cloud security alliance®

MIGGO

Survey Report

2026 State of Modern Application & AI Security

Acknowledgments

Lead Author

Hillary Baron

Contributors

Marina Bregkou
Josh Buker
Ryan Gifford

Special Thanks

Eliana Vuijsje
Arye Zaks
Noa Gur Arieh
Ofer Gayer

Design

Victoria Gavlinskaya

About the Sponsor

Miggo Security, leader in AI Runtime Security and ADR, delivers exploit mitigation for AI and modern applications. While attackers weaponize vulnerabilities at machine speed and patching takes weeks, Miggo closes this Patch Gap in minutes with precision mitigation engineered for the exact exploit path. Powered by patented DeepTracing™, Miggo reverse-engineers each exploit primitive and maps it to live runtime and then generates, validates, and deploys a targeted mitigation on the vulnerable path without interrupting a single engineering sprint. Security teams cut vulnerability backlog by over 95% and mitigate over 90% of exploitable risk in under an hour. Miggo has also been awarded among others, the Frost & Sullivan Product Innovation Award 2025 and Gartner Cool Vendor 2025.

Table of Contents

Executive Summary

Application security programs have become highly effective at finding vulnerabilities earlier in the software lifecycle. Yet production environments remain where risk consistently materializes. In 2025 alone, the National Vulnerability Database (NVD) recorded more than 40,000 CVEs, while VulnCheck reported exploitation activity increasingly following disclosure within days rather than weeks. As exploit timelines compress and AI accelerates both vulnerability discovery and exploit generation, the challenge is no longer whether organizations can identify risk — it is whether they can mitigate it before it becomes operational.

Built on survey data from more than 900 cybersecurity leaders and practitioners, this report examines where application security programs are breaking down in practice, why production environments continue to absorb incidents despite mature pre-production controls, and how organizations are adapting as runtime risk becomes harder to interpret, prioritize, and contain.

1. Known Vulnerabilities & the Patch Gap Are Driving Real-World Incidents

Application security incidents are overwhelmingly tied to vulnerabilities organizations already knew about rather than unknown threats. Eighty percent of organizations experienced at least one incident involving a known vulnerability in the past year, while most respondents report remediation timelines for critical vulnerabilities measured in days rather than hours. Organizations taking longer to remediate reported substantially higher incident rates, reinforcing that the primary challenge is not discovering vulnerabilities, but reducing the exposure window between identification and mitigation.

2. Runtime Is the Breach Battlefield — Incidents Slip Past Pre-Production Controls

Pre-production tooling and shift-left practices are widely adopted, yet production incidents remain pervasive. Nearly half of organizations report incidents tied to vulnerabilities that were identified before release but still reached production, while another large segment report vulnerabilities that pre-production controls failed to identify entirely. The findings suggest that expanding detection coverage upstream does not necessarily translate into reduced runtime exposure when mitigation and enforcement capabilities remain concentrated outside production environments.

3. AI Is in Production, Security Is in Post-Mortem

AI-powered application components are already operating in production across most organizations, but runtime oversight remains largely retrospective. Most respondents rely on post-incident auditability or incomplete logging rather than real-time runtime visibility. As AI systems introduce more dynamic and less predictable behavior into production environments, organizations are increasingly managing incidents after the fact rather than maintaining continuous enforcement and intervention capability while activity is occurring.

4. The Main Bottleneck for Protection is Proof of Exploitability

Organizations identify exploitability validation and runtime context as the most significant constraints in production security operations. The dominant challenge is distinguishing genuinely exploitable vulnerabilities from theoretical findings, while the most desired capability is proof that a vulnerability is reachable and exploitable within a specific production environment. Respondents consistently prioritize runtime evidence and contextual validation over additional staffing or broader scanning coverage, indicating that confidence in prioritization has become a larger operational issue than visibility alone.

5. The Will to Block Exists But Trusted Mitigation is Missing

Organizations broadly support stronger runtime mitigation and virtual patching capabilities, but most do not trust current controls to safely enforce automated blocking in production. While WAFs and similar technologies are widely deployed, the majority of organizations operate them conservatively due to concerns around false positives, lack of application context, and the risk of disrupting business-critical functionality. The findings point to a gap between enforcement intent and enforcement confidence rather than a lack of interest in runtime protection.

6. Investment Intent Is Turning Toward Runtime Security

Security investment remains weighted toward pre-production controls, but organizations are increasingly recognizing the need for runtime visibility and defense. A substantial share of respondents plan to increase investment in runtime monitoring and protection over the next two years, reflecting growing concern that existing security models were designed for a slower threat environment. As AI-assisted exploit development and vulnerability discovery continue to compress remediation timelines, organizations are beginning to shift focus from identifying vulnerabilities earlier to managing and mitigating them more effectively in production.

Takeaway

Application security programs have achieved broad maturity in vulnerability discovery and pre-production testing, but the findings in this report suggest that production environments remain the decisive layer where exposure turns into operational risk. The central challenge is no longer whether organizations can detect vulnerabilities, but whether they can validate, prioritize, and mitigate them quickly enough once deployed. As AI-driven application behavior becomes more dynamic and exploitation timelines continue to compress, runtime visibility, exploitability validation, and trusted enforcement are becoming foundational requirements for modern application security operations.

Key Findings

Organizations have built mature capabilities for earlier lifecycle phases — scanning code, testing before release, shifting security investment left — yet production environments remain where risk consistently materializes. The challenge is no longer whether vulnerabilities can be detected; it is whether teams can act on that knowledge before those vulnerabilities are weaponized into working exploits. That pressure is increasing: the National Vulnerability Database (NVD) recorded more than 40,000 CVEs in 2025 alone, while VulnCheck's State of Exploitation: 1H 2025 Vulnerability Exploitation & Threat Activity Trends found that exploitation activity increasingly follows disclosure within days rather than weeks or months.

Frontier AI is compressing that window further: autonomous exploit generation against production code has moved from research demonstration to operational use, and AI is increasingly being applied to vulnerability discovery at industrial scale. As application architectures grow more dynamic and AI-powered components become embedded in live environments, the conditions that determine remediation velocity and enforcement confidence are changing faster than most security programs have adapted. The questions at the center of application security are moving from discovery to execution.

Key Finding 1:

Known Vulnerabilities & the Patch Gap Are Driving Real-World Incidents

Incidents are not coming from the unknown; they are coming from the unresolved. The majority of organizations have experienced application security incidents tied to vulnerabilities their teams had already identified. Detection is not the failure point. The failure is in what happens between identification and remediation: a window that, for actively circulating vulnerabilities, stays open long enough to be exploited.

The data on incident prevalence leaves little ambiguity. In the past 12 months, 36% of organizations experienced an application security incident involving an already-known vulnerability multiple times, and another 44% experienced it at least once. Only 19% report no such incident. When more than four out of five organizations have been in this position within the last year, it describes a structural condition across the industry rather than a pattern of isolated failures.

Figure 1: Percentage of Orgs that Experienced an AppSec Incident Involving a Previously Known Vulnerability in the Past 12 Months

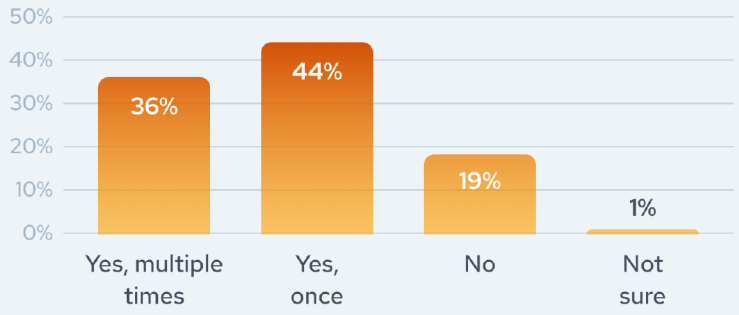

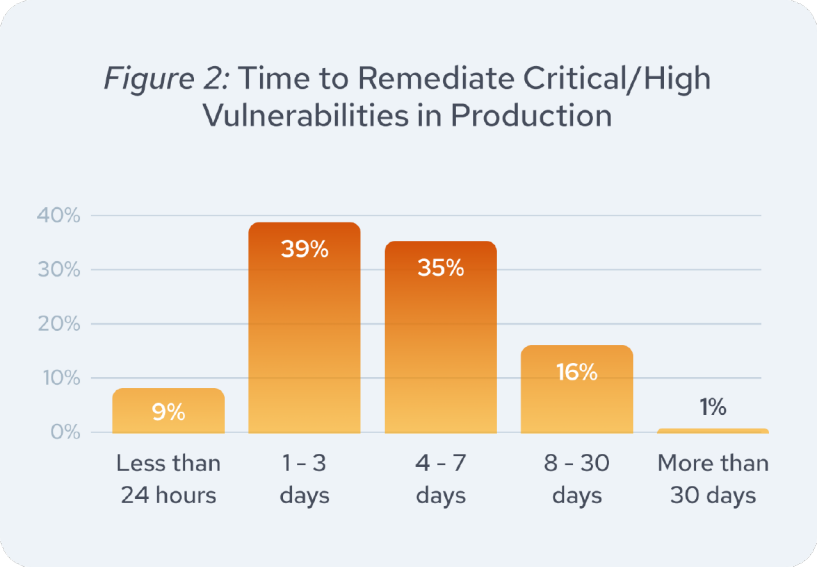

Remediation timelines, also referred to as the patch gap, explain the mechanism. For critical and high-severity vulnerabilities, 39% of organizations take 1–3 days to remediate in production and 35% take 4–7 days. Just 9% achieve remediation within 24 hours. At 1–7 days for 74% of respondents, even the faster end of that range leaves a meaningful window for exploitation, particularly for vulnerabilities already in active use by threat actors. Part of what extends those timelines is prioritization uncertainty — 34% of organizations cite disagreement on vulnerability relevance or exploitability as a contributing factor to delayed remediation, and 20% point to insufficient production context to assess impact safely (Figure 24).

The data also shows remediation speed is associated with incident outcomes. Among organizations taking 4–7 days to address critical and high-severity vulnerabilities in production — the second most common timeframe — 97% experienced a known-vulnerability incident in the past 12 months. That rate drops to 64% among those remediating within 1–3 days. Put differently: most organizations that take several days to close a known exposure end up experiencing it as an incident.

97%

of organizations that take 4–7 days to remediate critical vulnerabilities — the modal patch speed — **experienced a known-vulnerability incident in the past 12 months**

Security teams know what needs to be fixed; the gap is in how quickly they can close it and what to do between discovery and remediation to mitigate. This is an exposure window problem, not a detection problem. For organizations that have invested heavily in scanning, alerting, and discovery tooling, this finding reframes where additional investment changes outcomes: the risk does not live in what goes unnoticed, but in what remains unresolved. Where those incidents occur, and whether the controls designed to prevent them are positioned at the right lifecycle layer, is what the next finding examines.

Key Finding 2:

Runtime Is the Breach Battlefield – Incidents Slip Past Pre-Production Controls

Pre-production security, the current shift-left approach of application security, has not solved the production problem. Despite widespread adoption of build-time tooling and shift-left practices, the overwhelming majority of organizations experienced application security incidents that reached production in the past year — including cases where the underlying vulnerability had already been flagged before release.

Figure 3: Production AppSec Incidents Relative to Pre-Production Identification

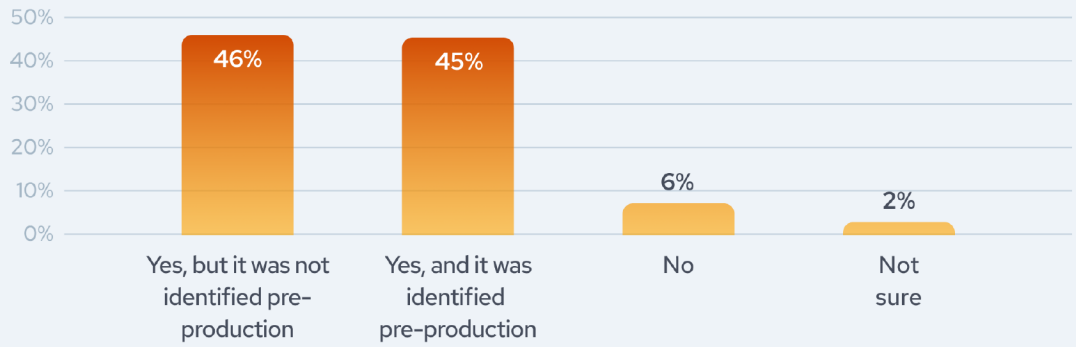

The scale of production bypass is difficult to attribute to gaps in pre-production coverage alone. Of organizations experiencing a production application security incident in the past 12 months, 46% report the issue was not identified in pre-production at all. Another 45% report the incident involved a vulnerability that was identified pre-production but still reached production. Six percent report no such incident. The two figures point to distinct failure modes. The 46% reflects vulnerabilities that pre-production controls did not flag — a coverage gap that shift-left approaches have not fully closed. The 45% represents a different problem: vulnerabilities identified before release that reached production anyway, a failure of remediation rather than detection.

Figure 4

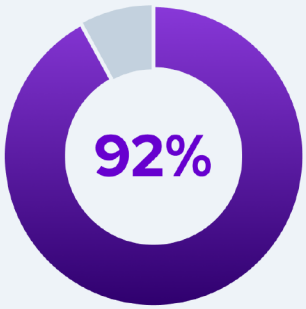

of organizations prioritizing pre-production identification had a known-vulnerability incident in the past 12 months

Figure 5

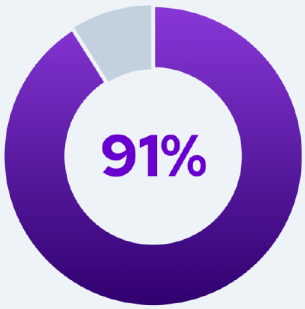

of 'Very Confident' respondents experienced a production incident in the past 12 months

The pattern holds across the organizations most invested in pre-production. Among respondents who prioritize identifying risks before deployment when evaluating new application security investments, 92% experienced a known-vulnerability incident in the past 12 months. Confidence in strategy does not change the picture: even among respondents who describe themselves as very confident their current application security approach will hold up as AI-driven attack surfaces expand, 91% experienced a production bypass incident in the past 12 months. Greater confidence and deeper investment in upstream tooling are not, on their own, associated with lower incident rates.

Conventional tooling adoption makes the incident rate harder to explain away. Sixty-five percent of organizations use Static Application Security Testing (SAST), 31% use Web Application Firewalls, and 26% use Dynamic Application Security Testing (DAST). These are mature, widely adopted controls. Yet 80% of respondents experienced at least one application security incident in the past year. Broad pre-production coverage and near-universal incident exposure are coexisting, which suggests that catching vulnerabilities earlier in the lifecycle does not translate directly into preventing their exploitation at runtime.

Risk is concentrated in production, but the controls and the investment are concentrated upstream — a misalignment the data on investment intent makes clear. For organizations deepening pre-production capabilities without a parallel commitment to runtime enforcement, additional investment may improve detection coverage without proportionally reducing exposure. The exposure gap is not a capability gap — it is a location gap. That gap is widening as AI-powered components enter the same production environments where conventional security investment has not yet followed.

Key Finding 3:

AI Is in Production, Security Is in Post-Mortem

AI-powered application components are no longer an emerging capability — they are an operational reality. **Seven in ten organizations already are running AI-powered application components in production,** though the deployment picture is bifurcated: of that group, an equal share report active security concerns as report limited or no concerns. The oversight infrastructure most organizations have built to manage these deployments, however, was not designed for AI's pace or behavior. Post-incident reconstruction has become the dominant mode of AI runtime governance at a moment when the speed and autonomy of these components make after-the-fact analysis increasingly insufficient as a control.

Figure 6: Current Use of AI-Powered Application Components in Production

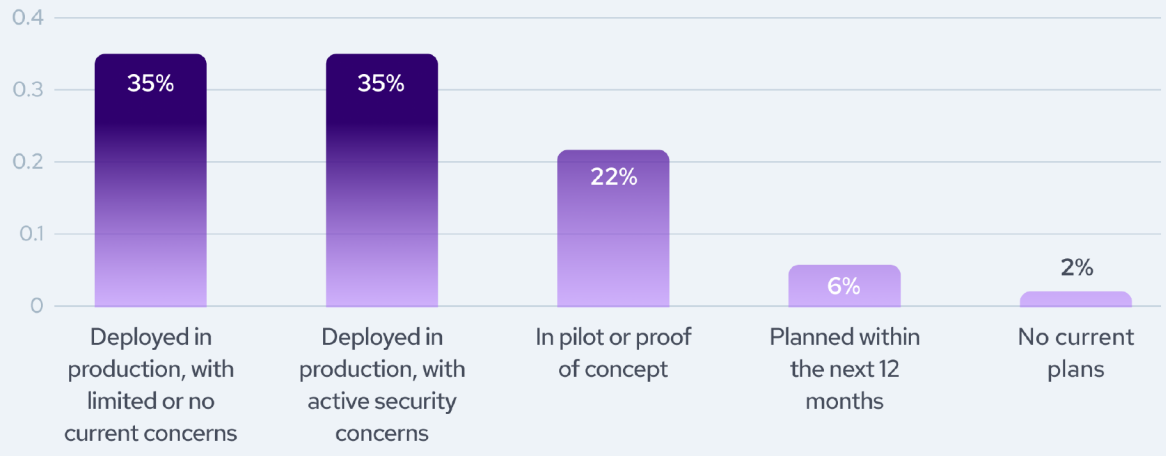

AI adoption in production is substantial and already bifurcated between confidence and concern. Thirty-five percent of organizations report AI components deployed with limited or no security concerns, while an equal 35% report production deployments with active security concerns. Another 22% are in pilot or proof-of-concept phases. As those pilot deployments mature into production, the governance gap the data identifies will apply to a broader share of enterprise environments.

Figure 7: Organizations' Visibility into Runtime Behavior of AI-Powered Application Components

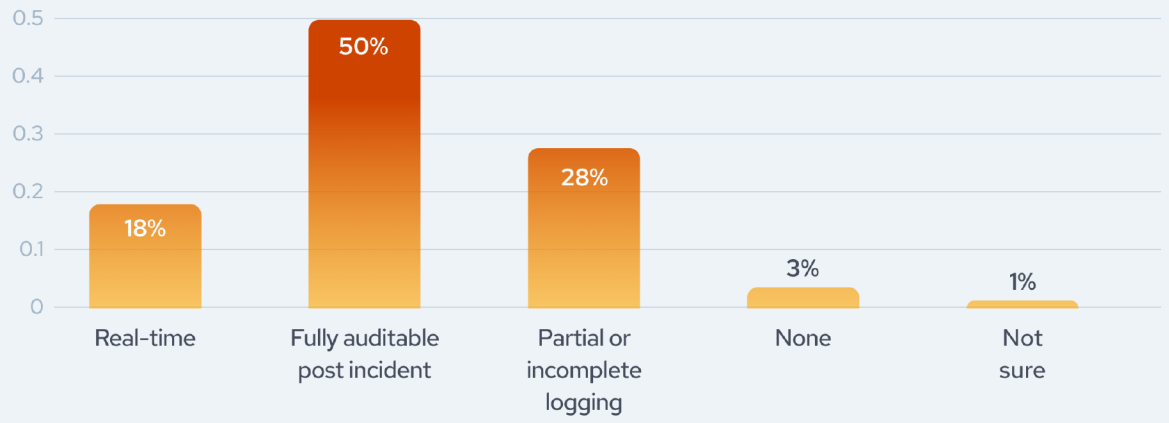

Visibility into how those components behave at runtime is primarily retrospective. Fifty percent of organizations describe their AI runtime visibility as "fully auditable post-incident." Another 28% report partial or incomplete logging, and only 18% have real-time visibility. There's a small portion (3%) who admit they have none in place. **For most organizations, current AI runtime oversight means being able to reconstruct what happened after an incident, not observe or intervene while it is occurring.** Given the speed at which AI components can execute decisions and propagate actions across connected systems, post-incident auditability closes the accountability loop without closing the exposure window.

only
18%
of organizations **have real-time visibility** into the runtime behavior of AI-powered application components

That context-and-mitigation problem from the prior finding is amplified in AI-driven environments. WAF enforcement was already constrained by insufficient application context; AI components that generate novel, non-signature-based behavior make that context harder to build and maintain. The enforcement trust gap that limits conventional runtime controls does not diminish with AI involvement — it widens, because the behavior these components produce is less predictable and less amenable to pattern-based classification.

Key Finding 4:

The Main Bottleneck for Protection is Proof of Exploitability

When security teams face a production investigation, the dominant friction is not a lack of people — it is a lack of signal. The hardest part of production security work, according to respondents, is distinguishing genuinely exploitable vulnerabilities from theoretical findings. What would most change remediation velocity is runtime evidence that confirms actual exploitability — whether a vulnerability is genuinely reachable in that specific environment. More staff working from incomplete context produces faster activity; it does not produce faster resolution.

54%

cite the inability to distinguish **real threats from non-exploitable findings** as their top investigation challenge

The investigation challenge is a signal quality problem. Asked to name their top challenge when investigating a suspected production security risk, 54% of respondents cite the ability to distinguish real threats from non-exploitable findings. Prioritizing vulnerabilities by actual risk level follows at 32%. Staffing or skill limitations appear at 4%. **These teams are not under-resourced; they are working from information that does not support confident, prioritized action.**

Figure 8: Top Challenges when Investigating Suspected AppSec Risks in Production

What would improve remediation velocity points to the same gap. Asked which resource would be most useful when addressing known application vulnerabilities, 41% want clear proof that a vulnerability can be exploited in production. The ability to mitigate or contain risk without requiring an immediate code change follows at 37%, and visibility into exact code paths and data flows at 33%. Additional staffing or resources: 4%. The pattern is consistent across both questions – runtime evidence and exploitability precision matter more than capacity.

Figure 9: Most Useful Resources for Remediating Known Application Vulnerabilities

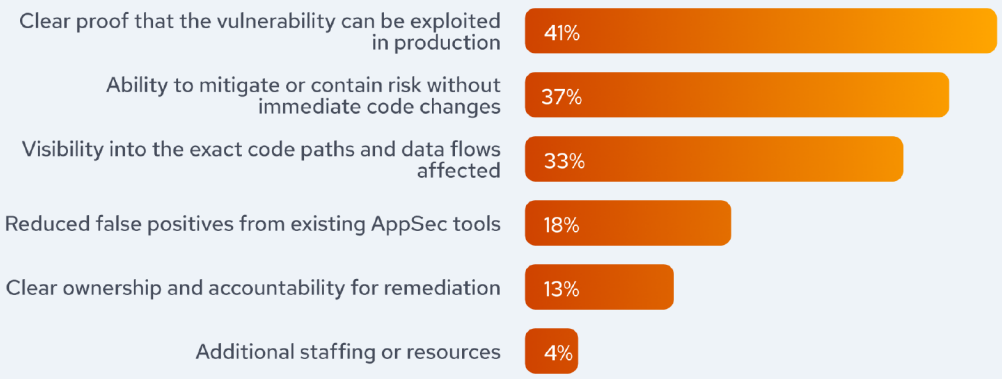

This is a meaningful finding for how security programs get resourced. **Organizations that have responded to persistent production incidents by adding headcount or expanding scanning coverage may be solving for the wrong constraint.** The bottleneck is the mechanism that converts a known vulnerability into a confirmed, actionable risk – runtime proof that this particular vulnerability, in this particular environment, is reachable. Without it, prioritization depends on severity ratings and theoretical exposure rather than evidence of exploitability, and remediation queues remain difficult to triage with confidence. As production environments incorporate more AI-driven workloads whose runtime behavior is harder to interpret, the demand for that validation signal becomes more acute.

Key Finding 5:

The Will to Block Exists But Trusted Mitigation is Missing

73%

of organizations would **adopt virtual patching** that could reliably block production exploits with minimal false positives

Organizations broadly want runtime blocking and mitigation capability – but the conditions under which they would use it reveal the constraint. When asked whether they would adopt virtual patching controls that could reliably block production exploits with minimal false positives, 73% say they are likely or very likely to do so. The conditional phrasing of that question is where the constraint surfaces: most organizations do not believe their current controls meet that standard, and the distance between enforcement appetite and enforcement reality is substantial.

Only 17% of organizations report their WAF or equivalent controls are configured to automatically block application-layer attacks. The majority operate in more conservative modes: 46% block well-understood patterns while alerting on complex ones, and 24% operate primarily in alert or monitoring mode. An additional 11% use logging only, citing false positive concerns as the explicit reason. Four in five organizations have the tooling but are not using it at full blocking capacity.

Figure 10: Likelihood of Adopting Virtual Patching with Minimal False Positives

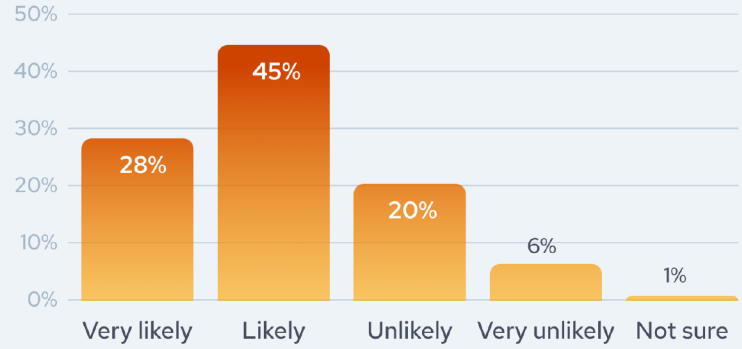

Figure 11: Current Use of WAF or Similar Mitigation Controls in Production

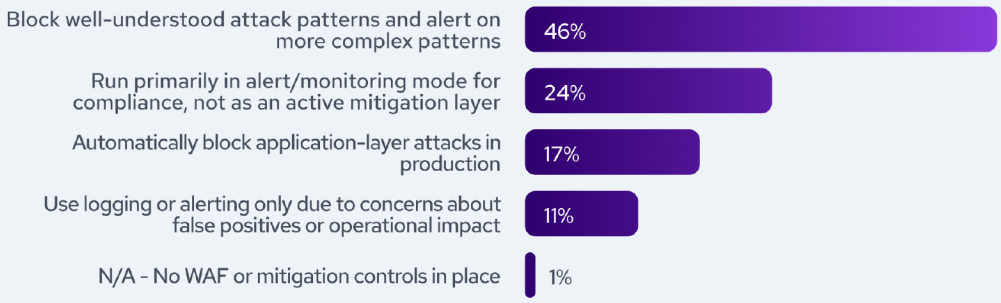

When asked about top challenges with WAFs as a mitigation layer, **56% of respondents cite lack of application-level context for safe blocking,** and 32% cite fear of disrupting business-critical functionality. High tuning and maintenance effort follows at 30%. These concerns share a root: without granular visibility into how an application behaves under normal conditions, security teams cannot confidently distinguish a malicious request from a legitimate one at the point of enforcement.

Figure 12: Top Challenges Using WAFs or Similar Controls as a Mitigation Layer

That visibility gap has compounding operational effects. When investigating unusual application behavior in production, 45% of security teams report struggling to explain it at least monthly — 11% say this happens weekly. Teams that cannot interpret what they are seeing are unlikely to trust automated blocking, and teams operating in alert-only mode accumulate unresolved signals rather than acting on them. Conservative WAF configuration and limited production visibility reinforce each other. **The issue is not intent, but capability: without the application context needed to confirm what is actually exploitable in production, organizations remain aware of risk but unable to reduce it in time.** That capability gap is increasingly reflected in where security leaders say they plan to invest next.

Figure 13: Frequency Your Security Team Struggles to Explain Unusual Application Behavior in Production

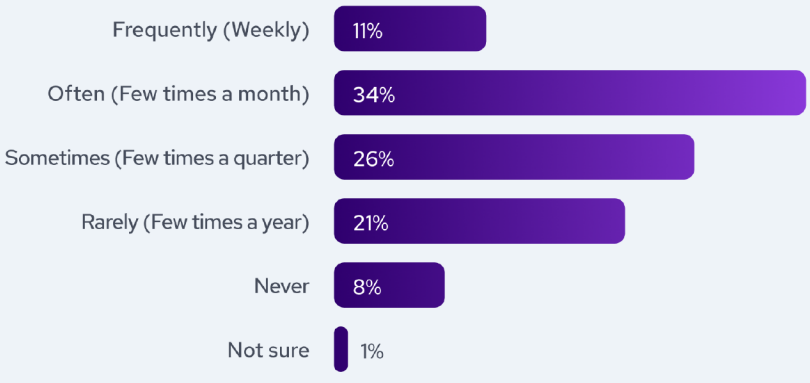

Key Finding 6:

Investment Intent Is Turning Toward Runtime Security

Figure 14: Planned Areas of Security Investment Over the Next 24 Months

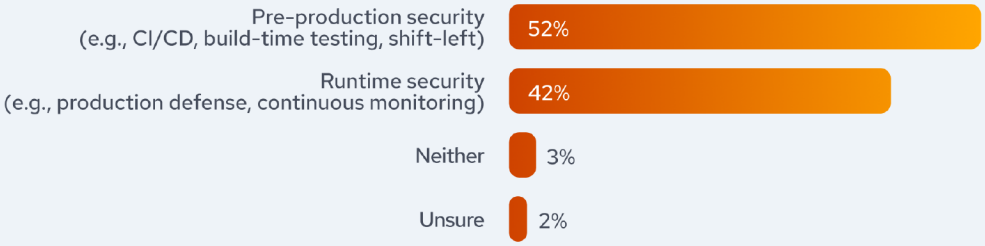

When asked where they plan to direct additional security investment over the next 24 months, 42% of organizations name runtime security — production defense and continuous monitoring — while 52% still point to pre-production. The gap reflects where security programs have historically concentrated, but the picture is evolving. **The survey was conducted before systems such as Mythos demonstrated that frontier AI could generate working exploits at machine speed, compressing the window between vulnerability identification and active attack.** Against that backdrop, the 42% runtime intent signal is notable — representing nearly as many organizations as plan to increase pre-production investment — though budget expectations for runtime remain mixed: 29% expect the category to grow over the next 12–24 months, while 35% expect a decrease. Investment patterns are beginning to reorient, but the data reflects a field still navigating where priorities need to go.

42%

of organizations plan to invest more in **runtime security** over the next 24 months

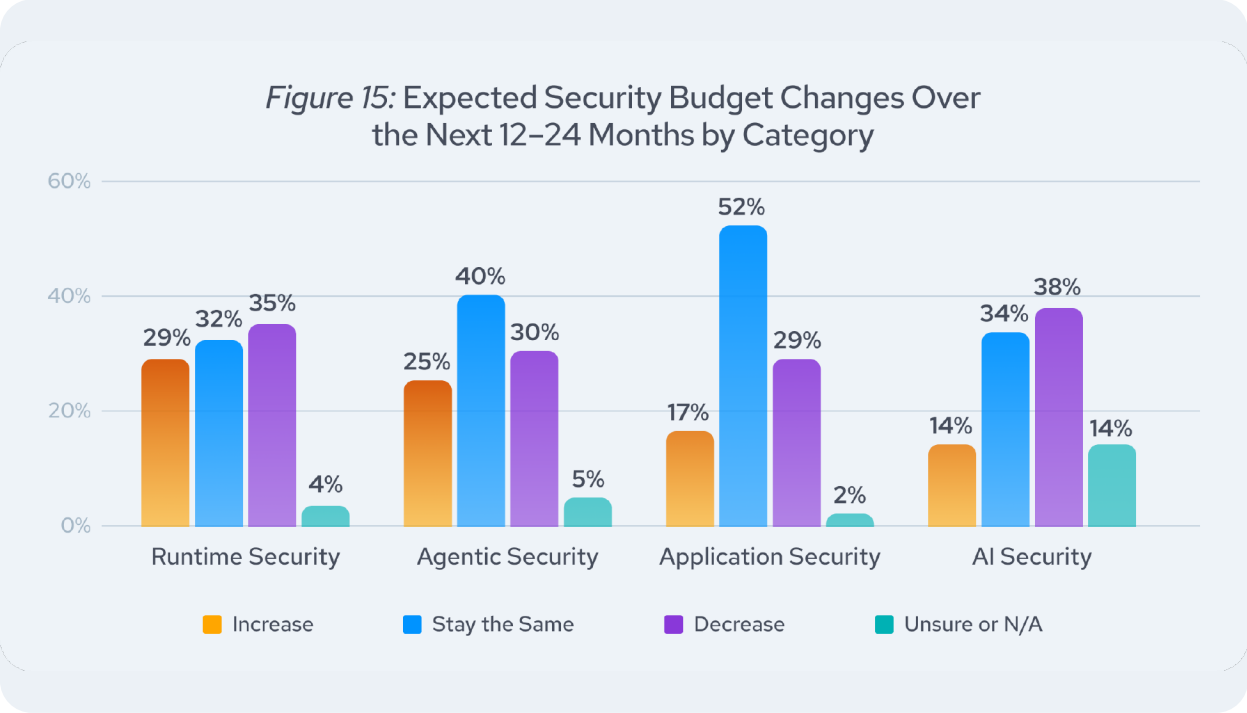

Investment signals do not indicate this gap is closing quickly. Thirty-eight percent of respondents expect their AI security budget to decrease over the next 12–24 months, and 34% expect it to stay flat. Fourteen percent expect an increase, and another 14% were unsure or indicated the question was not applicable. Current allocation levels suggest AI security is already a recognized priority — 69% allocate 6–20% of their security budget to AI and agentic security — yet future intent does not match current spend. The 14% unsure how their AI security budget will change may signal genuine uncertainty about which approaches will prove effective rather than a deliberate decision to reduce investment. For organizations carrying active production concerns about AI components without real-time visibility into their behavior, the gap between deployment and oversight represents a risk that existing governance models were not designed to manage.

Conclusion

Application security programs have matured significantly at the pre-production layer — the tooling, testing, and processes that identify vulnerabilities before release are broadly in place. Shift-left strategies have delivered detection maturity, but that maturity was built for a threat environment that no longer moves at the same speed. That gap is structural and consistent across the key findings: known vulnerabilities that remain unresolved long enough to be exploited, production environments absorbing incidents that shift-left controls were not designed to intercept, and runtime tools that operate in alert mode rather than active mitigation.

The missing capability is not additional coverage or headcount; it is runtime confidence, the ability to see and act on production behavior precisely enough to prioritize and mitigate.

As AI-driven components enter production and make behavioral interpretation harder still, and as frontier AI compresses the time between disclosure and exploitation, the question for organizations whose investment has concentrated upstream is no longer whether risk reaches production — it is whether they can act on it quickly enough once it does.

Final Thoughts from Miggo Security

Runtime Mitigation Stops What Patching Can't

At Miggo, we believe runtime mitigation is the missing layer for AppSec teams in the era of Mythos and AI enabled vulnerability exploitation. As the findings in this report make clear, the challenge is no longer detecting vulnerabilities; it is closing the window between identification and remediation before attackers weaponize them. While vulnerabilities reach production at scale and patching takes days or weeks, Miggo closes that Patch Gap in minutes with precision mitigation engineered for the exact exploit path, and without requiring engineers to write code.

Powered by patented DeepTracing™, Miggo reverse-engineers each exploit primitive and maps it to live runtime, then generates, validates, and deploys a targeted mitigation for any exploit without interrupting a single engineering sprint. Security teams using Miggo cut vulnerability backlog by over 95% and mitigate over 90% of exploitable risk in under an hour.

MIGGO

Miggo delivers trusted runtime mitigation in minutes – not days.

See how Miggo closes the Patch Gap in your environment with precision mitigation engineered for the exact exploit path.

Request a demo today

Full Results

Overview Application Security Visibility & Tooling

Figure 16: Which application security tools or controls does your organization currently use (Select all that apply)

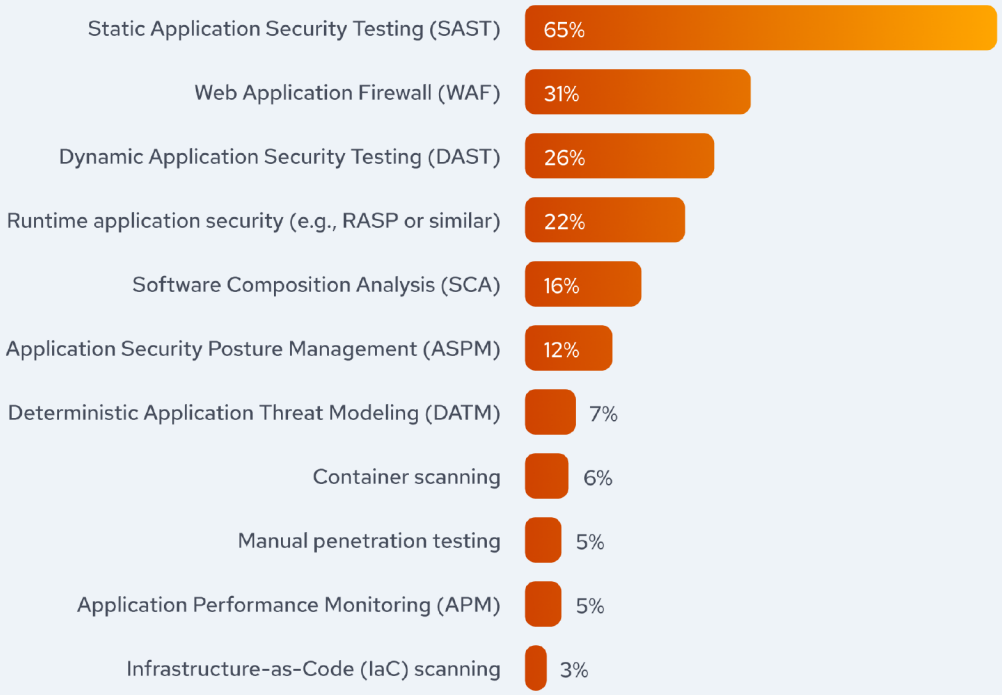

Figure 17: Rate the level of visibility your security team currently has into the following aspects of applications.

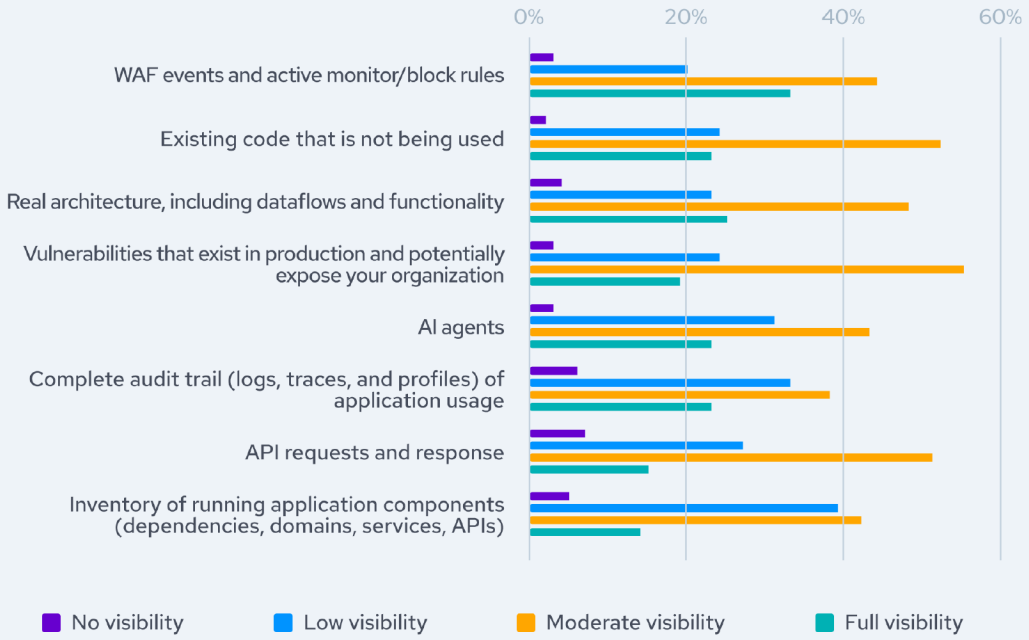

Figure 18: In the past 12 months, has your organization had a production application security incident that bypassed pre-production controls?

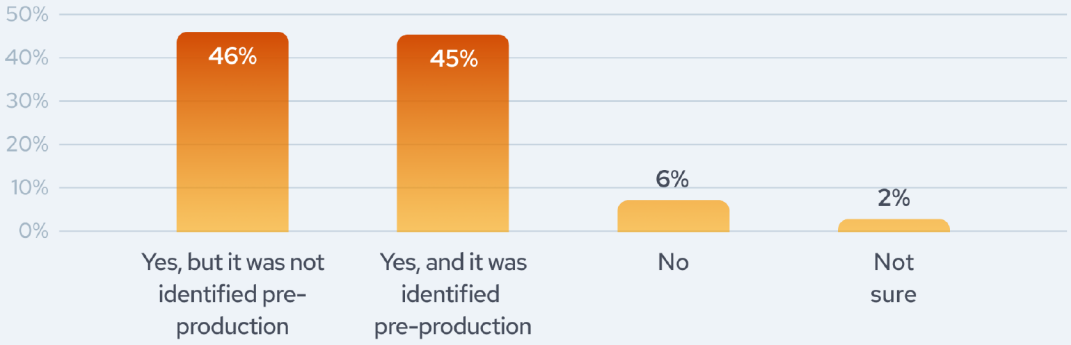

Figure 19: What are the top challenges your organization faces when investigating a suspected application security risk in production? (Select up to 2)

Runtime Behavior and Emerging Risk

Figure 20: How often does your security team struggle to explain unusual application behavior in production?

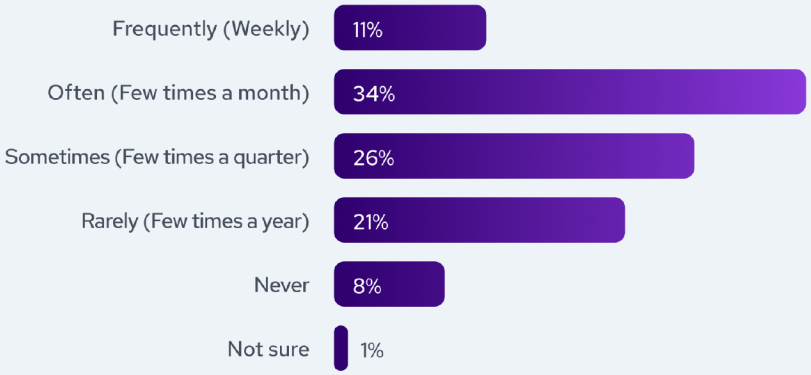

Figure 21: Which best describes your organization's current use of AI-powered or autonomous application components?

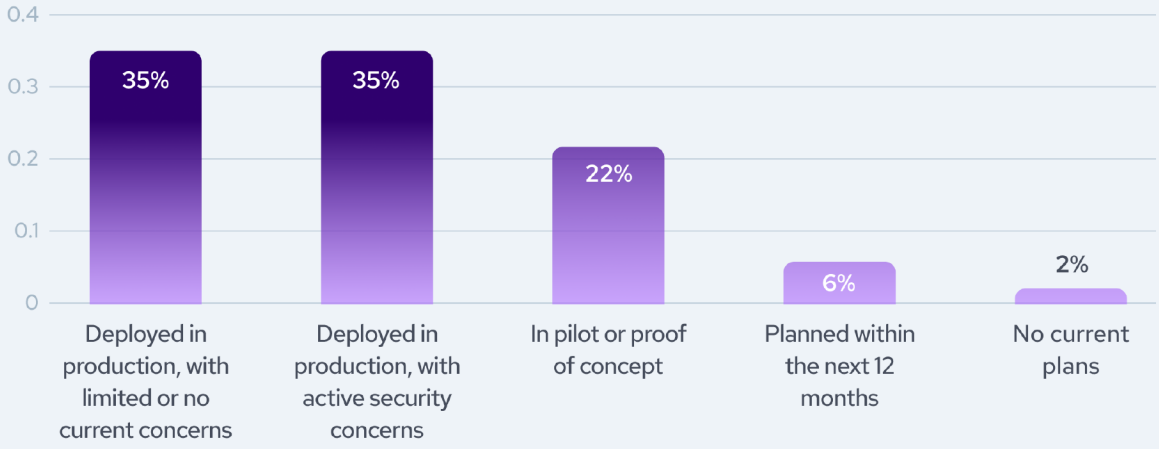

Figure 22: How effective are your current application security tools at detecting and assessing risk caused by AI applications and/or agents?

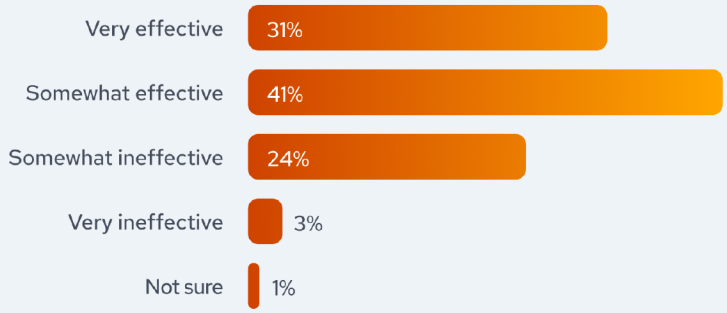

Patching

Figure 23: In the past 12 months, has your organization experienced an application security incident that involved a vulnerability already known to your security team?

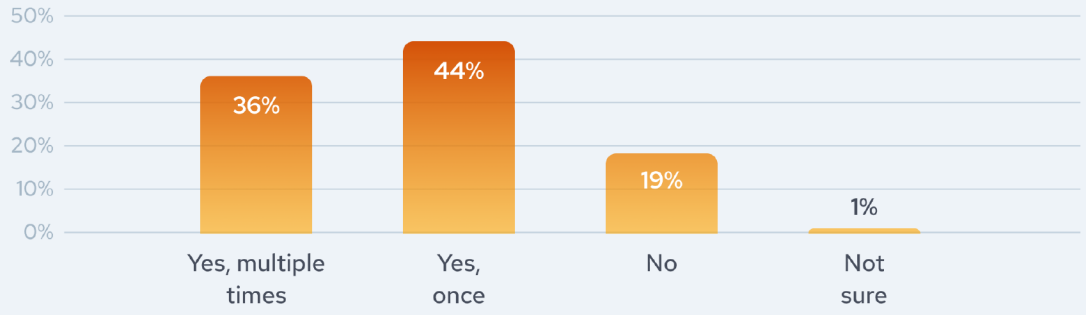

Figure 24: When a known application vulnerability is not remediated promptly in production, which factors most commonly contribute to the delay? (Select up to three)

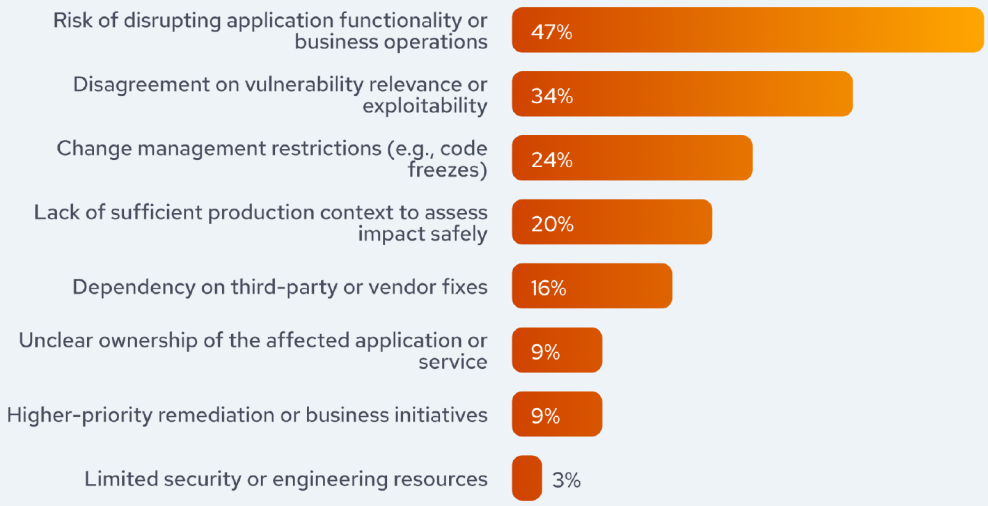

Figure 25: On average, how long does it take for an identified critical/high application vulnerability to be remediated in production?

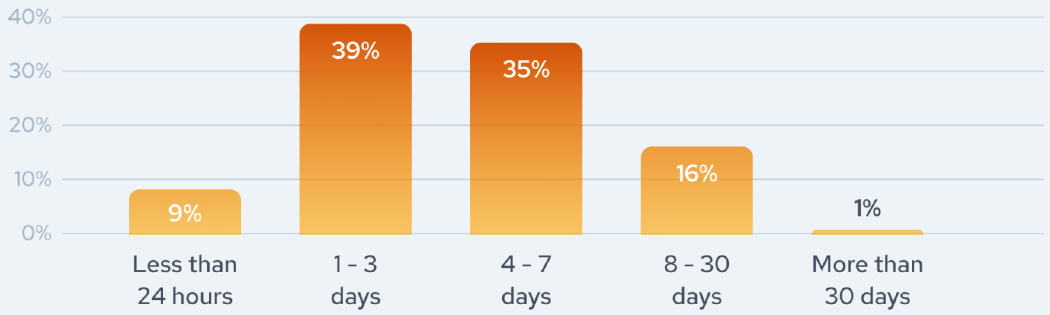

Figure 26: Which of the following would be most helpful when remediating known application vulnerabilities? (Select up to two)

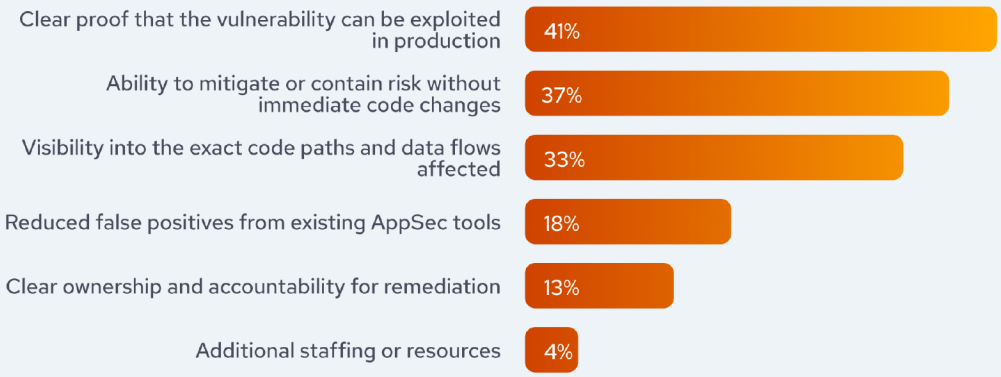

Figure 27: How often do application security findings lead to disagreement or pushback from engineering teams due to a lack of evidence or context?

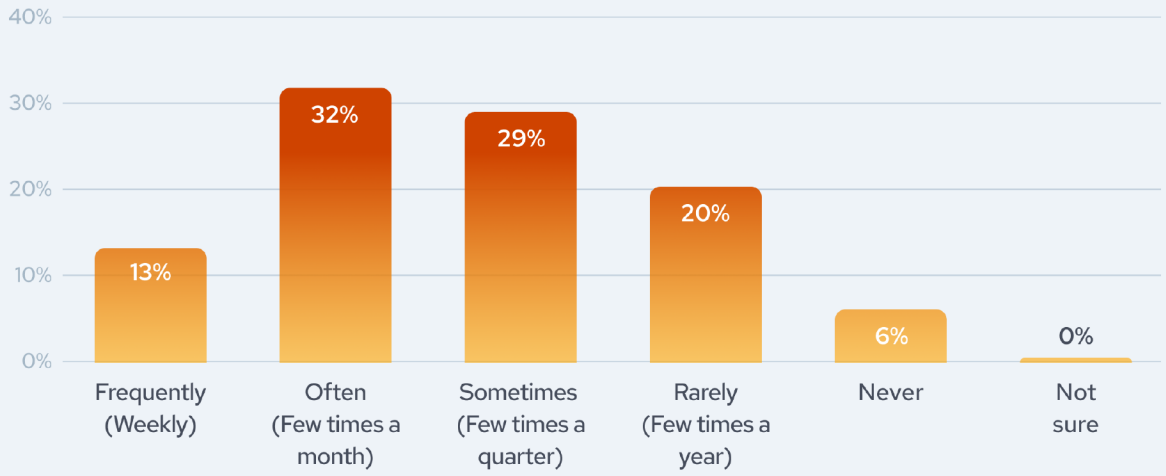

Mitigation and Virtual Patching in Production

Figure 28: Which statement best reflects your organization's current use of a Web Application Firewall (WAF) or similar mitigation controls in production?

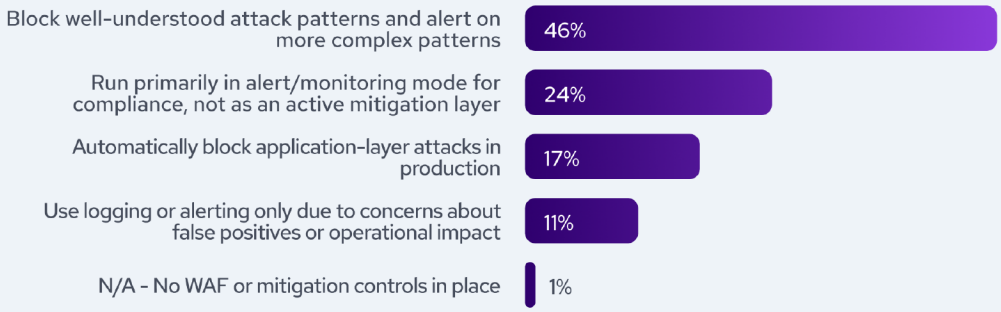

Figure 29: What are your top challenges when using WAFs (or similar controls) as a primary mitigation layer? (Select up to two)

Figure 30: If mitigation controls could reliably block production exploits with minimal false positives, how likely would your organization be to use them for virtual patching?

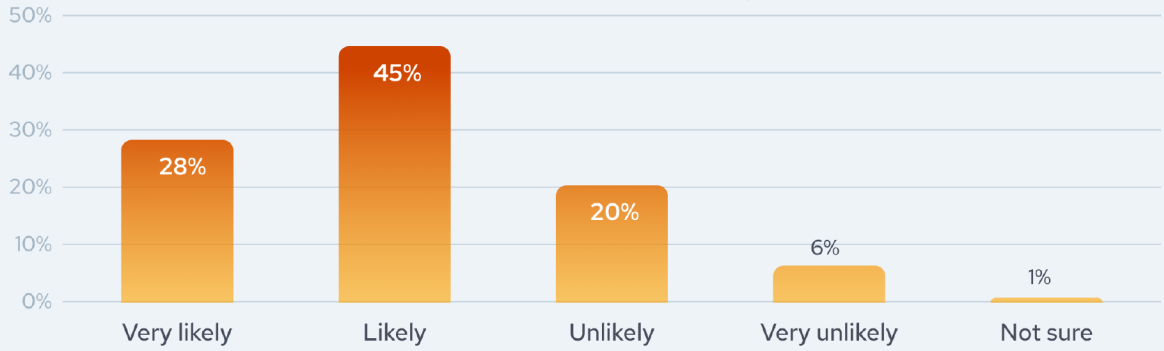

Future Investment and Priorities

Figure 31: Which stage of the application lifecycle does your organization plan to invest more in over the next 24 months?

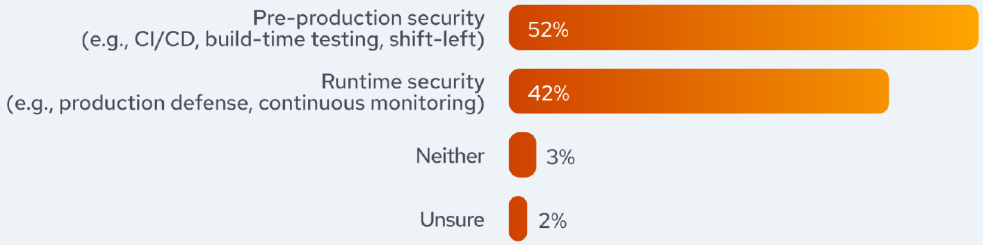

Figure 32: Which capability would be most important to your organization when evaluating new application security investments? (Select up to 2)

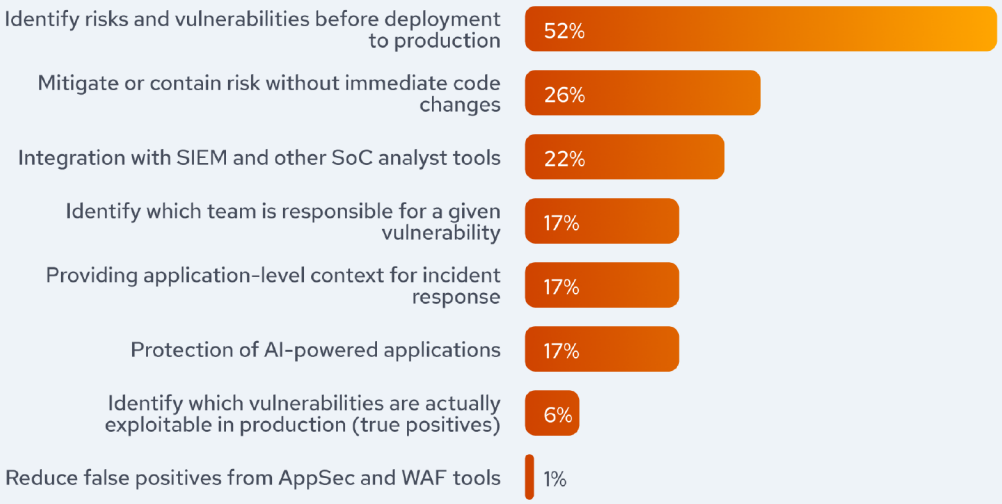

Figure 33: As attack surfaces increase from dynamic and AI-driven workloads, how confident are you that your organization's current application security strategy will be reliable?

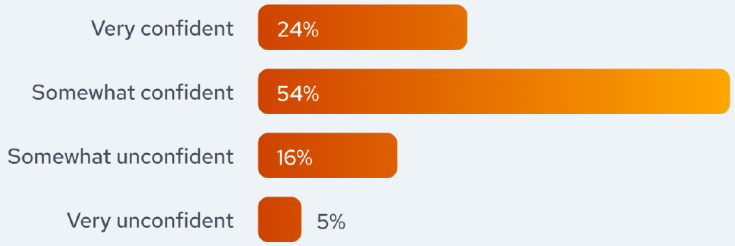

Figure 34: Which of the following best describes your organization's visibility into the runtime behavior of AI-powered application components?

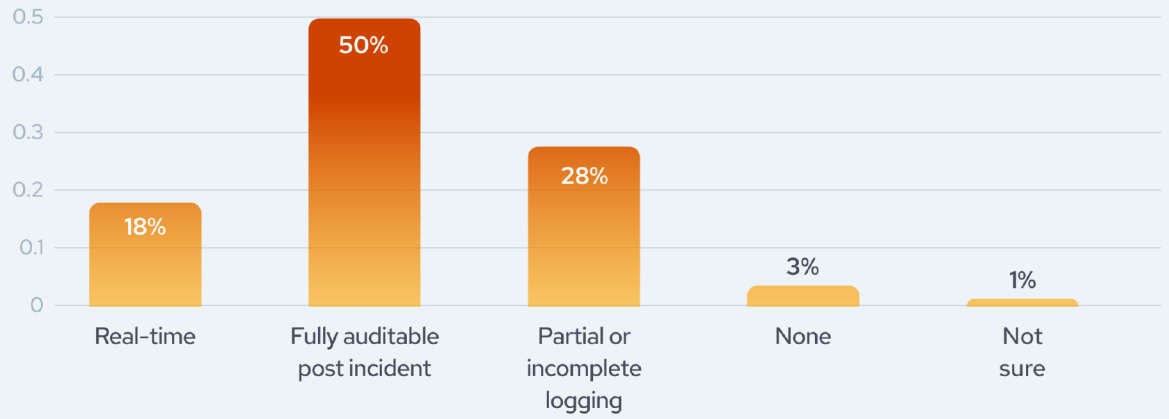

Figure 35: What percentage of your 2026 security budget is allocated to application or runtime security?

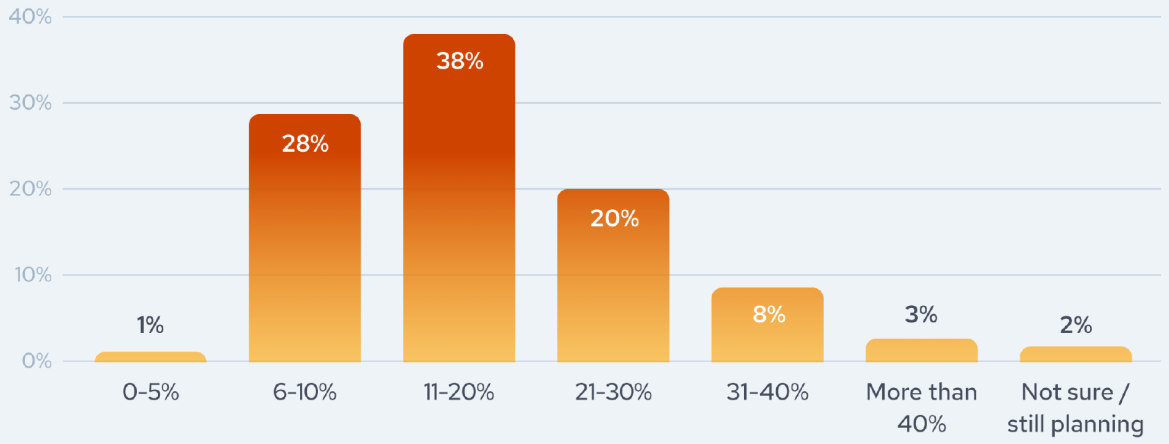

Figure 36: What percentage of your 2026 security budget is allocated to AI and agentic security?

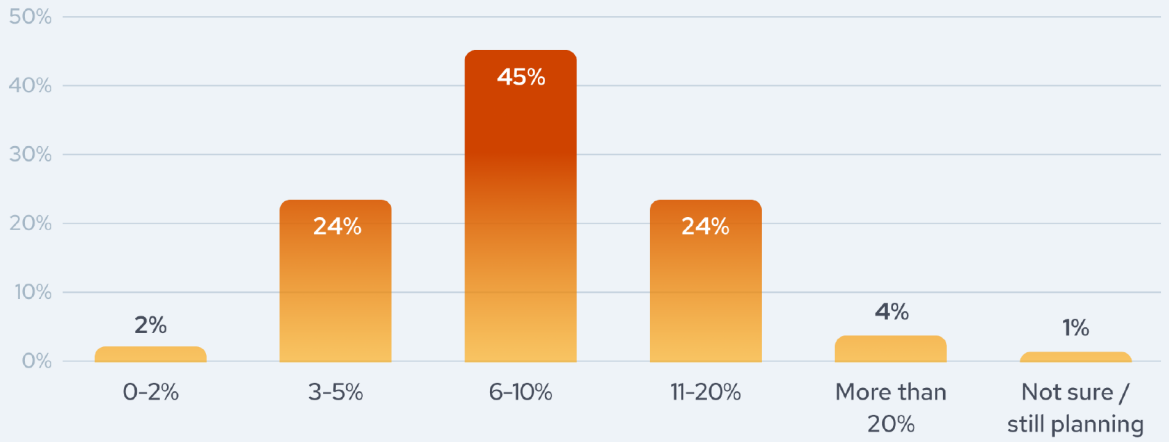

Figure 37: What changes do you expect in the following areas of your security budget over the next 12-24 months? (Please select one option for each category)

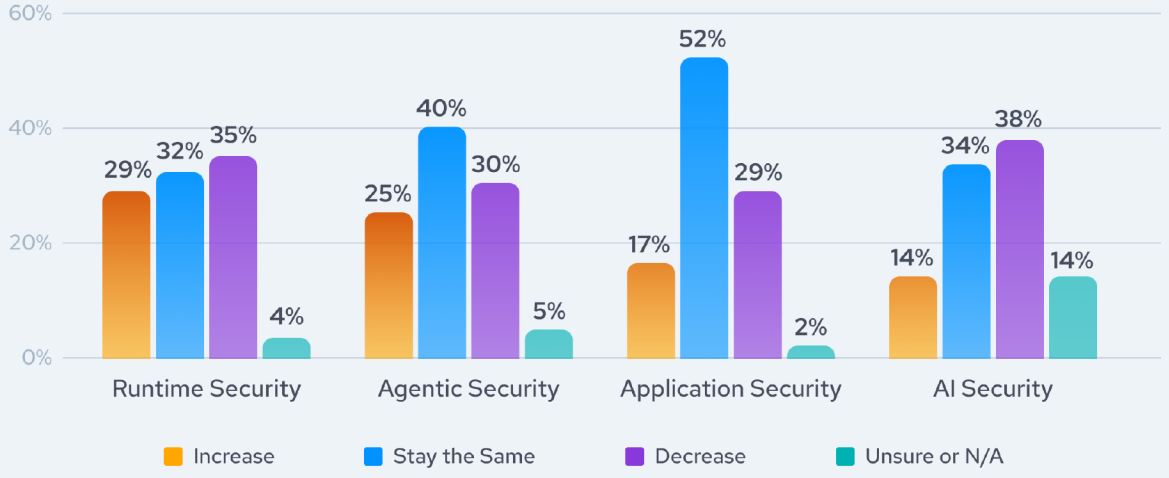

Demographics

Figure 38: What region of the world are you located in?

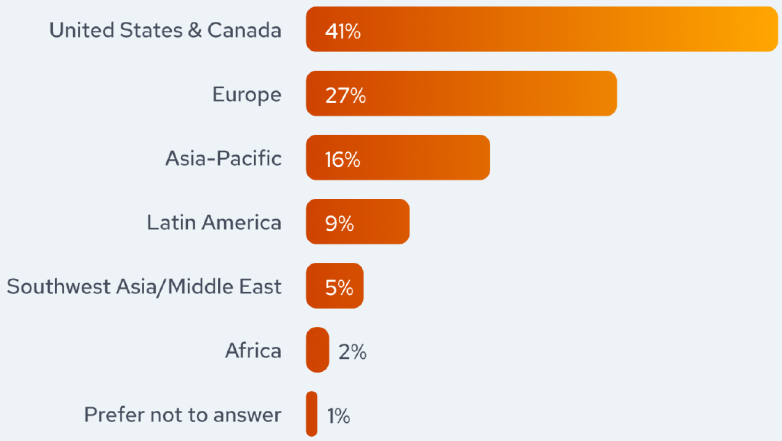

Figure 39: What is the size of your organization?

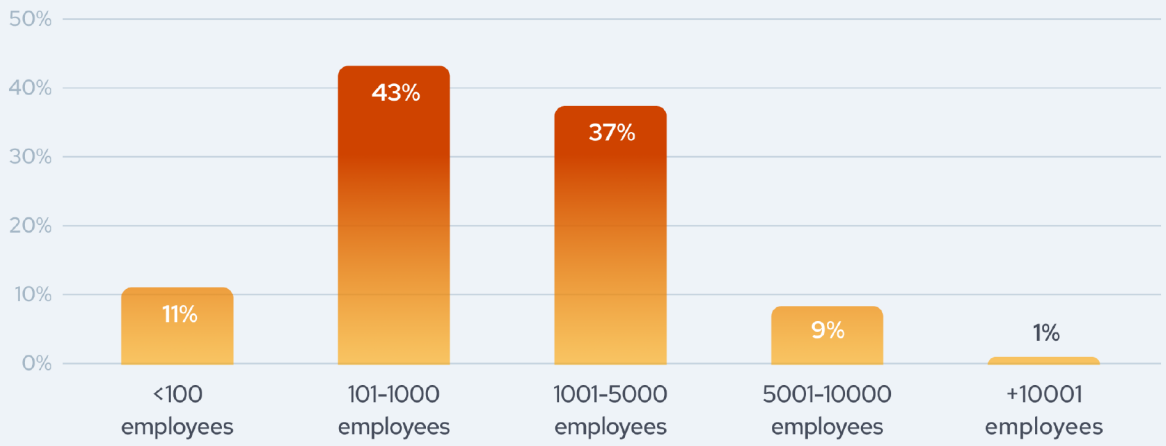

Figure 40: What best describes your current role?

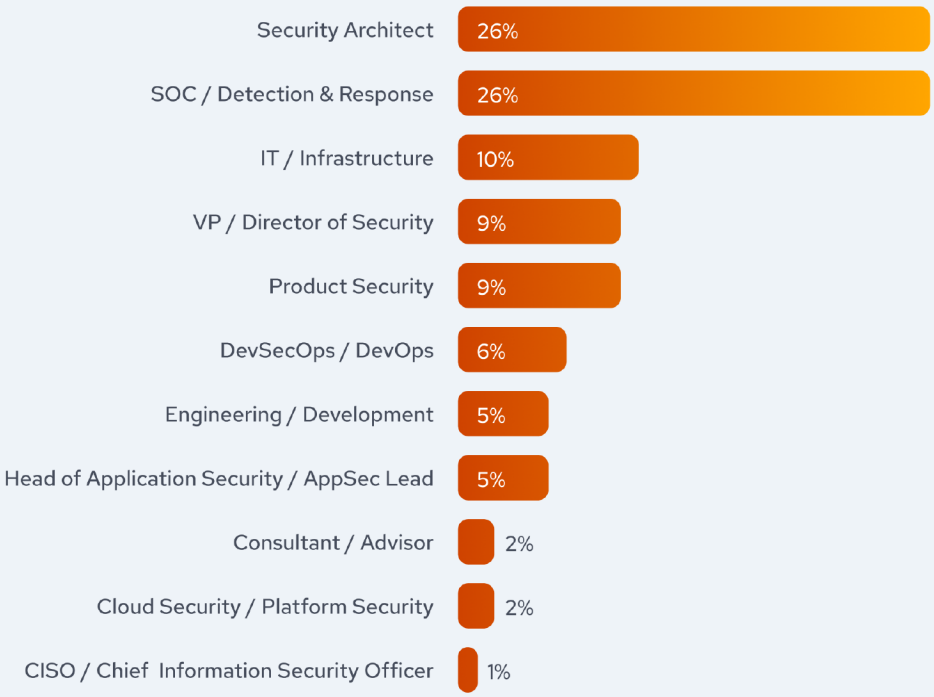

Survey Methodology

The Cloud Security Alliance (CSA) is a not-for-profit organization with a mission to widely promote best practices and ensure cybersecurity in cloud computing and IT technologies. CSA also educates various stakeholders within these industries about security concerns in all other forms of computing. CSA's membership is a broad coalition of industry practitioners, corporations, and professional associations. One of CSA's primary goals is to conduct surveys that assess information security trends. These surveys provide information on organizations' current maturity, opinions, interests, and intentions regarding information security and technology.

Miggo commissioned CSA to develop a survey and report to better understand the industry's knowledge, attitudes, and opinions regarding application security. Miggo financed the project and co-developed the questionnaire with CSA research analysts. The survey was conducted online by CSA in January 2026, and it received 902 responses from IT and security professionals from organizations of various sizes and locations. CSA's research analysts performed the data analysis and interpretation for this report.

Goals of the Study

This research examines how organizations manage application security risk in environments where vulnerabilities are continuously identified but not always immediately remediated, and where threats increasingly materialize in production. The goal is to establish a clear view of how detection, prioritization, and runtime controls interact to influence real-world security outcomes.

Specifically, the survey aims to:

- Measure how organizations detect, track, and remediate application vulnerabilities
- Evaluate the effectiveness of pre-production and runtime security controls in preventing incidents and reducing exposure
- Assess how organizations approach exploit mitigation
- Examine challenges in investigation and remediation, particularly around prioritization, exploitability validation, and operational decision-making
- Explore how emerging technologies, including AI-powered application components, are affecting visibility, oversight, and security investment priorities