Latio

Application Security Market Report

2026

Table of Contents

Executive Summary

Whether you call it application security, product security, or DevSecOps, securing software is complicated. Today, practitioners are expected to manage a growing set of scanners, reduce large vulnerability backlogs, coordinate remediation across teams, and participate in architecture and threat modeling, often with limited headcount and little tolerance for noise.

AI is adding to this complexity, amplifying both the risks and opportunities in application security. AI assisted coding is reshaping how applications are built, deployed, and maintained. In parallel, the capabilities of platforms themselves are evolving with AI: features from autofix workflows, to false positive analysis, to scanning itself, are all radically changing product expectations.

This report is designed to help practitioners and buyers navigate the current application security landscape. It covers the transitions in application security over time, from waterfall development to DevOps to emerging AI code generation workflows. The report then breaks down every subcategory of scanner, the development of modern features, as well as how AI capabilities are changing functionalities we use today. We conclude with actionable buyer guidance that spans across SMB, mid-market, and enterprise environments.

Key Takeaways

- **Application security has largely consolidated into platform players.** The capability differences have more to do with user, integration and developer experiences than pure scanning functionalities.
- **AI-native static analysis and business logic detection are the most immediately meaningful changes in Application Security detection capabilities.** These new scanners are capable of detecting entirely new categories of vulnerabilities which have traditionally been reserved for manual review.
- **Application security evaluations should focus on usability and backlog reduction more than specific scanner functionalities.** Tool evaluations should be guided by the time to fix an issue, rather than the number of issues detected.
- **ASPM as "management without scanning" has largely collapsed into broader vulnerability management and exposure programs.** ASPM is shifting into continuous threat exposure management, or universal vulnerability management.
- **Securing AI-generated code is still an open market with unclear best practices.** General approaches involve giving organizational context to agents, and having secure code reviews in pipelines, but this field is rapidly changing.
- **Supply chain security is expanding towards malware, package health, and secure-by-default consumption patterns.** CVE detection alone is not enough for modern supply chain security.

2026

APPLICATION SECURITY SURVEY RESULTS

The Gap is Runtime Context

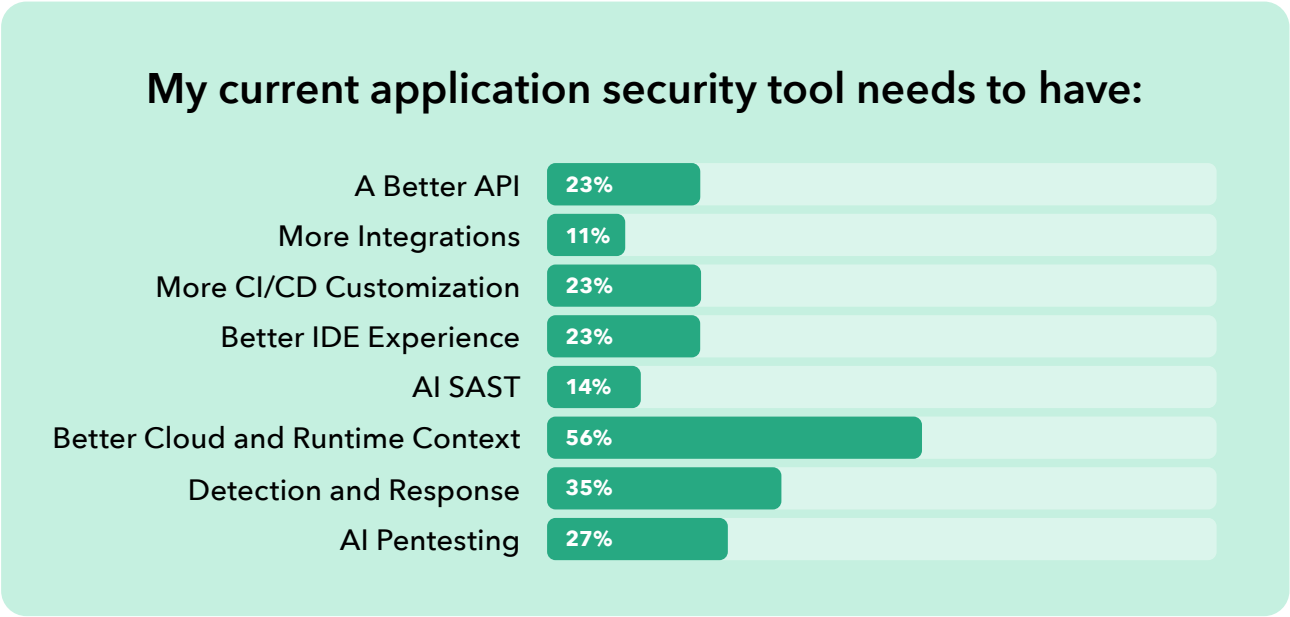

AI may be driving industry excitement, but usability remains the primary focus for immediate feature enhancements. Specifically, better APIs, IDE state tracking, and integration experiences were highly requested from practitioners. The most requested feature by far was better cloud and runtime context, because teams want to better prioritize and determine the truth of a particular finding.

When reviewing the results alongside the preference to separate runtime and code experiences, this finding highlights the need for strong integration between tools managed by different teams to support vulnerability triage, false positive analysis, and delivering fixes.

Open Source or Not - Results Matter More

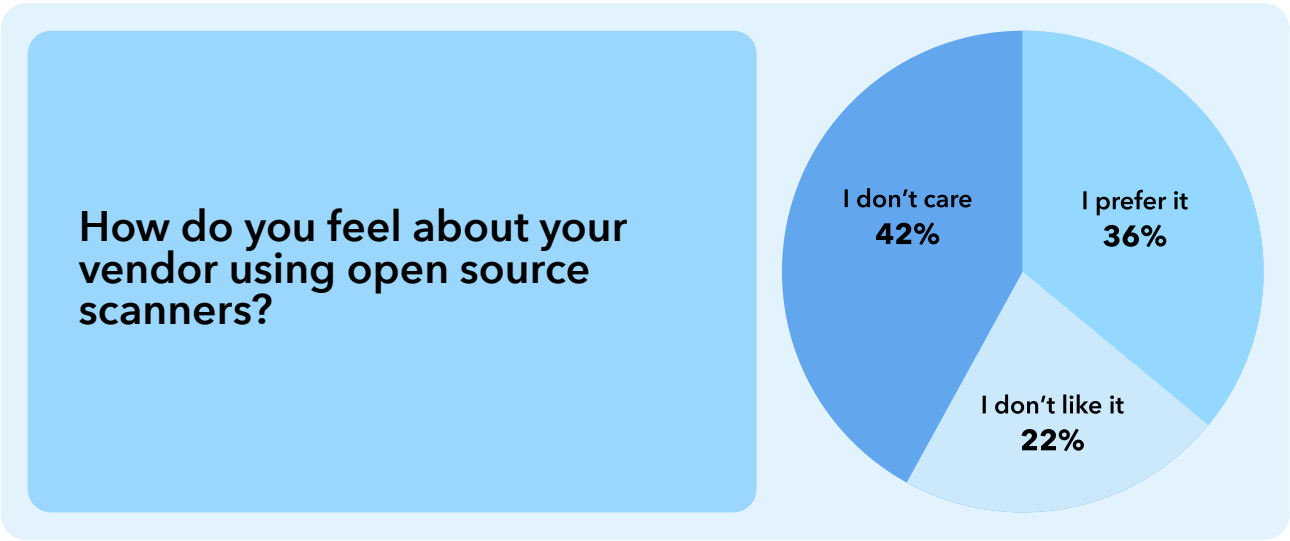

Increasingly, security operations teams are being tasked with handling runtime application security alerts, though ownership remains mixed. Cloud and product security teams are also responsible for these alerts in many organizations, and in some cases developers are involved as well. It is worth highlighting that these management functions no longer sit with traditional network security teams, and have instead shifted toward broader security operations and cloud teams.

This speaks to the growing importance of application-layer visibility for security operations teams, as they are increasingly expected to handle more contextual and technical alerts.

The AI Features That Matter

Which AI feature are you the most excited about?

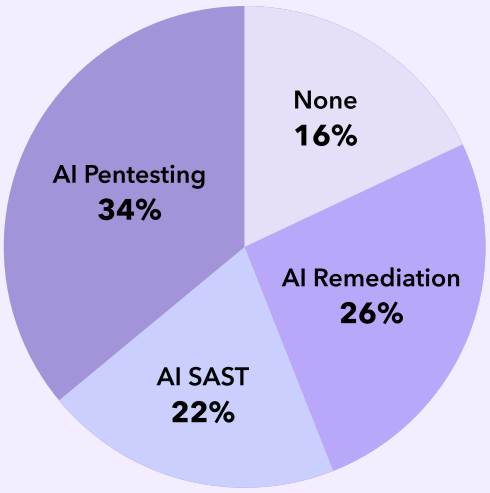

Application security practitioners are ready to adopt the latest AI features - from prioritization to static code analysis. The relatively low number of “none” responses is surprising, because it suggests that practitioners in this area are less skeptical about AI’s ability to transform day to day workflows than in other categories. Seeing the power of AI code generation directly makes practitioners in this category more open to AI adoption.

While most AI features garnered similar levels of excitement, AI pentesting pulled a slight lead, indicating that teams are excited about being able to make continuous pentesting a reality. Overall it's clear that AI is transforming the capabilities of existing testing methodologies and teams are ready for it.

Cloud and Infrastructure

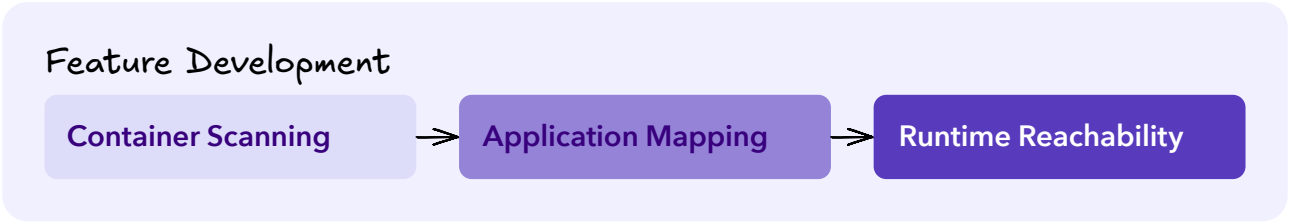

When tools start scanning containers, they open Pandora's box of complexity. Once you scan a container, you're quickly involved in its entire lifecycle - scanning the Dockerfile, the build process, the container registry, and the running container. Additionally, many teams don't realize they're getting duplicative results from their container scanner and their SCA scanner, adding another layer of confusion.

The responsibility of scanning containers has landed squarely on the cloud security side of the fence; however, typically application security tools do a better job of delivering the results to the teams that can actually do the fixing. This tension has led most large application security providers to provide container scanning in some form or another as part of their platform.

Infrastructure as code has the same problem, landing as an infrastructure team responsibility, but being more native to an application security integration workflow. This functionality is generally the most underbaked, because while basic scanning is widely accessible in open source tooling, making it work for your environment requires a high degree of customization. Scaling IaC analysis quickly leads to high levels of complexity, as building drift detection and support for custom modules become must-have capabilities. Other tools also include basic CSPM scanning as part of their platform in order to function as an all in one platform for smaller companies.

All in all, cloud security capabilities have long faced an awkward overlap with application security. On one hand, for organizations that are heavily invested in infrastructure as code and have strong deployment guardrails, an application-security-first approach works well. However, many companies operate in far more cobbled-together cloud environments, consisting of a mix of manually deployed and infrastructure-as-code-managed assets. These organizations tend to value the runtime monitoring capabilities of cloud security providers more than treating everything as code. Nonetheless, having these capabilities closely tied together creates many benefits to prioritization, drift detection, and contextualizing findings.

Management Tools

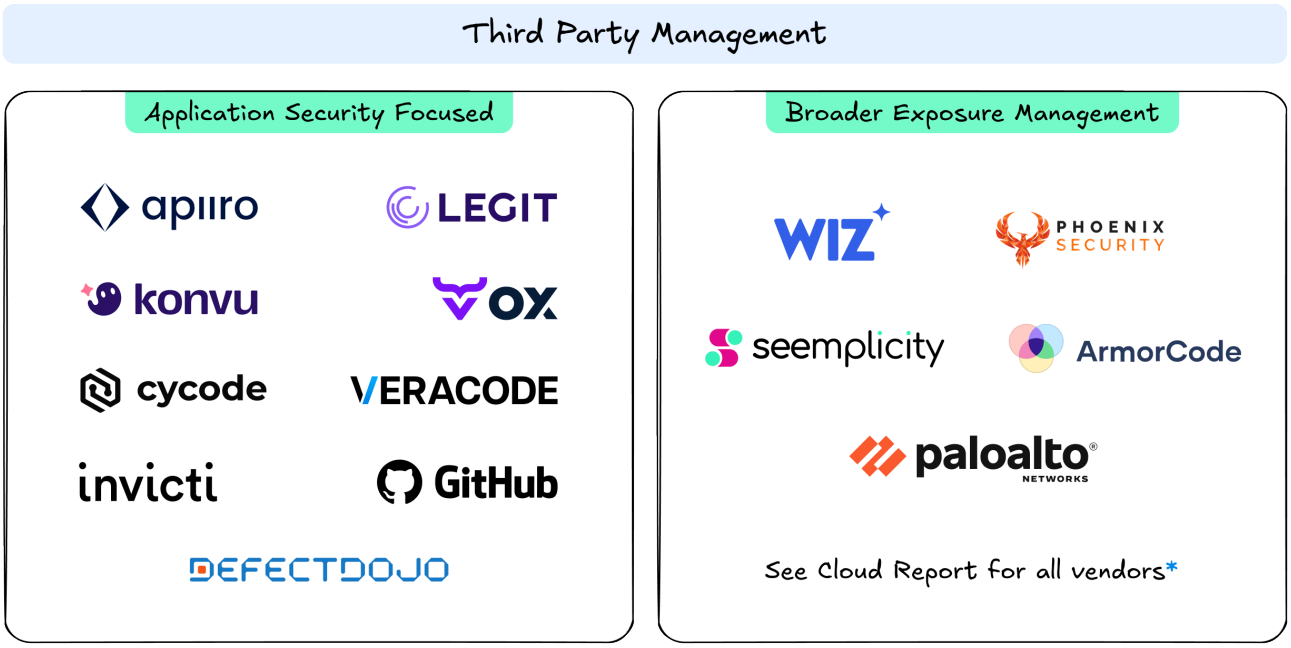

In our cloud security report, we made it clear that vulnerability management solutions are consolidating under CTEM functionalities. ASPM vendors, as traditional analyst firms described them, were a transitional tool for consolidating the findings of different application scanners. Consolidating findings across containers and SCA scanners was important, but the broader de-duplication and contextual project has more to do with infrastructure than code.

The threat exposure management category (vulnerability management 2.0) remains top of mind for security leaders entering 2026, but they don't necessarily expect their application security scanner to be that same tool. Developer experience and finding quality remain paramount, with broader vulnerability management goals typically owned by different teams.

Even as the market shifts toward more unified exposure and vulnerability management tools, there can be meaningful benefits having an orchestration layer on top of your application security scanners. For organizations with tens of thousands of GitHub repositories, monitoring pipeline coverage, understanding which applications map to which teams, and orchestrating scanning tools are massive challenges.

Tools in this category take different approaches to solving these enterprise problems. Some, like Apiiro, focus on a more application-centric model. Others, such as Phoenix Security, lean more heavily into cloud and container-centric views. Legit Security emphasizes CI/CD coverage, while Palo Alto Networks brings a cloud-centric perspective through its broader platform. Each of these approaches deliver value for large organizations that need clear visibility into security coverage across sprawling development environments.

BUYER'S GUIDE

Picking the right tool varies and depends on individual tech stacks and priorities. The Latio Buyer's Guides exist to provide clear and actionable considerations when searching for the best available tools.

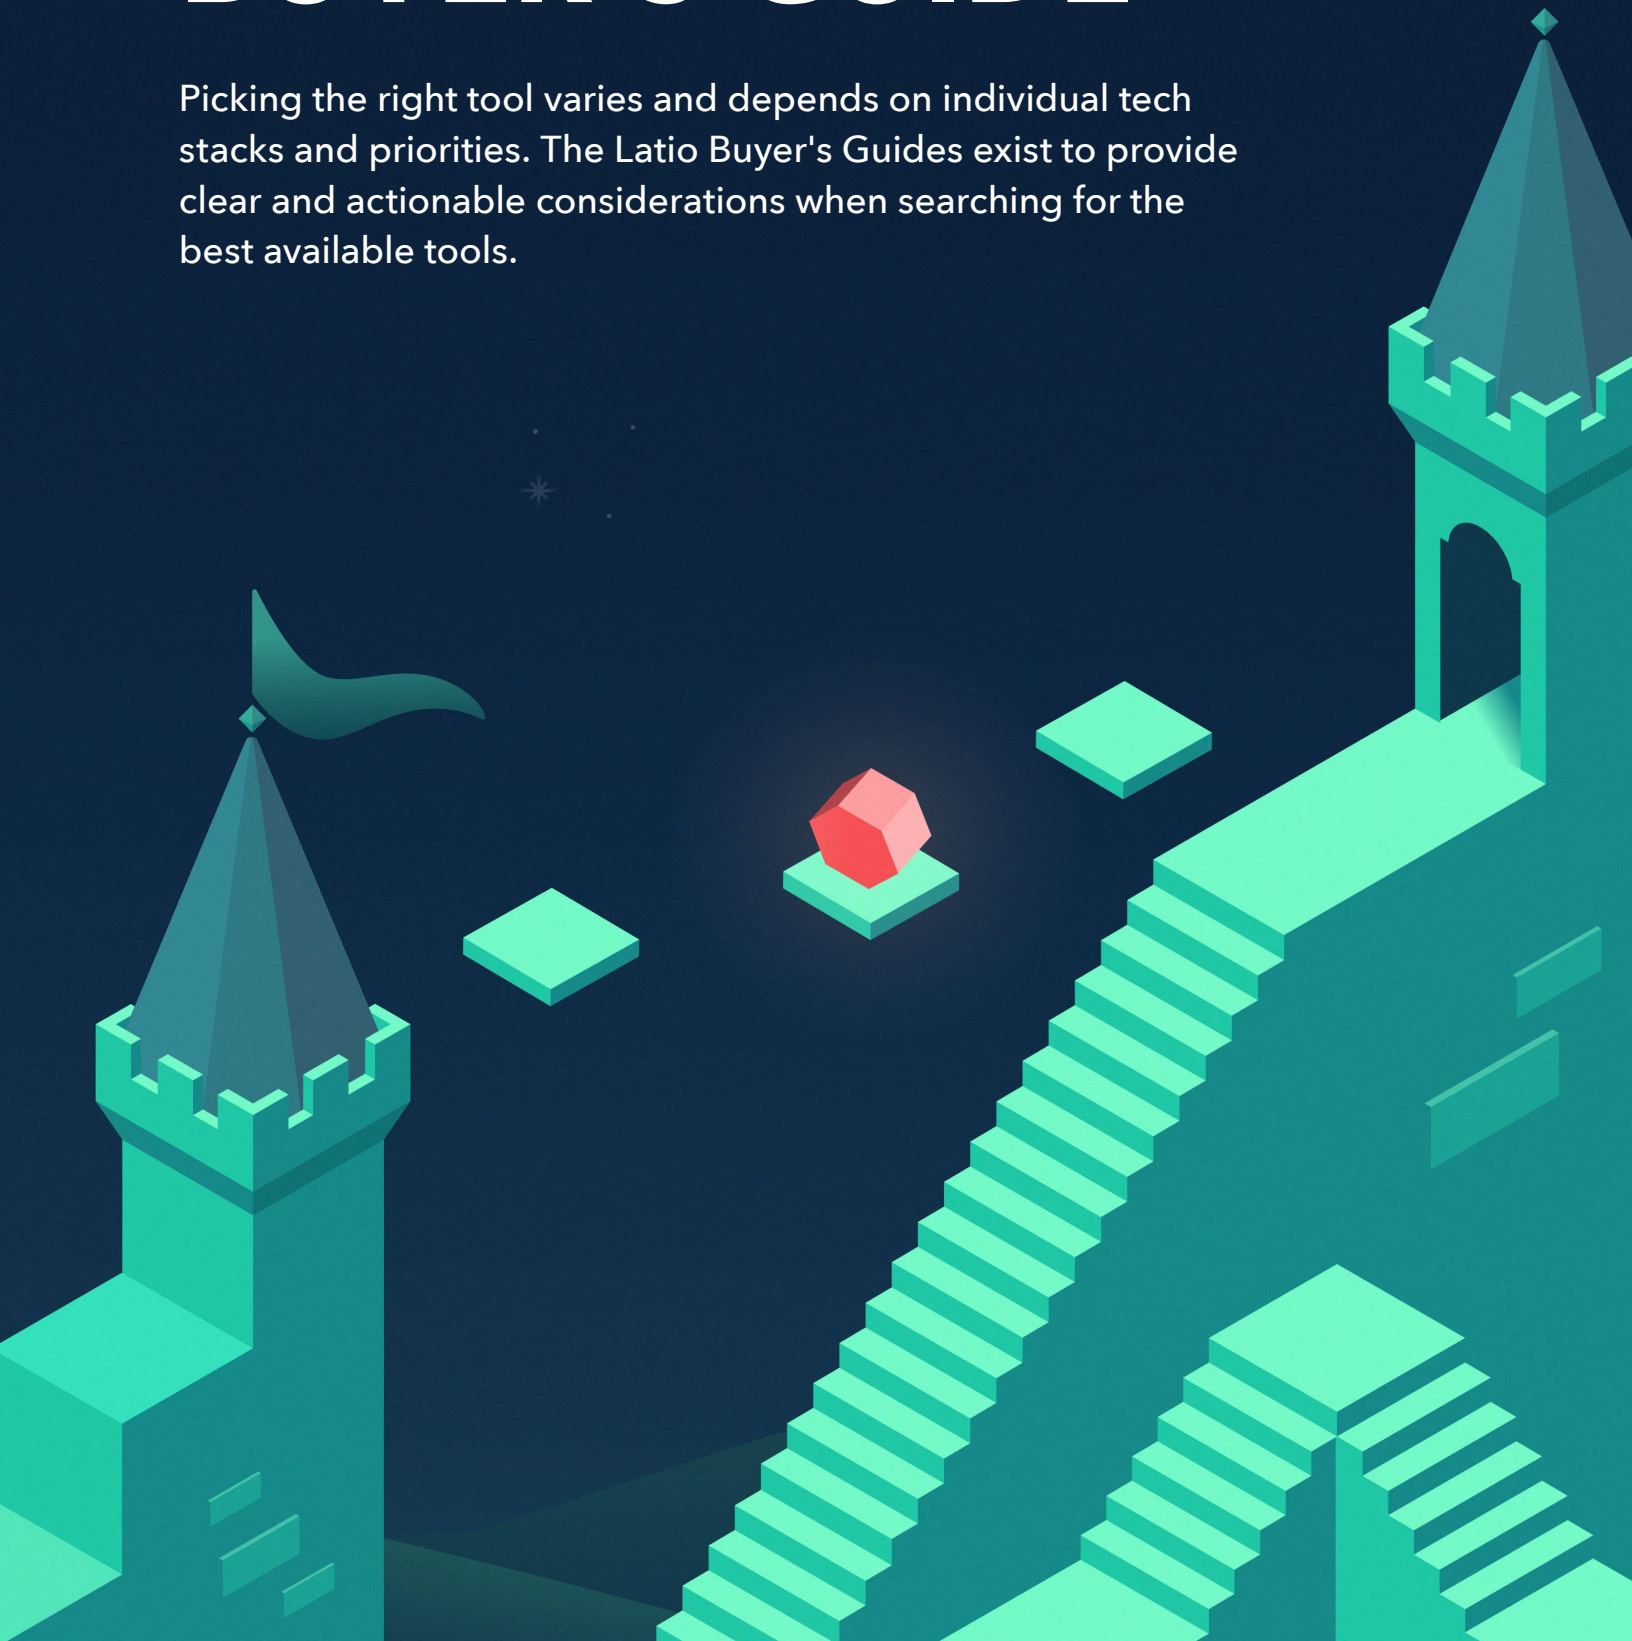

SMB and Mid-market Buyer's Guide

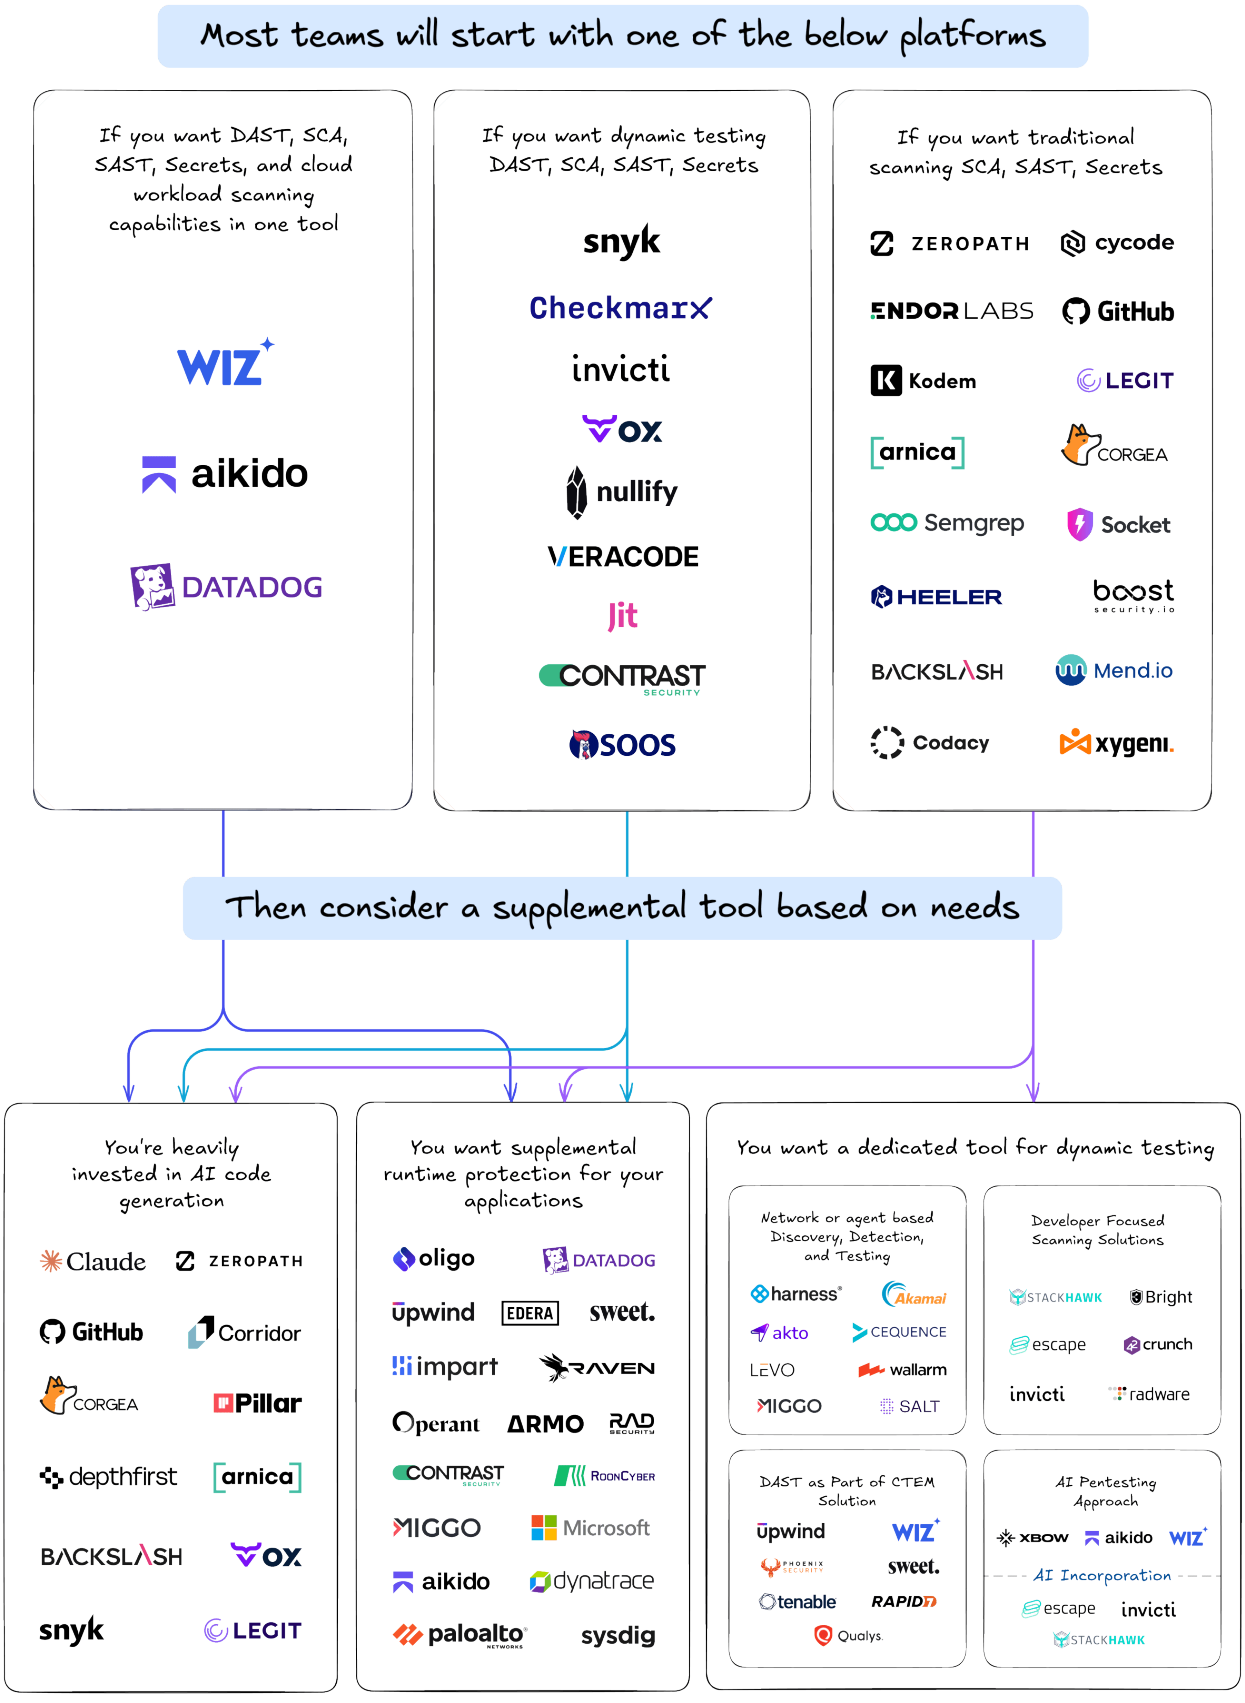

Enterprise Buyer's Guide

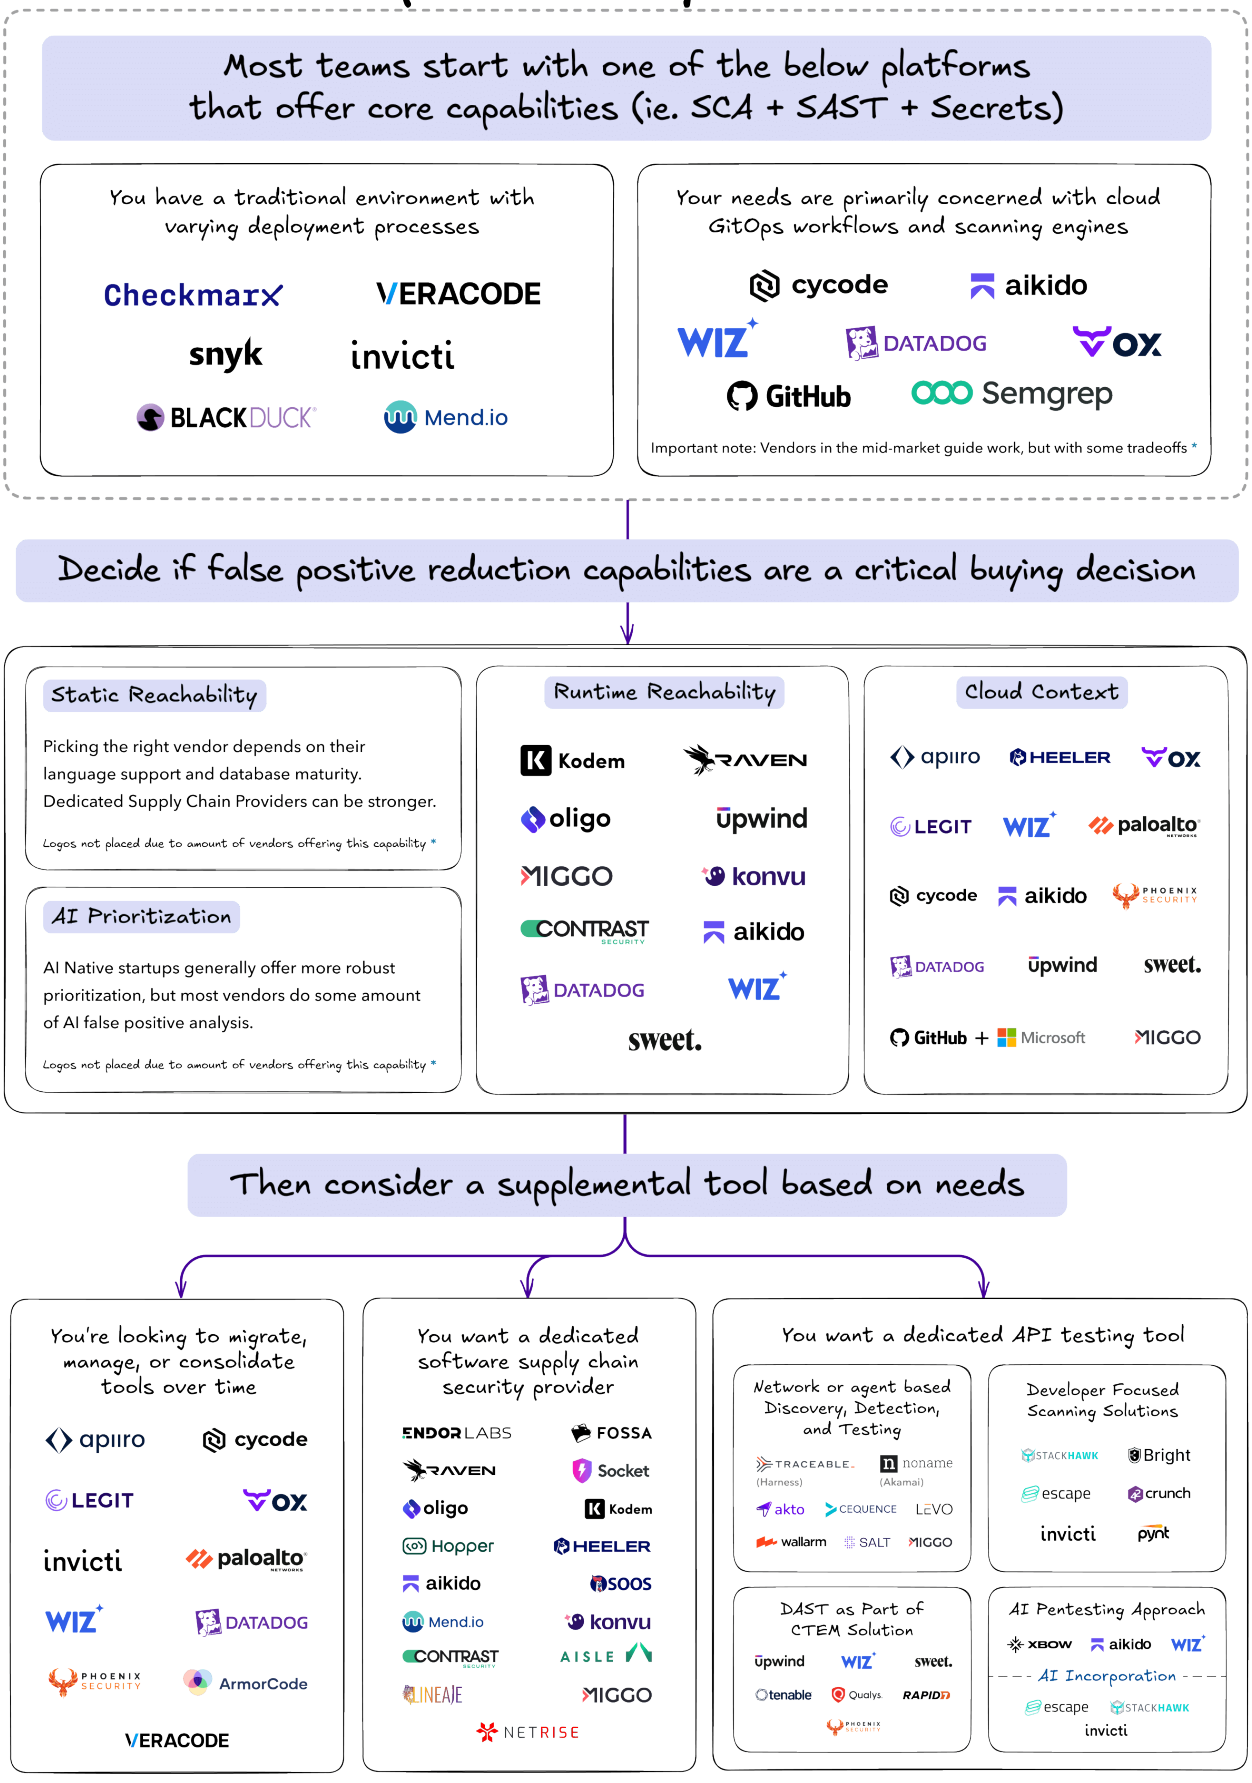

Enterprise

Choosing a Platform

Application security in large enterprises looks significantly different due to the high number of legacy applications, diverse coding languages, deployment models, and development teams that have built and managed their own pipeline solutions over time. Organizations with these distributed environments often also have strict compliance requirements and long-standing vendor relationships, which require customizations of SBOMs and multiple scanning types as binaries are built and deployed.

These enterprises differ fundamentally from cloud-native startups and SaaS companies, where standardized developer platforms are core to the organization. They typically operate on more unified cloud-native architectures, with deployment processes built from the ground up or fully migrated using consistent, standardized approaches.

The first step of our buyer's guide is intended to highlight this distinction between sprawling developer environments and fully modernized ones. On one side, there are vendors that support a wide range of testing methods and workflow requirements, designed for large enterprises and their often complex compliance needs. On the other side are vendors that function well as all-in-one application security solutions, either through deep customization, strong extensibility into existing platforms, or broad coverage of the many use cases enterprises require.

Additionally, the vendors listed in the cloud native architecture section are not meant to be exhaustive. Nearly every company shown in the mid-market diagram can fit into an enterprise stack, with different trade-offs depending on organizational priorities.

Approaches to False Positive Reduction

Vulnerability prioritization and reduction is a massive challenge on its own for large enterprises. For this reason, static function-level reachability and AI-driven prioritization are important differentiators, but their maturity varies widely by vendor and language, and often depends on proprietary vulnerability databases that are difficult to assess externally. In general, static reachability requires longer CLI-based scans to produce strong results, particularly for statically typed languages.

Vendors have also taken different approaches to AI prioritization. Some build deep, proprietary indexes of customer codebases and run sophisticated prioritization logic, while others rely on far more superficial techniques, such as prompting general-purpose LLMs with a finding and its surrounding code.

PHOENIX SECURITY

Latio
APPLICATION SECURITY MANAGEMENT LEADER
2026

Enterprise vulnerability management programs fail because scanners rarely answer the three questions that drive remediation: who owns it, where is it running, and what is the fastest, lowest-impact fix. For organizations that prioritize actionable attribution and operate in regulated industries, many platforms overlook the customization details required to operationalize remediation. Phoenix Security has consistently stood out for its attention to enterprise-level details that make vulnerability management work at scale.

One of Phoenix's primary differentiators is its ability to align asset attribution, code-to-cloud correlation, and reachability analysis with business goals, regardless of which scanning tools feed into the platform. This allows teams to maximize the value of their existing scanners rather than being forced to adopt entirely new ones. These benefits extend to providing AI prioritization and remediation.

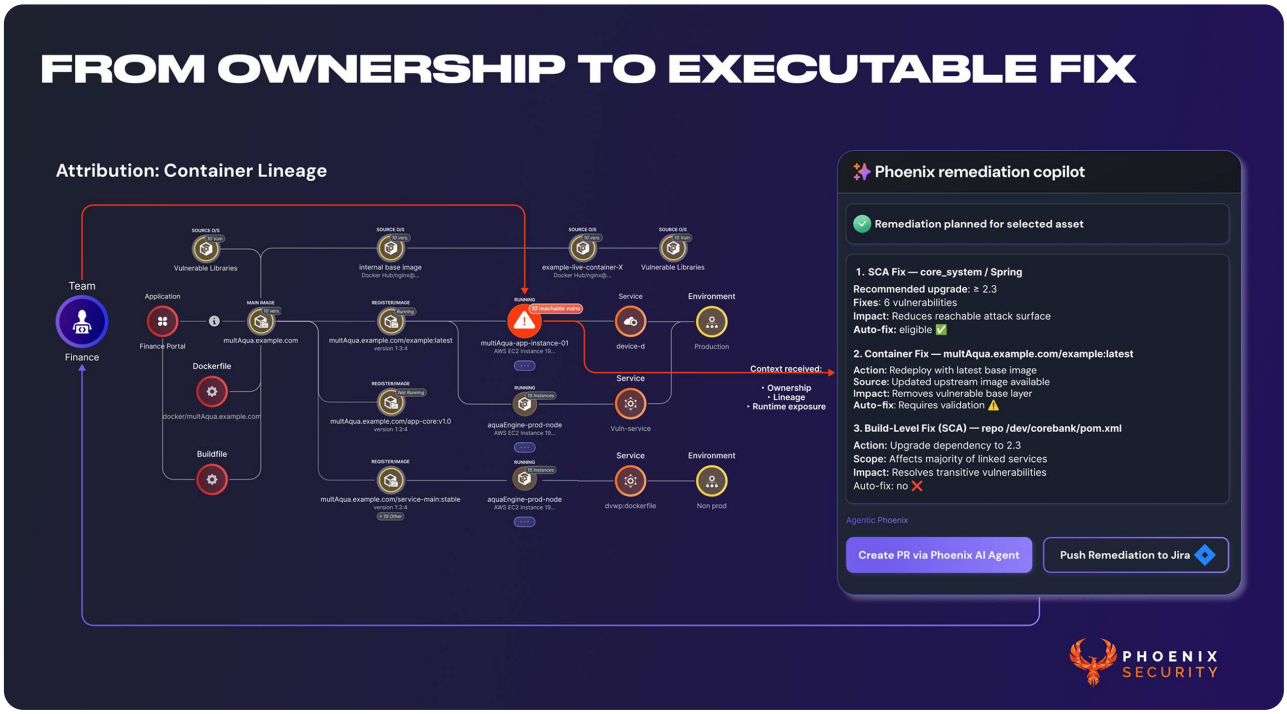

Phoenix's attribution model is designed for distributed enterprises where ownership changes, services are ephemeral, and data resides across multiple systems. Phoenix supports PYRUS CMDB-as-code patterns and integrates with enterprise sources of truth to ensure accurate ownership. Phoenix also uses AI across the platform, from vulnerability enrichment to remediation. These features make Phoenix a strong option for enterprises seeking a scalable, configurable vulnerability management solution.

The Benefits of Phoenix Security:

Attribution at Enterprise Scale

Phoenix's attribution CMDB enables teams to manage asset ownership across their entire lifecycle programmatically.

Tool-Agnostic Reachability

Reduce false positives by applying native provenance, lineage, and reachability analysis on top of existing scanners, drastically reducing vulnerability counts.

AI Prioritization and Remediation

Phoenix's AI systems analyze threat intelligence data to predict the threat types most likely to lead to exploitation, and provide precise remediation guidance.

WIZ GitHub DATADOG

cycode aikido

apiiro VOX

invicti Checkmarx snyk

Semgrep VERACODE

apiiro WIZ LEGIT

cycode PHOENIX SECURITY ArmorCode

Latio

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

Lativ

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