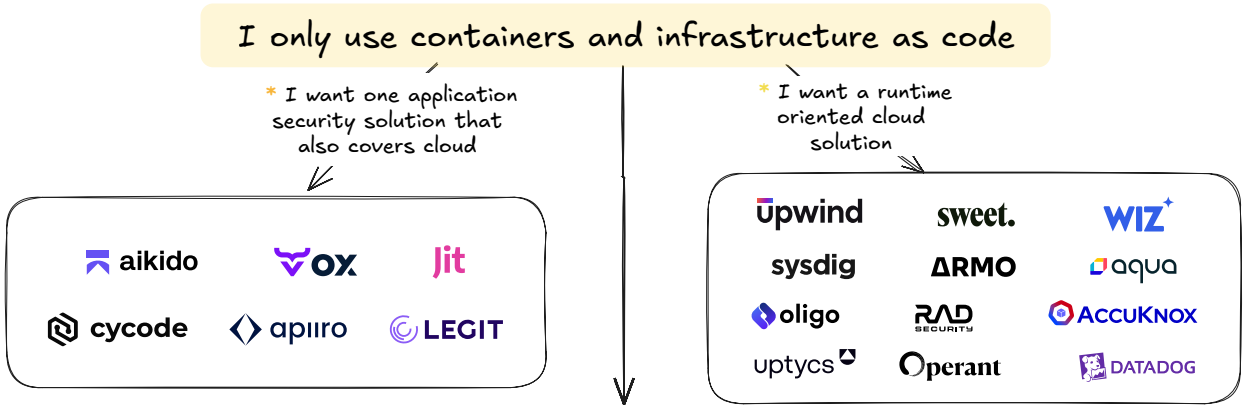

If your cloud infrastructure is highly focused on containers and infrastructure as code, you have the freedom to skip CNAPP altogether and instead choose a solution emphasizing code scanning, and a solution emphasizing runtime security.

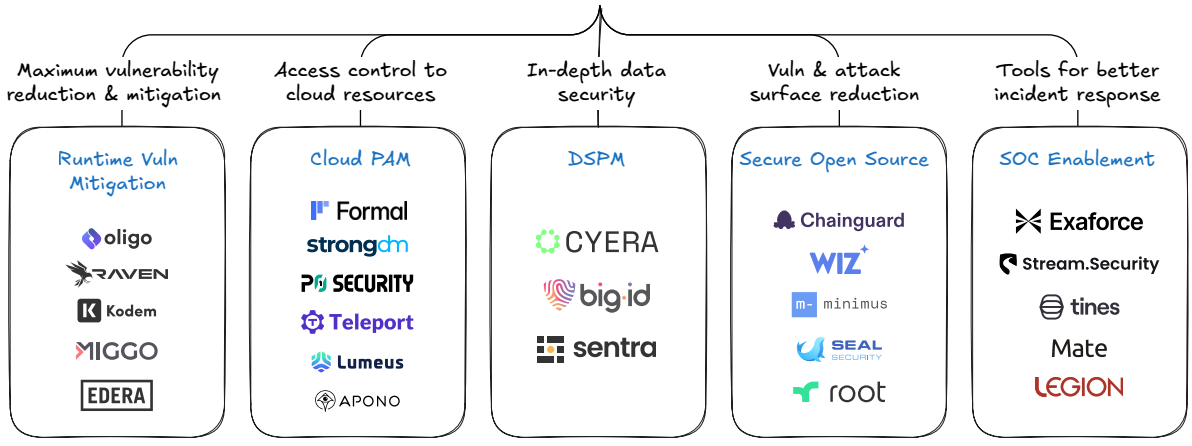

Ultimately, the larger the company, the more complex this security stack may get, increasing the importance of a CTEM solution as the glue that holds the systems together. At the bottom of the flow chart, we highlight specific developments that are highly requested features of CNAPP platforms alongside representative vendors in each category.

While we address each of these sections in more detail earlier in the report, when should you actually look to buy one of these solutions? Here's a brief guide for each:

Runtime Vulnerability Mitigation

These solutions excel in large regulated environments where there are a massive number of vulnerabilities to be managed. By using advanced function level runtime technologies to mitigate entire categories of vulnerabilities, each of these vendors offer ways to drastically reduce your attack surface while creating compliance ready attestation for why certain categories of vulnerabilities can be ignored.

Cloud PAM

Getting developers access to cloud workloads continues to be really complicated. Many organizations use customized AWS CLI workflows alongside okta in order to access AWS roles, and then a combination of bastion hosts and/or VPNs to access workloads. The vendors we highlighted in this section all focus on providing best in class access to cloud services and workloads. This allows teams to create logical access flows for their environments beyond complicated SSH key rotation nightmares.

DSPM

At one point in time, I've been alongside those who have said DSPM is a feature of CNAPP. However, the data protection needs of some organizations goes beyond what CNAPPs offer. For these companies, with large data stores across cloud, SaaS, and on-premise systems, dedicated DSPM solutions can help discover, classify, and protect distributed data environments.

Secure Open Source

A variety of solutions exist to make the vulnerability management process easier for organizations by owning a part of the process for them. Chainguard began with minimal container images but has expanded into both backporting patches and mirroring open source registries. Minimus offers minimal container images alongside deployment workflows. Root and Seal have focused on patch backporting for open source.

All of these solutions offer teams ways of standardizing projects that DevOps or Platform teams normally maintain. Purchasing these solutions make sense for organizations that are looking to remove this extra work off their development teams for maintenance upstream.

SOC Enablement

These solutions focus on enabling security operations teams to address incidents faster, but each highlighted vendor represents a different approach. Exaforce is a full data platform built with AI throughout the product and can function as either better AI assistance for existing alerts, or a fully fledged SIEM augmentation or replacement. Tines represents most "AI-SOC" products which are doing SOAR operations with more LLMs on top. Mate and Legion are using AI browser plugin automations to automate repetitive SOC tasks. Stream provides cloud context for traditional EDR tools to provide teams better cloud contexts.

Mid-Market Considerations

For companies under 500 developers, getting a high ROI with a limited team is essential. It's important for these teams to have a single solution that covers as many bases as possible for the security program. Additionally, smaller companies tend to have more consolidated infrastructure, allowing them to purchase tooling that is specialized for their infrastructure.

Two companies worth highlighting for their midmarket focus, and how they exemplify the decision making process based on architecture are Aikido and Intruder. Aikido provides all-in-one application and cloud security scanning, while Intruder provides all-in-one network based vulnerability scanning. For companies with a focus on IaC and cloud infrastructure, vendors like Aikido make a lot of sense, compared to companies with a hybrid Microsoft footprint, who would get more mileage out of Intruder.

Another path for mid-market companies is relying on open source options where they're the best fit. In application and cloud security, some great open source options exist for building your own security program such as Prowler, Falco, Opengrep, Trivy, and Kubescape. These

tools are also often supported as data ingestion points from other vendors, allowing you to operationalize them alongside other tools.

When purchasing a tool in the mid-market, the key factor to consider is how difficult it will be to operationalize. There are a million security tools out there able to generate thousands of daily alerts for teams to spin their wheels on. For smaller companies, this can be absolutely crippling to developer and engineering progress. Instead, teams should focus on solutions that consolidate and reduce alert overload through toxic combinations, and automated workflows.

Enterprise Considerations

For the largest companies, security tooling can quickly become a bloated mess to navigate, as different teams purchase different solutions to meet their own goals. Combined with complex regulatory requirements, there's often a vast amount of tool overlap. For these companies, having a mature CTEM strategy becomes essential, as the single data layer that unifies their vulnerability programs.

These larger companies can also afford to invest in best in class tooling such as choosing a CADR that emphasizes runtime application insights and protection, or advanced reachability features. CTEM then becomes the unifying tool that ingests relevant code and runtime contexts in order to determine the exploitability and priority of remediating a specific risk.

Another primary concern for enterprises should be support for standard data models, whether via API or OCSF standards. Enterprise tooling needs to be flexible enough to support dynamic use cases, tools, and data transformations that more standardized companies don't need to worry as much about.

For organizations with high regulatory requirements, minimal base images, self-hosting, and reporting all become critical features. These help standardize and reduce their attack surface, while meeting complex timelines around vulnerability remediation and incident response.

Conclusion

This report has explored the evolution of CNAPP tooling, the past, present, and future, and highlights a key insight: the future of cloud security will not be defined by a single platform offering a set of capabilities. Instead, organizations will move towards more consolidated cloud vulnerability management programs which sit alongside their security operations programs. These two programs will drive results focused respectively on proactive risk mitigation, and fast reactions to ongoing security incidents. The future of CNAPP lies in the broader operationalization of vulnerability management through Continuous Threat Exposure Management (CTEM) and incident response through Cloud Application Detection and Response (CADR).

VENDOR SPOTLIGHTS

All badge winners were given the opportunity to spotlight their product by having the Latio team author a dedicated page explaining why they were awarded in this report.

WIZ

Latio CLOUD SECURITY LEADER 2025

Latio CTEM LEADER 2025

Latio CODE TO CLOUD LEADER 2025

Wiz has defined the cloud security market since their founding by meeting their customers every step of the way. As teams have moved from visibility to runtime protection and code security, Wiz has launched robust product offerings to compete with the best of point solutions in these areas, offering formidable competition across the board.

When it comes to creating actionable cloud security alerts, Wiz continues to be a leader by extending toxic combinations from code, to cloud, and now even on-premise. Wiz's launch of an on-premise sensor, and the ability to normalize third party data are easily the most underappreciated product launches of 2025, as it brings the power and elegance of using Wiz as a cloud security solution into diverse on-premise environments. These capabilities alongside their newly launched incident response capabilities offer cloud expertise to the most diverse of organization types.

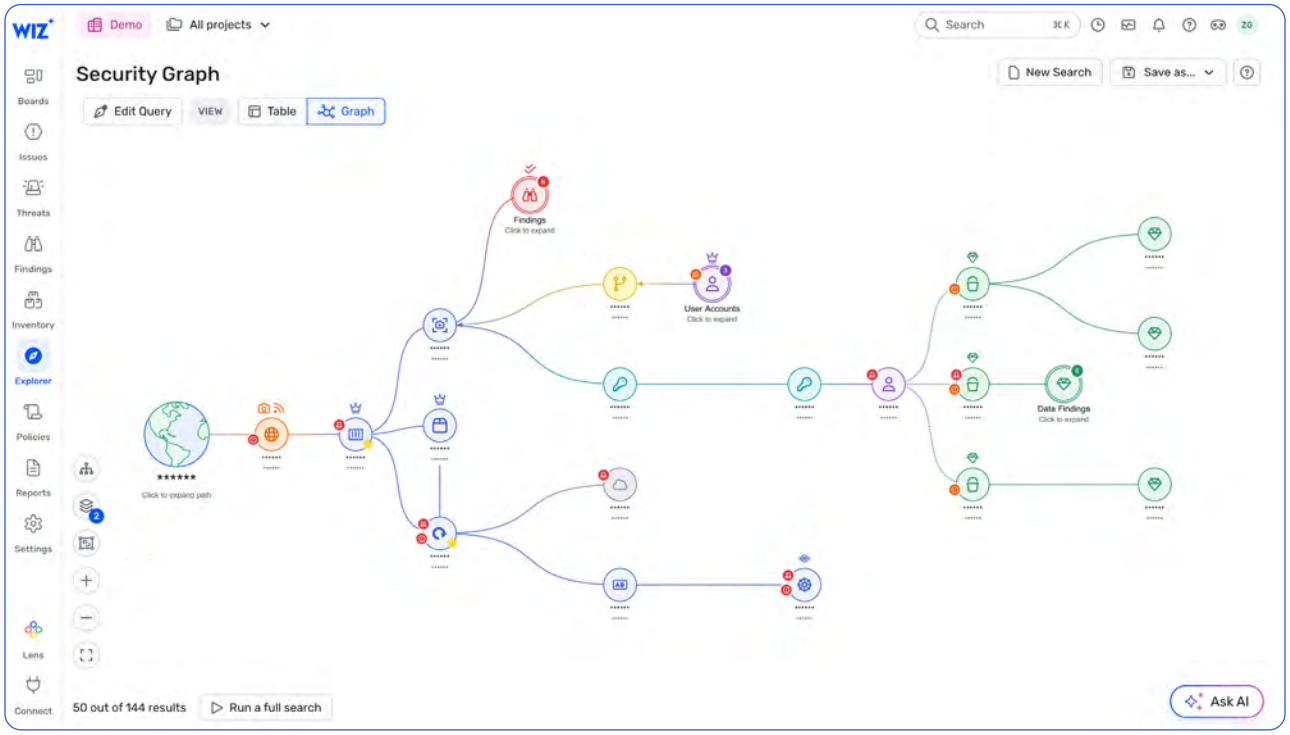

After their recent acquisition by Google, I was skeptical of CEO Assaf Rappaport's claim that Google would "enable us to execute and innovate even faster;" however, Wiz has once again proven the outlier by showing me at least 10 different major functionalities in preview at a time. The team is continuing to innovate at an alarming rate which has enabled them to be a major competitor in code, runtime, cloud, and on-premise environments all at the same time. Wiz not only defined the first ten years of cloud security, but they're setting up to be a major player for the next ten years as well.

The Benefits of Wiz

Best in Class Vulnerability and Posture Management

Wiz defined modern cloud security vulnerability and posture management by combining agentless vulnerability scanning, their graph search, and toxic combinations.

Unified Platform

Wiz offers competitive standalone solutions for application and runtime security, able to stand alone or as part of their larger CNAPP offering.

Rapid Innovation

Wiz continues to rapidly launch the latest in competitive features, from API discovery to AI investigations and MCP servers.

DATADOG

As a long time Datadog fan, I've always wanted to use their security solutions. Early on, I felt that other tools were more mature, but in the last year I've seen improvements that clearly came from their team acting on the feedback and needs from teams. I have no problem declaring Datadog a standalone leader in cloud security offerings. Instead of being only a developer tool with a security offering, it's a security offering that developers actually want to use.

Datadog offers security insights wherever your developers use them for performance ones - whether running SCA scanning, SAST, IaC, or the acronym soup of application security scanners as part of their pipeline monitoring, or catching application layer attacks against your runtime environment. When it comes to the security features teams need from a single platform, Datadog is one of the most comprehensive, offering everything from agentless vulnerability scanning through to SIEM ingestion and tuning.

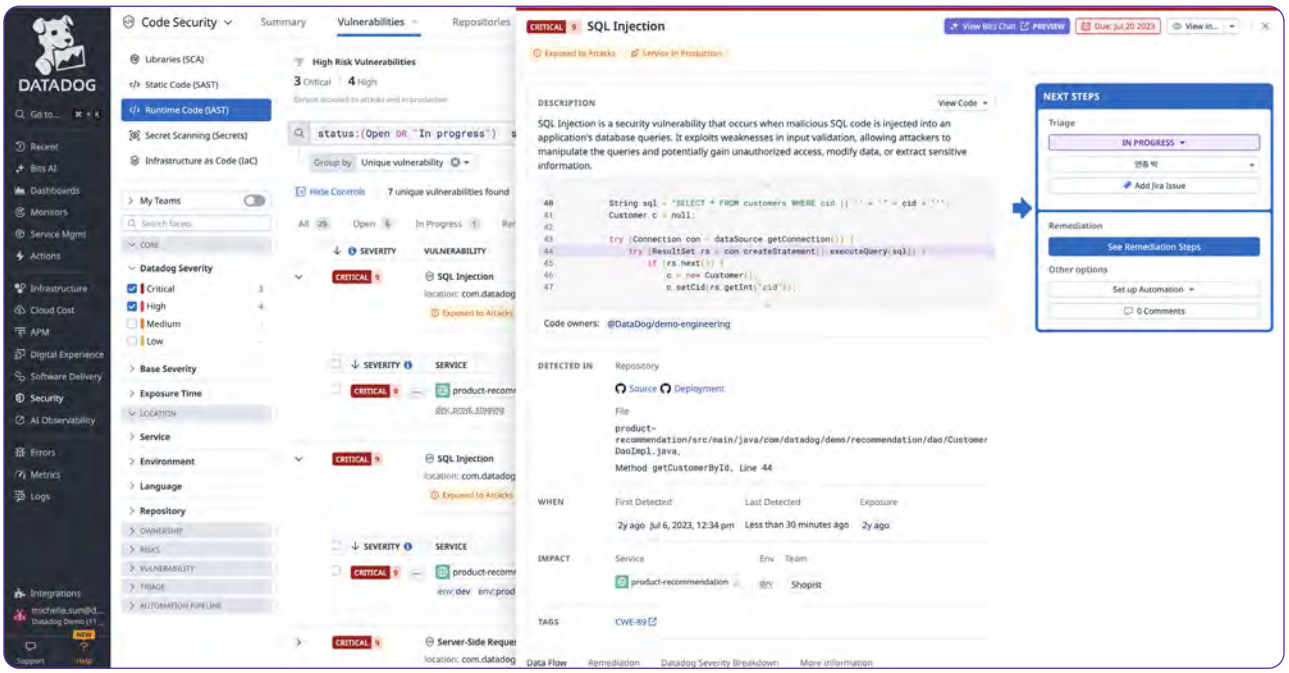

Where Datadog excels is their ability to contextualize cloud alerts to your application services, rather than merely the infrastructure. Whether it's giving web application detections the full traces of the attack, or prioritizing vulnerabilities based on function executions, Datadog makes it easy to get findings into the relevant developer workflows. Most cloud security tools don't have anything close to the level of application visibility Datadog is able to provide due to their roots in observability. This application layer visibility at runtime gives security teams the essential tools the next 5 years of cloud security demands.

The Benefits of Datadog

Meet Developers in the Tool They Use

Rather than cause more friction with development teams, leverage capabilities they also want to create unified outcomes.

Deep Application Insights

Gain meaningful insights into how your application normally functions, and get early alerts when an attack is underway.

Holistic Platform

Manage your entire infrastructure's health and security from a single source of truth, enabling joint investigations and reducing operational costs.

upwind

By focusing on a runtime approach to cloud security, Upwind offers deep visibility, scanning, detection, and prioritization capabilities across diverse cloud environments. Upwind's product is profoundly enjoyable to use, with a fast and simple onboarding process that I was able to get up and running in my environment in under 5 minutes.

Upwind has been on the frontline of runtime innovation, being the first CNAPP to include API security capabilities by inspecting layer 7 network traffic, launching ADR capabilities through their acquisition of Nyx, and using differential scanning in pipelines to determine if a vulnerability executes at runtime.

Upwind extends their runtime visibility to deploy leading posture capabilities as well, by using runtime-level intelligence to maximize the value of cloud inventory and configuration scanning. With its native agentless scanning module, customers gain unified visibility across inventory, configurations, vulnerabilities and exploit-paths, bridging runtime context and preventive posture into one platform.

Upwind's commitment to runtime has also rapidly expanded into both AI and API security testing more broadly, alongside robust runtime detection and response capabilities that have caught most attacks I've thrown at it. Their approach to network exposure is especially robust, going far beyond some other solutions by clearly indicating publicly exposed ports across services. Combined with their managed incident response offering, Upwind is a great choice for enterprises looking to secure their cloud environments at scale with precision.

The Benefits of Upwind

Easy Onboarding

Onboard your entire cloud infrastructure in minutes, using terraform or cloudformation.

Deep Runtime Detection

Don't miss a single process or packet that runs through your environment with end to end visibility into what your cloud workloads are up to.

Runtime Prioritization

Use advanced prioritization capabilities, from network to function level reachability to focus on securing what matters most.

sweet.

Sweet is all about discovering actionable findings in your cloud environment. Built for cloud and security operations teams, Sweet provides robust workload and application protection alongside general cloud posture and vulnerability management capabilities. Where Sweet really excels is in their AI capabilities; providing their own prioritization, 'Sweet Score™', and leveraging a unique detection approach that looks for suspicious application behavior in your environment. This allows them to turn early warning signals into full attack stories, tracing the activity of an attacker throughout your environment. By combining these signals into a single attack story across different layers of your cloud environment, they give operations teams actionable insight into protecting their entire cloud.

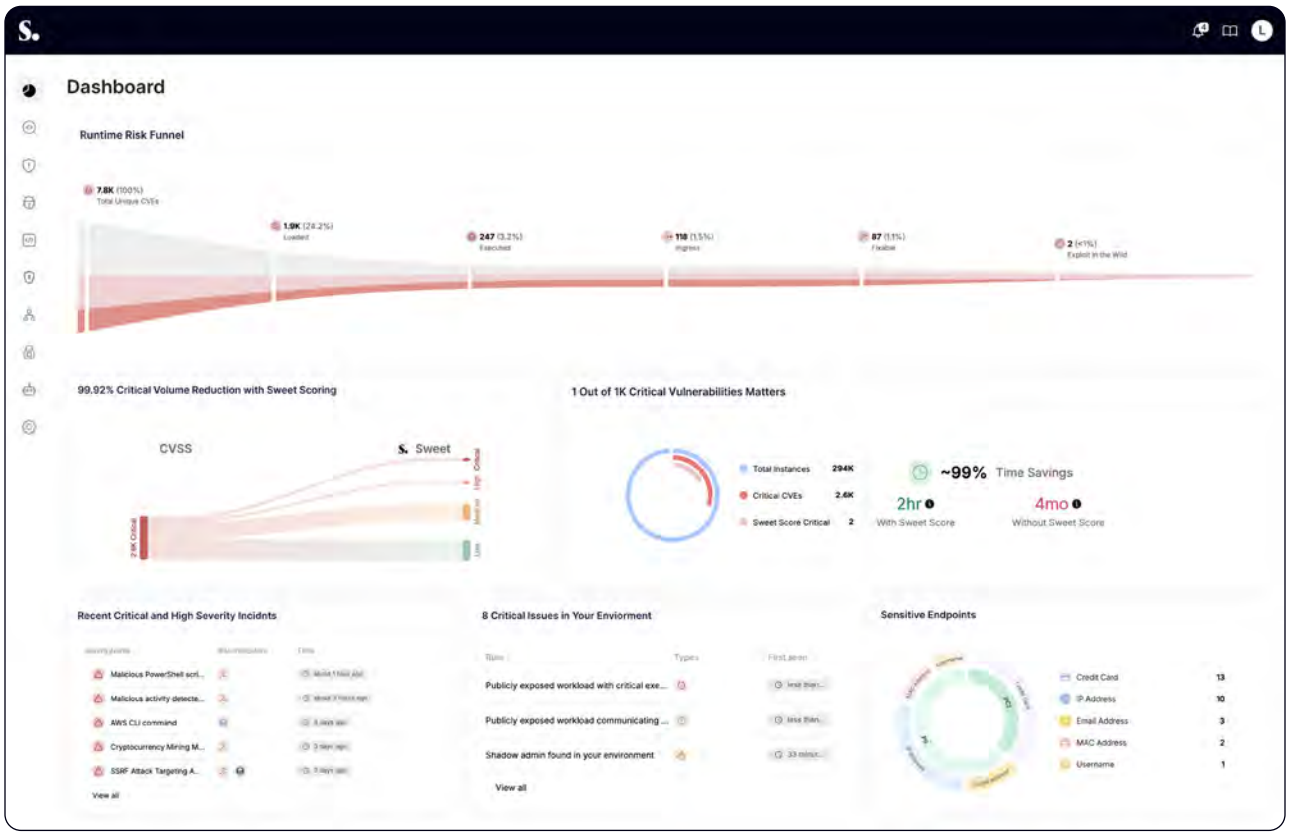

Sweet has also led expansions into AI SOC workflows and API layer threat detection, building impressive attack stories across all layers of a cloud attack. By extending to these layers, teams can see their real time attack surface across their cloud, getting a sense of all the data entering and leaving their environment. Sweet's runtime focus continues into other CNAPP features like Cloud Misconfiguration, Human and Non-Human Identity management and Secret Management, looking at what's actually running in your environment in order to determine risk and priority. Sweet also leverages runtime signals to help stop vulnerabilities before they reach production.

The Benefits of Sweet

Agentic Investigator - 'SweetX™'

Ingests data from thousands of organizational tools and 3rd party data sources to infer the root cause of an incident, enabling rapid response.

Robust Detections

Detect attacks as soon as they happen in your cloud environments, no matter where the attack starts.

Consolidate Threats

Combine high severity posture alerts with realtime data to gain a consolidated view of your posture and what to fix.

sysdig

Latio CLOUD SECURITY LEADER 2025

Sysdig helped define the modern era of runtime cloud security with Falco, their open source runtime detection engine. Today, Falco remains the most battle-tested runtime security solution, strengthened by adoption across the open source community. This makes Sysdig a great recommendation for teams that need real-time detection, customizability, and reliable protection in diverse and highly regulated environments.`

From a CNAPP perspective, Sysdig builds on that runtime security foundation and has all of the posture and vulnerability management features expected from a leader in the CNAPP category, with an emphasis on making runtime alerts actionable for teams.

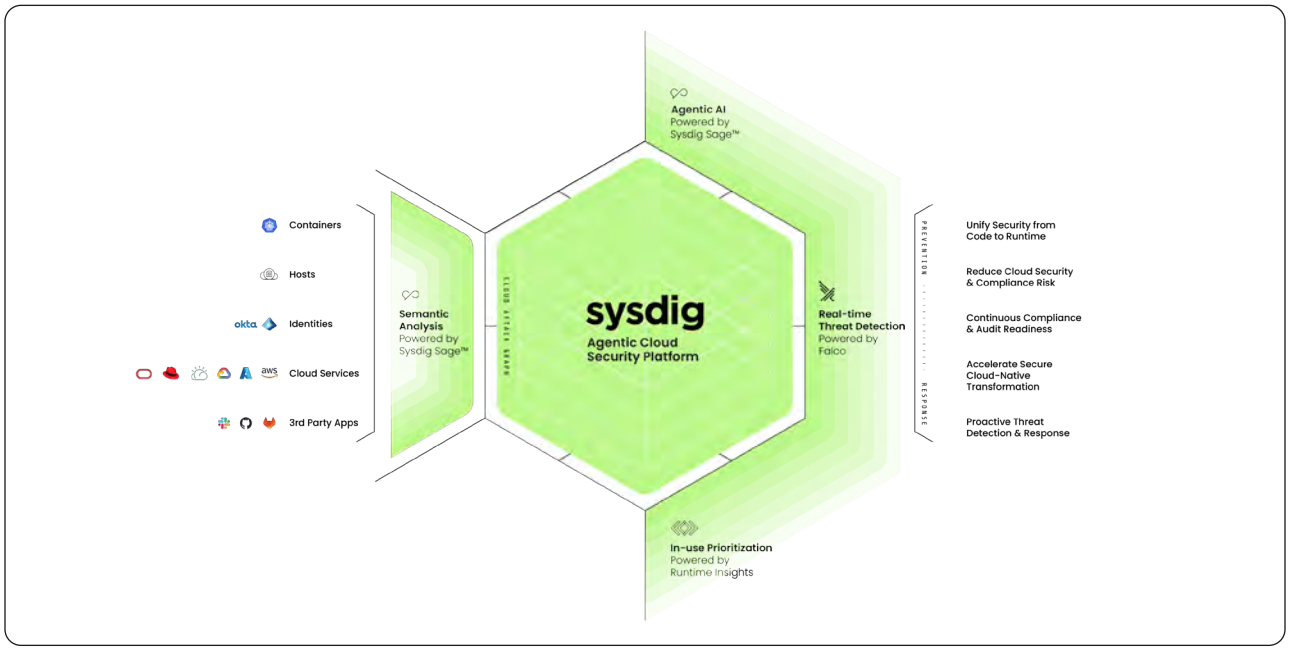

When it comes to runtime security, Sysdig continues to be a leader by delivering a depth of telemetry, customization, and AI-driven correlation that teams need to operationalize their cloud incident response programs. The granular detection & querying capabilities Sysdig provides are necessary to respond to real incidents.

Sysdig's pillars are open source, AI, and runtime security, and the product exemplifies leadership across these various points. Their Sydig Sage™ AI analyst enables accessing the data from the platform to drive meaningful insights via prompts. By empowering AI tooling with rich runtime telemetry, Sysdig is able to drive meaningful insights, from detections to reporting to remediation.

The Benefits of Sysdig

Open Innovation

With roots in open source, Sysdig's tools are always engineering first, allowing teams to adapt to the specifics of their environment.

Agentic AI

Tightly integrated with Sysdig's runtime attack graph, Sysdig Sage™ accelerates security workflows. Sage can be used to triage findings and guide decisions in real-time.

Runtime Insights

Sysdig's ability to gather deep telemetry allows them to understand business critical context which can provide deeper insights and more actionable decisions.

maze

Latio
CLOUD SECURITY SEGMENT LEADER
2025

Maze has built an agentic approach to vulnerability management by giving AI agents resources to investigate CVEs and your environmental context to determine the true priority of a vulnerability. A struggle in vulnerability management is the unique nature of every CVE - no single method of reachability can ultimately account for every edge case in the way AI can.

Maze adapts to the specifics of your environment, investigating CVEs on a case by case basis, the way a highly skilled product security engineer or developer would. Whether investigating containers or other cloud workloads, Maze dives into your environment in a way that traditional tooling categorically cannot, capturing the nuance and context needed to understand real exploitability and severity.

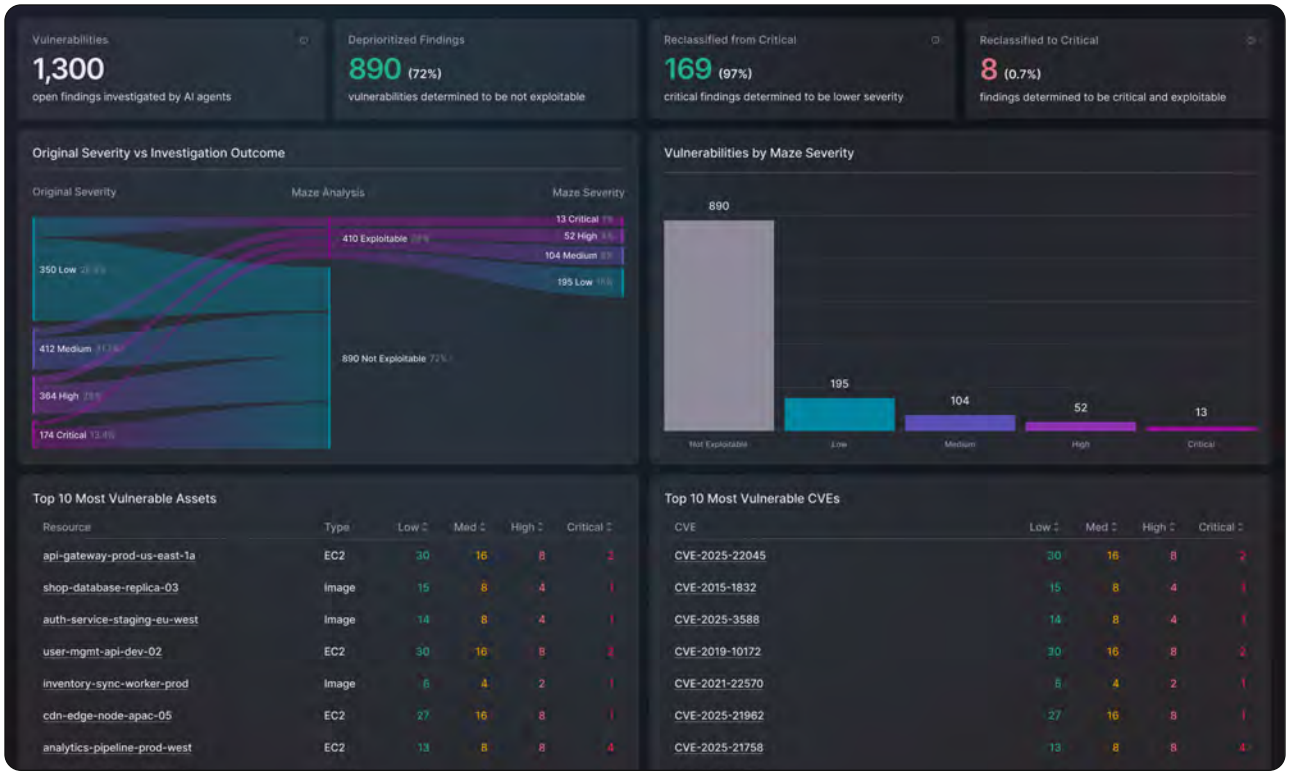

By giving AI the freedom to investigate every vulnerability, teams are unconstrained by the limits of traditional approaches to vulnerability research. Each CVE receives a thorough analysis of how it would be exploited, alongside evidence to attest if the vulnerability is a true or false positive. Maze is focused on the main pain points vulnerability management teams face every day: understanding what findings mean, separating the true positives from all the noise, analyzing the remaining exploitable issues to prioritize what matters, and attesting to the false positives for auditors.

The Benefits of Maze

Save Developers & Security Time

by enabling them to only focus on issues that matter, with steps to reproduce the attack, full vulnerability analysis, and prioritization based on risk to your environment.

Reduce Vulnerability Workloads

with best-in-class false positive detection, rule out massive numbers of false positives.

Understand Every Weakness

with thorough exploitability analysis and control assessments, you can accurately prioritize every potential threat vector.

cogent

Lalio
CLOUD SECURITY SEGMENT LEADER
2025

Cogent is an AI-native vulnerability management platform that helps security practitioners resolve risk more efficiently across the entire lifecycle: code to cloud, asset mapping, CVE prioritization, and remediation.

Cogent's multi-agent system ingests security data and layers in business context like ownership, criticality, and exploitability to prioritize and group vulnerabilities for remediation. It delivers clear, business-aware, remediation guidance and executes workflows end to end through ticket creation, assignment, and resolution, either autonomously or with a human in the loop.

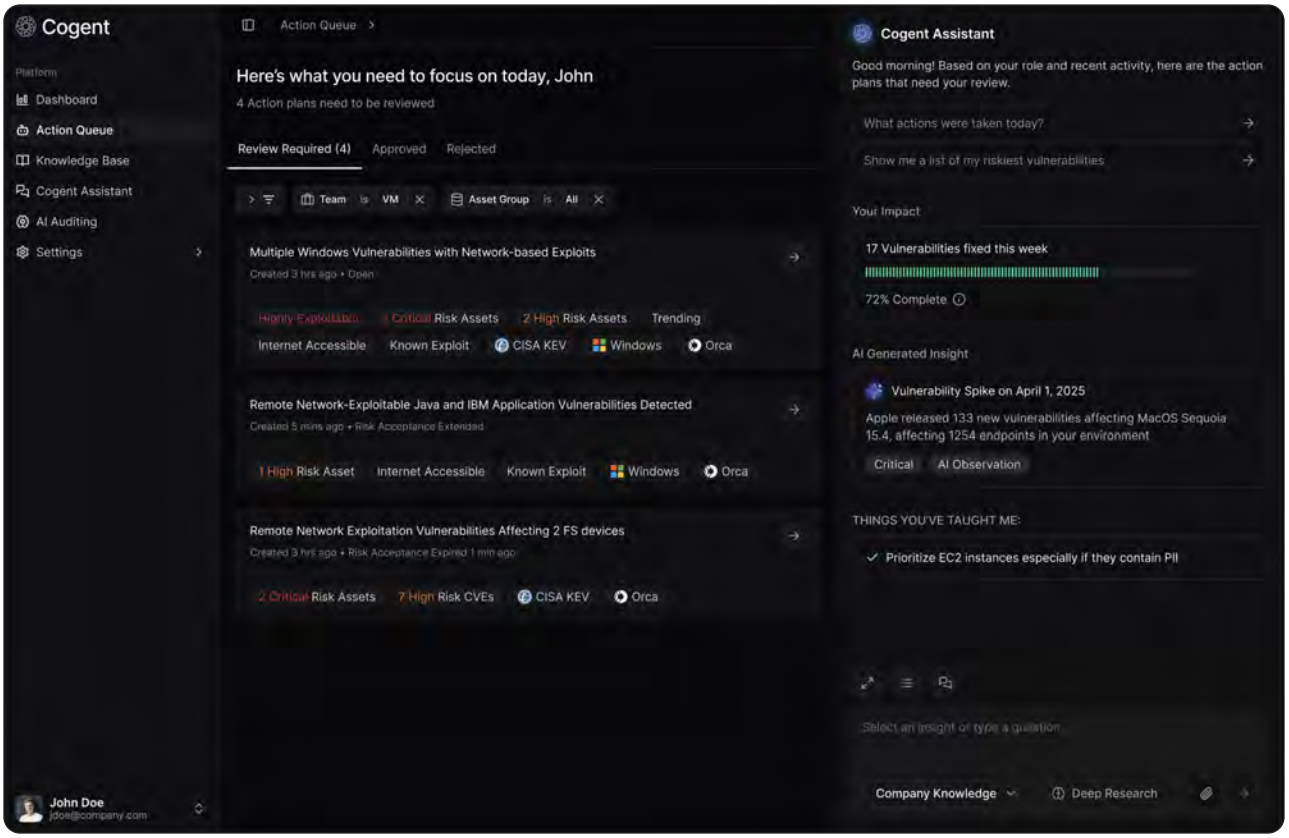

AI-native approaches to vulnerability management are still novel. Cogent has built a holistic approach with an AI knowledge platform that serves as the foundation for agents to guide remediation at scale. This enables teams to respond faster and get accurate updates on risk resolution.

The Benefits of Cogent

Integrate with Your Tools

Works with scanners, cloud, identity, EDR, CMDB, and ITSM to ingest findings and push tickets, updates, and verifications.

Prioritize with Accuracy

AI ranks CVEs using exploit intelligence and your environment, explains the why, and bundles fixes by asset and owner.

Remediate with Confidence

Step-by-step plans, automatic ownership assignment, SLA tracking, and closed-loop verification.

Operant

Latio CLOUD SECURITY SEGMENT LEADER 2025

Latio CLOUD SECURITY INNOVATOR 2025

Operant leverages deep insights across your runtime environment to create meaningful protections for applications - from AI to APIs. Operant gives unparalleled visibility into how traffic flows in your application, enabling detection capabilities across all kinds of applications, clouds, and visibility into sensitive data flows. Beyond visibility, Operant also takes action, aliasing sensitive data and stopping AI attacks in place.

Operant has expanded capabilities in securing AI applications, especially with complex MCP usage in production. Operant detects and prevents various runtime exploits of AI applications, aliases sensitive data, and open sourced a red-teaming suite for detecting data leakage to AI systems.

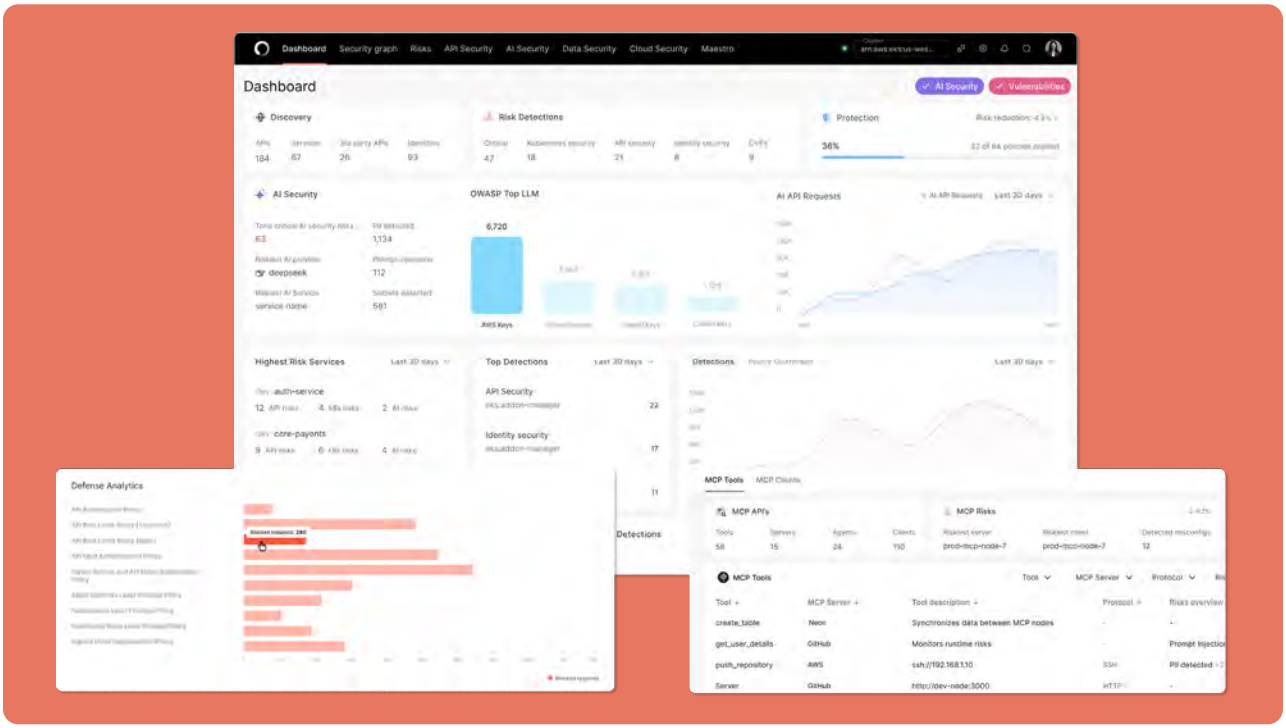

As software supply chains grow more inter-connected than ever, and applications chain MCP tool calls across various providers, deep runtime visibility into how your application is actually behaving is key to understanding and securing your applications.

Operant's approach to runtime AI application security is truly holistic by defending APIs, applications, agents, and containers all at the same time, offering one stop for meaningful workload protection.

The Benefits of Operant

Security at Runtime

View all of your data flows at runtime, not letting a single packet go overlooked.

Deploy AI with Confidence

Protect your runtime applications from leaking sensitive data, or getting manipulated by attackers.

Holistic Protection

Secure your containerized workloads in production with leading detection and response capabilities.

Raven has built a leading Application Detection Response (ADR) solution for companies to protect critical applications. Raven unlocks the missing piece of workload protection for software teams by using deep application inspection at runtime to tie functions that execute back to the original code that produced them. Developers get the insights they need to separate action from noise by contextualizing alerts to an application's environment, finally allowing security and developers to talk the same language.

Without ADR, doing vulnerability management and incident response for SaaS applications is full of challenges. Every time a CVE is investigated, nothing but wildly unfeasible exploit conditions arise, leading to developer complaints. Every incident leads to confusion as teams try to trace process trees back to the code that produces them. Raven solves these problems by giving actionable data to security and developer teams to only fix real problems.

Raven helps with incident response and vulnerability management by offering teams the critical application layer insights they've been missing. They do this with low overhead, offering out of the box performance dashboarding with common developer tools. Beyond insights alone, Raven even offers an elegant proactive protection solution by permissioning libraries into known categories of system calls, enabling true zero day protection. For enterprises looking to bring their applications into their security program, and mitigate advanced attacks, Raven is a great solution.

The Benefits of Raven

Eliminate Vulnerabilities

Respond only to known vulnerable function executions, prioritizing what matters.

Stop Attacks

As application layer attacks continue to increase, block the latest attacks while giving your team time to patch.

Simple Deployment

Deploy an extremely efficient sensor in minutes to immediately reduce your vulnerability counts by 99%.

Formal

Latio CLOUD SECURITY SEGMENT LEADER 2025

Formal is an example of the next generation of privileged access management solutions, giving security leaders complete control over sensitive systems and visibility into access patterns and data flows. Formal secures infrastructure (DBs, K8s, VMs, MCP, etc.) access across on-prem and cloud environments, and governs agentic access. Key capabilities include just-in-time data access, real-time data flow visualization, session logging, and on-the-fly sensitive data redaction - even for data utilized by AI.

Formal also extends detection through to runtime anomaly detection, and can alert and take action on insider threats. This allows security and engineering teams to move fast while maintaining a single point of enforcement for access policies, logging, and just-in-time production access. These capabilities extend seamlessly to AI agents, modern MCP environments, and traditional infrastructure.

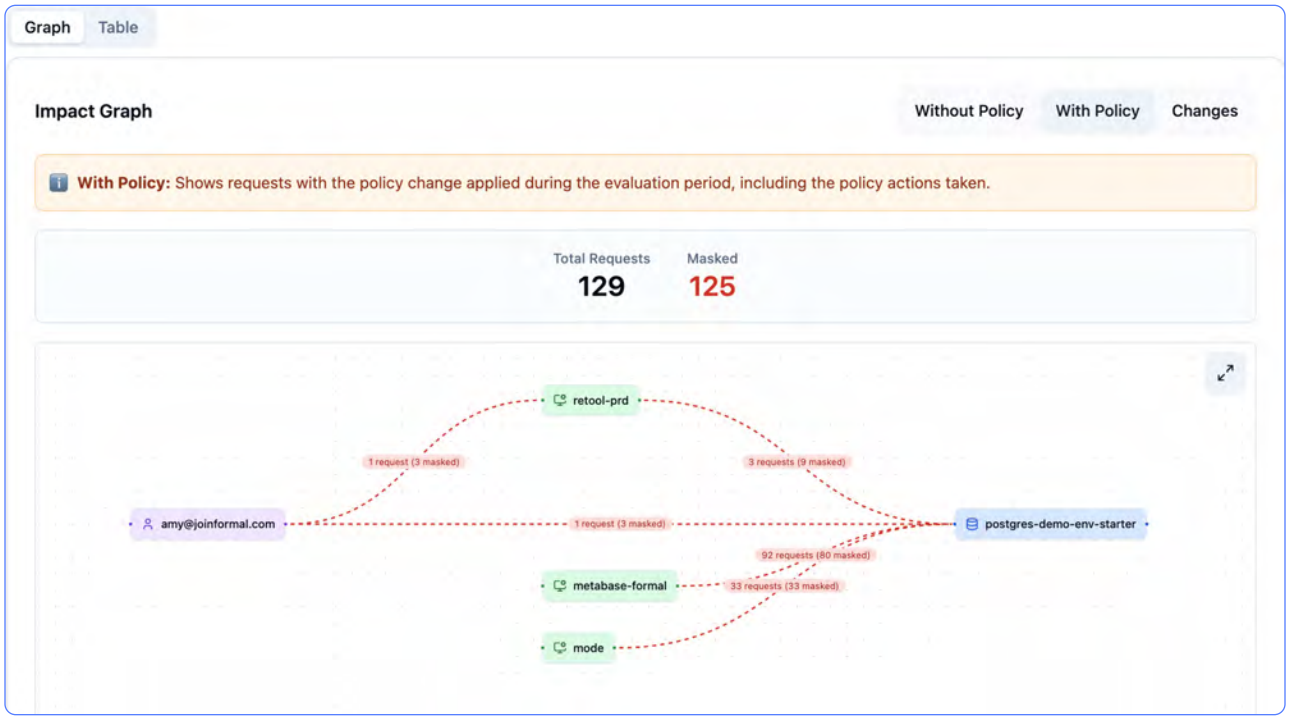

By combining all of these features into a single tool, teams have everything they need to build secure access patterns no matter where their data lives, and which humans or AI agents are interacting with it.

The Benefits of Formal

Visibility

Comprehensive, end-to-end visibility into sensitive data flows - illuminating how services interact and precisely who is accessing critical information.

Workflows

Implement streamlined, policy-driven workflows that allow users + services to securely request and receive access to sensitive resources with full auditability.

Enforce

Enforce real-time context-aware data masking and filtering to safeguard sensitive information and detect any unauthorized access or misuse.

CYERA

Latio
CLOUD SECURITY SEGMENT LEADER
2025

Cyera offers robust data security and protection capabilities for all types of data stores, whether cloud or on-prem, extending into DLP and real-time AI protection. Using Cyera starts with agentless integrations into your data environments to create robust visibility into various data types - from structured to unstructured data. While these data tagging systems can be imprecise in other tools, Cyera has created a robust classification system that goes beyond basic PII/PHI/PCI detection to include learned classification unique to your business.

Once the discovery is done, Cyera tracks access privileges and identity behavior, tying user accounts back to their root identity provider. This enforces granular tracking and authorization mechanisms giving teams real insights into how their data is used by humans, machines, and AI agents across the organization. With Cyera's approach to DLP, teams can automatically alert, redact, or remove access to sensitive data - from AI to databases - with AI analysis taking the first pass.

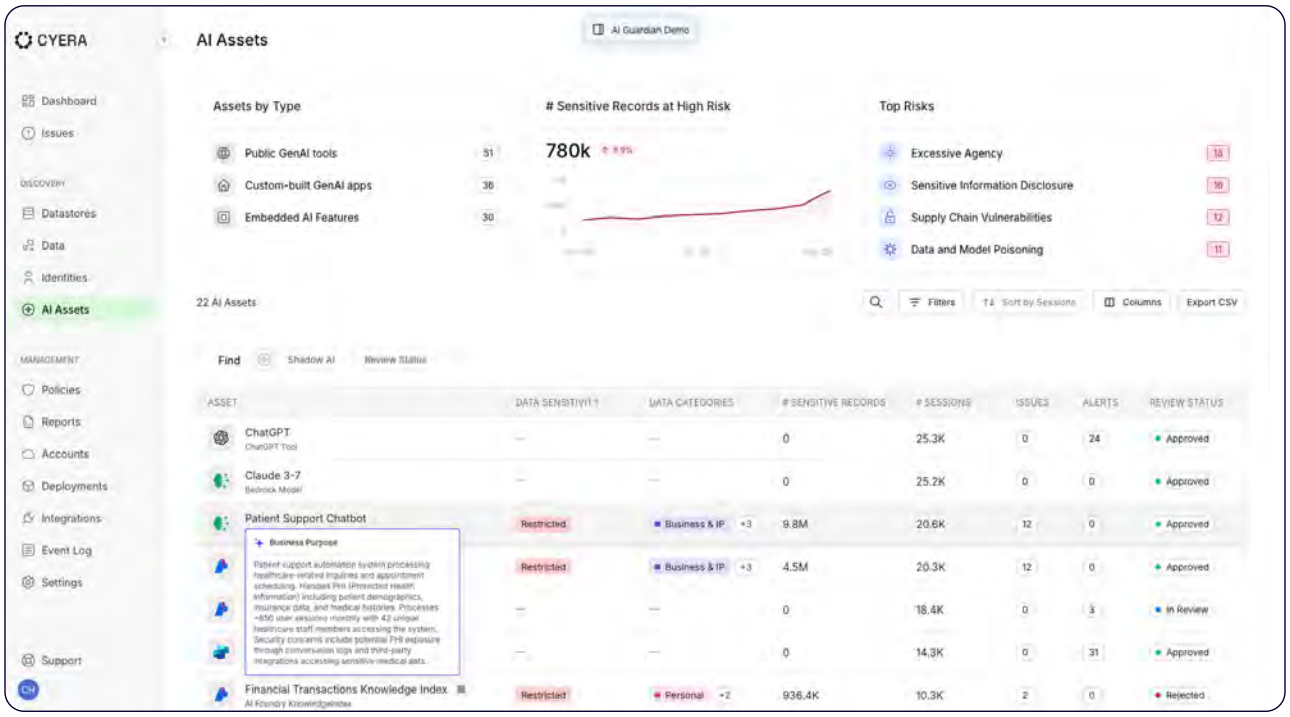

Cyera also utilizes their robust data discovery and classification capabilities to show where AI access exists across environments (whether public AI like Perplexity, embedded AI like Microsoft Copilot, or homegrown apps using AWS Bedrock), what types of data the agents have access to, and now even provide real-time protection capabilities. This covers everything from AI-SPM to runtime threats, making them a holistic AI security provider.

The Benefits of Cyera

Unified Data and AI Security

Cyera secures data at rest and in motion across IaaS, PaaS, SaaS, and on-premise, giving security teams visibility into where sensitive data resides, who can access it, and what actions humans or non-human identities take.

AI-Native, Business-Aware Classification

Cyera combines pattern-based and AI-driven methods to classify sensitive data, identifying business-specific data classes across structured, semi-structured, and unstructured assets.

Remediation and Action

Cyera extends security beyond the cloud, delivering integrated and native actions that remediate risks, prevent data loss, and safeguard information everywhere it resides.

Chainguard

At its core, Chainguard offloads necessary but time consuming vulnerability management toil from your platform, security, and devops teams. Many cloud security organizations spend years partnering with their platform teams attempting to build golden images that follow the latest compliance standards, and reduce their attack surface by eliminating unneeded packages. Chainguard Containers exist to offload this process, maintaining always up-to-date and minimal images, eliminating vast amounts of CVEs.

Chainguard's recent expansion beyond container images and into virtual machine images and code libraries is a huge step forward in security capabilities for organizations. By building common open source packages from source and simplifying the distribution, Chainguard takes on much of the attack surface of using language libraries. Standardizing on a single, trusted source for open source also makes complex mirroring setups used by advanced organizations more obtainable.

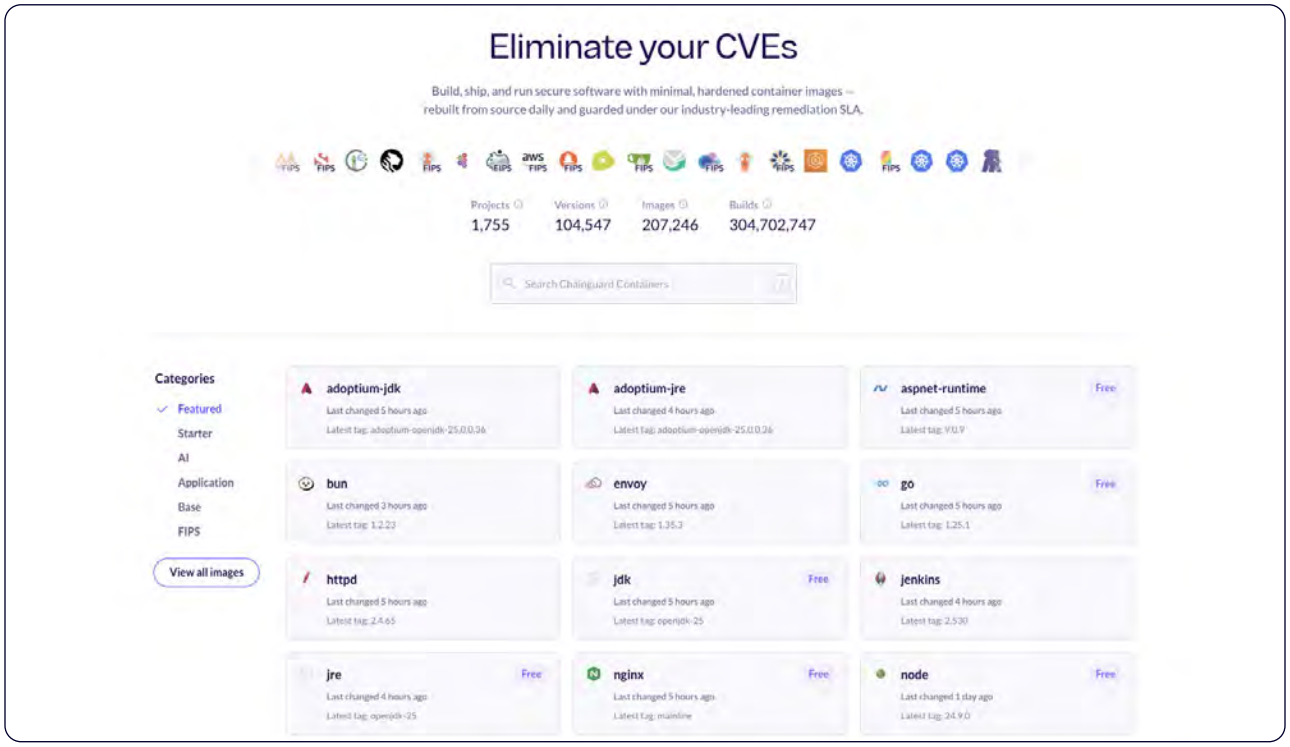

Finally, Chainguard's backporting of patches to older versions of language libraries like Python helps teams mitigate vulnerabilities without the major migration challenges that often occur with large changes. Altogether, Chainguard has expanded beyond minimal base images, and into providing a reduced attack surface for anyone using open source.

The Benefits of Chainguard

Minimal Images

Secure your infrastructure with always up to date minimized base images built from source to reduce your attack surface and remove vulnerabilities.

Secure Code Libraries

Reduce the threat of open source malware by adopting code libraries built from source in Chainguard's ecosystem with backported patches.

Reduce Vulnerabilities

Backported patches for language libraries and minimal, continuously updated images remove hundreds of vulnerability counts from production environments.

There's no platform that has more features for application security practitioners than OX. Their product provides every kind of scanner, asset mapping, and integration that a team could want. Most recently they also announced VibeSec, a collection of tools for bringing security policies and cloud context into AI coding tools. Whether you're using OX to ingest third party vulnerability findings to layer in code to cloud visibility, or as an all-in-one application security testing solution, OX has the use case covered.

OX offers deep integrations across your application security stack, monitoring code from developer IDE out to production. They offer every kind of application security scanning, mapping findings out to their deployed state, and prioritizing based on numerous details about the finding itself and the runtime environment.

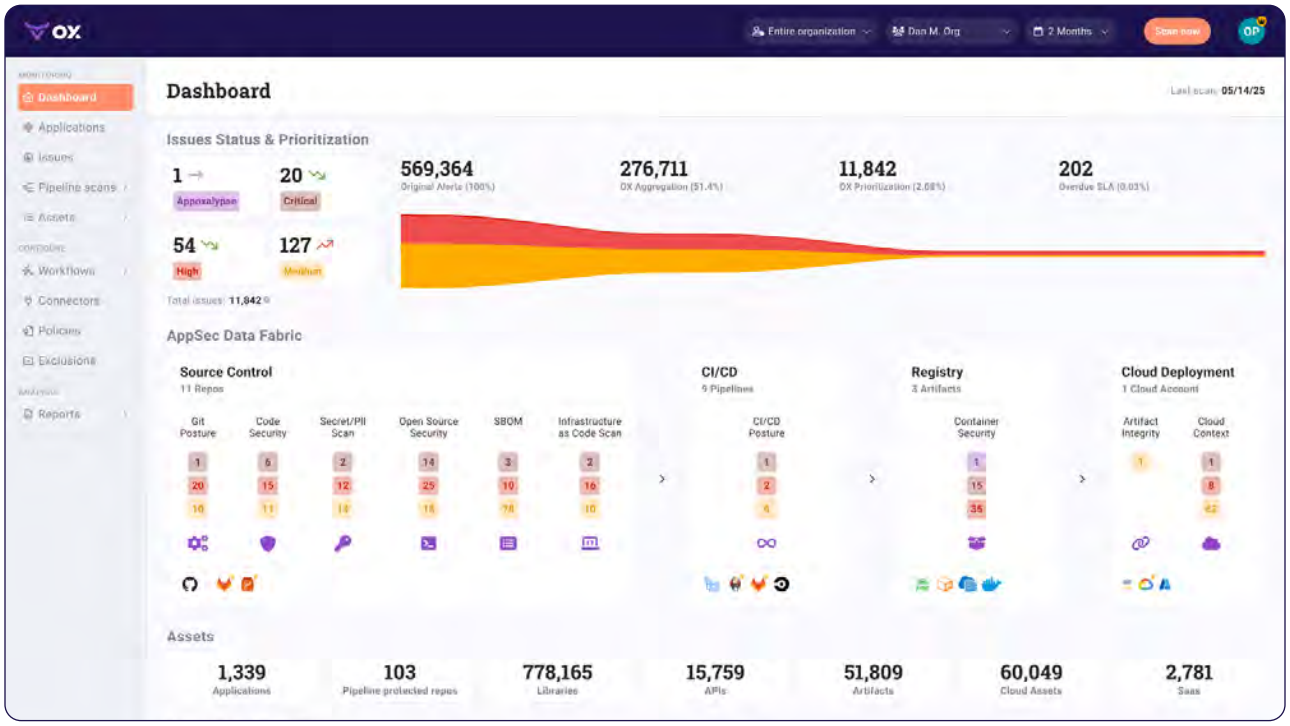

Prioritization is where OX's feature set really shines, using details such as what databases you're using alongside AI exploitation simulation in order to determine the priority of patching a finding. For cloud security teams, OX also provides agentless cloud asset scanning, vulnerability management, and CSPM scanning. OX extends additional runtime context through integrations with dedicated CADR providers, allowing even more robust prioritization insights.

The Benefits of OX

Vibe Code with Context

OX embeds security context - runtime, infrastructure, and policies - into AI coding environments, preventing vulnerabilities from being generated.

Secure Your Entire Stack

OX provides every scanner and integration an application security team needs to ensure the entire environment is covered.

Map Your Application to Deployment

OX maps pieces of code through to their containers and running APIs to enable accurate prioritization and remediation.

ARMO

ARMO has always been a leader in giving cloud security teams actionable visibility, findings, and detections in cloud native environments. ARMO's security solutions provide end to end insights for cloud systems - whether it's vulnerability reachability, posture misconfiguration, or detecting threats across your environment. ARMO's visibility into Kubernetes-based workloads and applications, highlighting risky workloads, and distinguishing theoretical from real risks based on runtime context.

ARMO tells the complete CADR detection story, discovering attackers as they traverse from your web application, to your workload, and to other cloud assets. By unifying these layers of detection, security operations teams are able to fully understand and respond to threats in your environment, while tuning out false positives that don't mirror real attacker behavior. This method of attack detection enables tools to trace most modern attacks which start at your application, and pivot deeper into your environment.

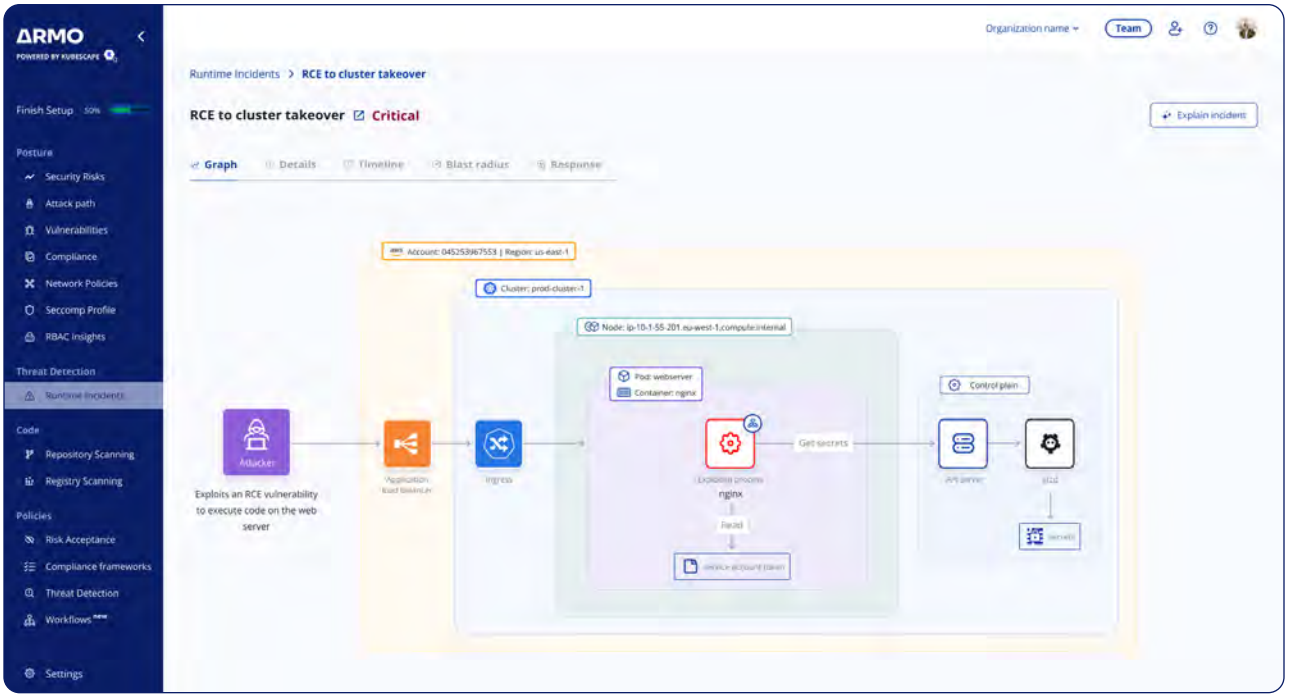

ARMO's actionable insights continue through to their approach to posture, as they provide robust insights into cluster RBAC permissions, network policy, and security policy generation based on how workloads function. By baselining workload behavior in your environment, ARMO is able to help deploy advanced security features that are typically extremely time consuming to implement. One of my favorite features of the platform is smart remediation, which lets you know if it's safe to change a workload to tighten the security without interfering with its normal behavior during production.

The Benefits of ARMO

Deep Application and Workload Visibility

ARMO builds full application profiles, covering APIs, linux capabilities, file access, networking, and syscalls to understand the full risk associated with each workload and application.

Beyond Just Detection

ARMO offers advanced protection and prevention capabilities based on runtime context, such as network policies and seccomp profiles.

Open Source Foundations

ARMO is powered by the open source Kubescape engine, a leading cloud-native open source security project with more than 11K starts, and used by 50K organizations.

Axonius

By building a robust discovery solution across devices, identities, assets, and software, Axonius positioned themselves well to evolve into a CTEM solution. Axonius creates a rich relational graph with asset insights across all of your devices, from user laptops to containers, enabling customizable risk scores and patching processes along the way.

The flexibility of Axonius' data platform ensures that whether you're assessing the impact of a vulnerability against your user's laptops or your cloud environment, you have the data you need to make an actionable decision.

Axonius is an especially powerful tool for teams with a large number of security solutions deployed across their endpoints, as having a single data lake to normalize and prioritize data becomes the key outcome teams need.

As various security vendors race to create the single vulnerability data lake to drive actionable remediations, Axonius is leading due to their robust mapping of vulnerabilities across various asset and identity types.

The Benefits of Axonius

Gain Complete Control of Your Environment

Aggregate, correlate, and normalize data from every tool your team is using across on-premise and cloud systems.

Orchestrate Actions

Automate where it makes sense, coordinate workflows when things require human intervention, or trigger tickets across systems.

Prioritize Fndings

Enrich asset data with business and security context to uncover exposures, compliance gaps, and infrastructure overspend.

oligo

Latio CADR LEADER 2025

Latio CLOUD SECURITY INNOVATOR 2025

Oligo provides robust runtime protection for cloud services, workloads and applications of all kinds - whether AI, cloud, web, on-premise, or third party. By baselining and monitoring all aspects of your environment, from code packages down to the function layer, Oligo is able to thoroughly secure workloads. Insight into what Oligo is capable of can be seen in the Application Attack Matrix, which introduces application layer attacks to the MITRE framework.

Oligo has been named a leader in CADR due to their ability to catch attackers no matter how complicated or obfuscated they are. First, Oligo is able to prioritize the vulnerabilities that matter most by looking for runtime behavior at the function level. Then, Oligo is able to make sure your applications and cloud workloads function only as intended, blocking attacks without killing the application.

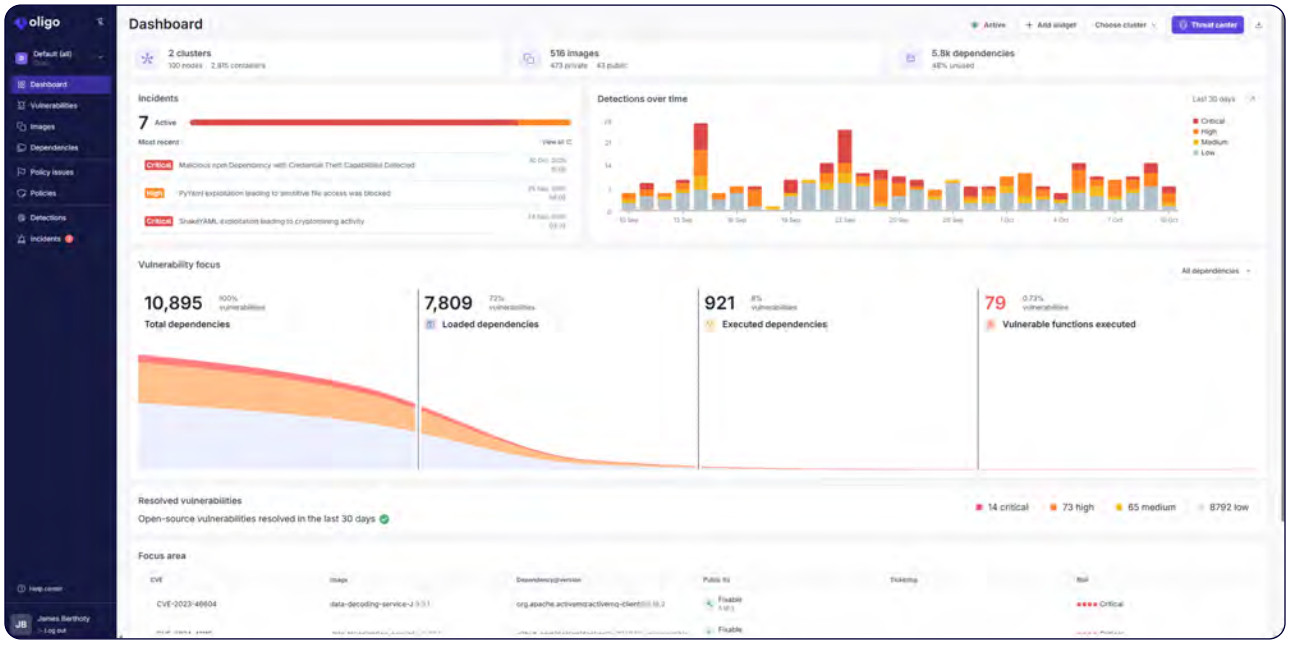

Oligo has demonstrated the effectiveness of their approach by revealing several real world attacks, especially in AI applications, that were detected with the assistance of their tool. As applications continue to evolve as the front door of an organization's attack surface, holistic cloud and workload protection that starts at the application layer is essential to the future of security.

Whether you're looking to prioritize only the vulnerabilities that matter, or looking for the best runtime workload defense you can get, Oligo is a great solution.

The Benefits of Oligo

Prioritize What Matters

With runtime function level reachability, Oligo tunes out the 99% of vulnerabilities that aren't applicable to your environment, and provides security operations teams the data they need to speak the same language with developers.

Stop Advanced Attacks

Oligo can block the most advanced attacks against cloud workloads, all without stopping the application.

Embrace AI with Confidence

Oligo extends runtime protection to AI applications, models, and agents, ensuring they operate securely and as intended.

The evolution of Palo Alto Network's cloud solution from Prisma Cloud to Cortex Cloud has proven to be a successful one, and far more than a rebrand. While Palo Alto Networks has long proven to be visionary in their acquisition strategies, providing features that customers are looking for as they begin looking for them, before Cortex Cloud there wasn't a cohesive glue uniting the products.

Over the past few years, Cortex Cloud has emerged as exactly that unifying solution, covering from code to cloud to SOC, presenting a strong argument for companies to buy the complete Palo Alto Networks security stack.

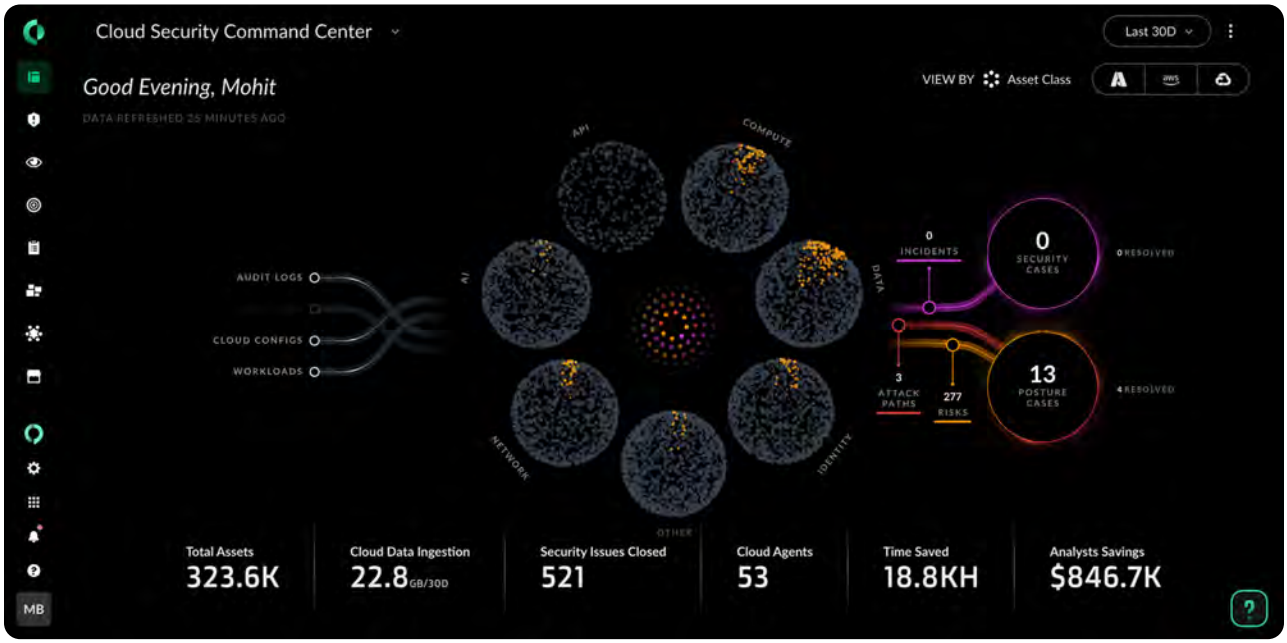

The transition to Cortex Cloud has radically changed how Palo Alto Networks approaches the cloud security market, creating a true better together story for the platform - from networks, to endpoints, to cloud, to applications, and now AI. With Cortex Cloud, Palo Alto Networks has recognized that security operations teams are the frontline for preventing, detecting, investigating, and responding to cloud threats. By offering CNAPP as a native extension of Cortex, Palo Alto Networks is delivering all of the relevant context from disparate tools to the teams who benefit from it the most. Cortex Cloud is one of the only CNAPP solutions that delivers a complete operating model for security across both cloud and SOC.

Cortex Cloud makes a strong case for teams to adopt the entire Palo Alto Networks ecosystem of products rather than one at a time, as every type of user benefits by having a unified data layer. CISOs get the reports they're looking for, the SOC gets the context they need to respond to threats in real time, AppSec teams gain centralized context, and vulnerability management teams know precisely what's critical and can be patched.

The Benefits of Palo Alto Network:

Threat Prevention

Intelligently block threats with complete AI-driven context from code to cloud to SOC.

Fewer Alerts

Effectively prioritize and fix issues across your environment by reducing duplicative scanners.

Reduction in MTTR

Immediately detect active incidents, and give responders all of the context they need to take immediate action.

tenable

Latio CTEM LEADER 2025

Latio CLOUD SECURITY ECOSYSTEM LEADER 2025

Tenable's Cloud Security solution, part of Tenable One, provides a comprehensive approach to discovering and managing risk across hybrid and multi-cloud environments. Tenable discovers toxic combinations across first and third party scanners, creating effective prioritization and remediation workflows allowing teams to fix what matters most - whether it lives on-premise, in the cloud, or on a workstation.

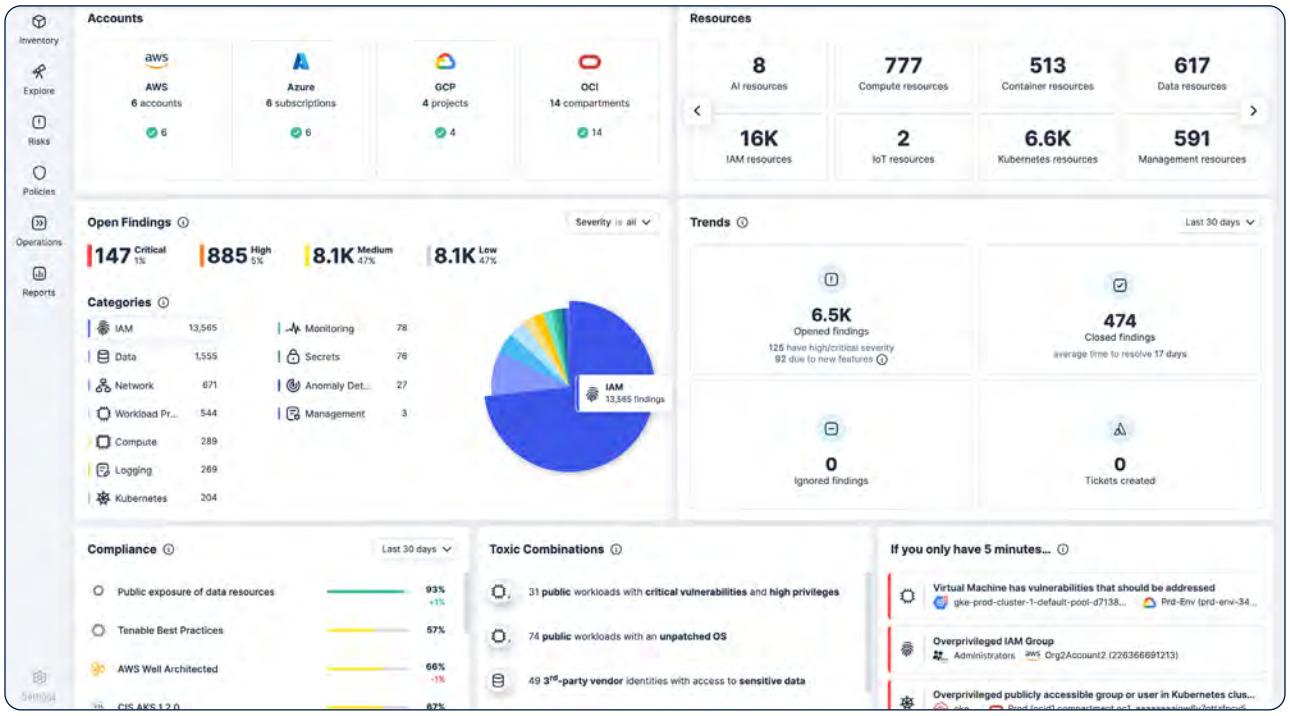

Organizations can accurately track vulnerability exposure timelines by combining all of their vulnerability data and SLAs into a single place, such as showing how quickly findings are getting resolved by various teams. Robust querying and reporting systems allow users to create custom insights, alerts, and workflows tailored to their unique context.

Tenable's integration of attack path analysis across different types of infrastructure enables teams to understand how attackers could exploit vulnerabilities no matter where they start. By combining deep telemetry with business context, Tenable drives meaningful insights, from detection through to remediation.

The Benefits of Tenable

Exposure Management

Tenable allows teams to find and reduce risk across their entire attack surface - from identities to workloads.

Leading Prioritization

Tenable's AI engine analyzes cloud security findings in the context of your environments, business criticality, and threat intelligence to surface risks that truly matter.

Unified Cloud Visibility

Tenable creates a unified asset inventory by continuously discovering resources no matter where they're hosted, or what they're hosted on.

WIZ · Upwind · DATADOG · orca SECURITY · sysdig · sweet.

oligo · Upwind · ARMO · sweet. · DATADOG

WIZ · ZAFRAN · JupiterOne · Qualys. · tenable · Axonius

cycode · ox · apiiro · WIZ · HEELER · aikido · DATADOG

oligo ARMO RAVEN dynatrace ZEST
aikido apiiro OX cycode
ZAFRAN RAD SECURITY HEELER JupiterOne
PROWLER uptycs sola intruder
Axonius PHOENIX SECURITY aqua Operant
Kodem MIGGO RoonCyber AccuKnox

maze CYERA Exaforce
Stream.Security Chainguard minimus
NOMA Pillar LEGION Formal
EDERA PJ SECURITY Operant strongdm
cogent tamnoon tines Mate

paloalto NETWORKS tenable Qualys.
SentinelOne CROWDSTRIKE Microsoft

Definitions

Cloud Security Leader

Cloud security leaders represent standalone cloud security solutions which attempt to be a single platform for end-to-end cloud security - from code to cloud to runtime. This requires support for all platform environments, from Windows to Containers, a variety of posture and vulnerability management capabilities, runtime protection, and some level of integration with CI/CD pipelines.

Cloud Security Ecosystem Leader

Cloud security ecosystem leaders represent competitive cloud security offerings that are meant to be purchased as part of a wider set of security solutions. These products are valued for how they fit into a customer's overall security strategy broader than what's traditionally considered "cloud security" alone.

Cloud Security Innovator

Cloud Security innovators represent companies who either currently have or are building towards innovations guiding the future of cloud security. These awards are for solutions that provide features that are technically innovative, empower teams with especially helpful user experiences, or go to market bundling of their services.

Cloud Security Segment Leader

Cloud security segment leaders represent solutions that are especially helpful as part of a larger cloud security stack, but don't align exactly as features of current cloud focused platforms. These leaders are selected as representatives of their segments due to either their focus on capabilities relevant to cloud security practitioners, or their leadership of a segment.

CADR Leaders

CADR Leaders are vendors building the future of workload security in the cloud through innovations combining deep visibility across all layers of a workload. These vendors are leading the charge making security events actionable for cloud security operations teams.

CTEM Leaders

CTEM leaders are building the future of doing vulnerability management at scale by taking data from multiple sources and unifying it into a single vulnerability management tool. These tools provide teams the ability to ingest, prioritize, and deliver remediations across several teams.

Code to Cloud Leaders

Code to cloud leaders use a combination of advanced data models to relate code to how it's deployed in production, mapping individual functions to their ultimate endpoints in production. These leaders map source code to deployed assets across various IaC types, from helm to terraform.

Latio

The only analyst firm that tests products, so you can find the right one.

Ever wonder: Am I using the right security tools for my business, or am I building the right product for the market?

Everyday companies are making decisions based on the information that is available to them, which is often incomplete and based on vibes rather than usage.

That's where Latio comes in.

Founded in 2023 by James Berthoty, Latio was built to solve a critical problem James was facing: there was no reliable, credible way to evaluate a vendor's capabilities until after an agreement was signed. Latio exists to make the buying and building processes better by getting accurate information to the most relevant teams.

We focus on the product, the practitioner, and the market rather than slides and hype cycles. We believe the greatest predictor of a great security tool and program is finding the right product fit for both vendors and buyers.

We are creating a future where every decision is based on tests, market insights, experience, and hard work, where it's easy to find the right product you're looking for.

Our mission is to help every team find the right security product. So we test every product, to make it easier for you to pick the right one.

A special thank you to everyone who has supported this mission, without you, none of this would be possible.

Learn more:

- latio.com
- Schedule a product briefing
- Schedule a security program sync
- Follow us