Red Hat

2026

# The state of **cloud-native security**

Protecting full cloud-native ecosystems

# Contents


## Introduction

Page 3

Chapter 1

## Security incidents and their cost

Page 8

Chapter 2

## Governance and maturity in cloud-native security

Page 10

Chapter 3

## Emerging investment trends: Focus on automation and supply chain

Page 14

Chapter 4

## The emerging risk frontier: AI and cloud security

Page 19

## Conclusion

Page 23

# Introduction

The **2026 State of Cloud Native Security report** builds on previous editions, expanding its focus beyond Kubernetes to reflect the broader enterprise security landscape. The research explores how organizations put a security focus into code, infrastructure, and workloads across hybrid and multicloud environments, with added emphasis on governance, automation, and the impact of AI.

This year's report draws from 600 completed surveys, each taking approximately 20 minutes, conducted online between August 25 and September 23, 2025. Respondents included IT professionals responsible for applications, security, platforms, and development at companies with 100 or more employees, sourced through expert networks and online panels.

## Key findings

Incidents affect nearly all organizations regardless of size or region

97%

experienced at least one issue.

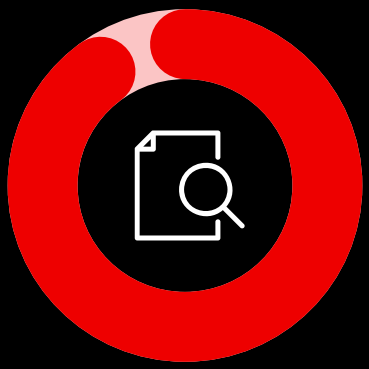

## Incidents reported

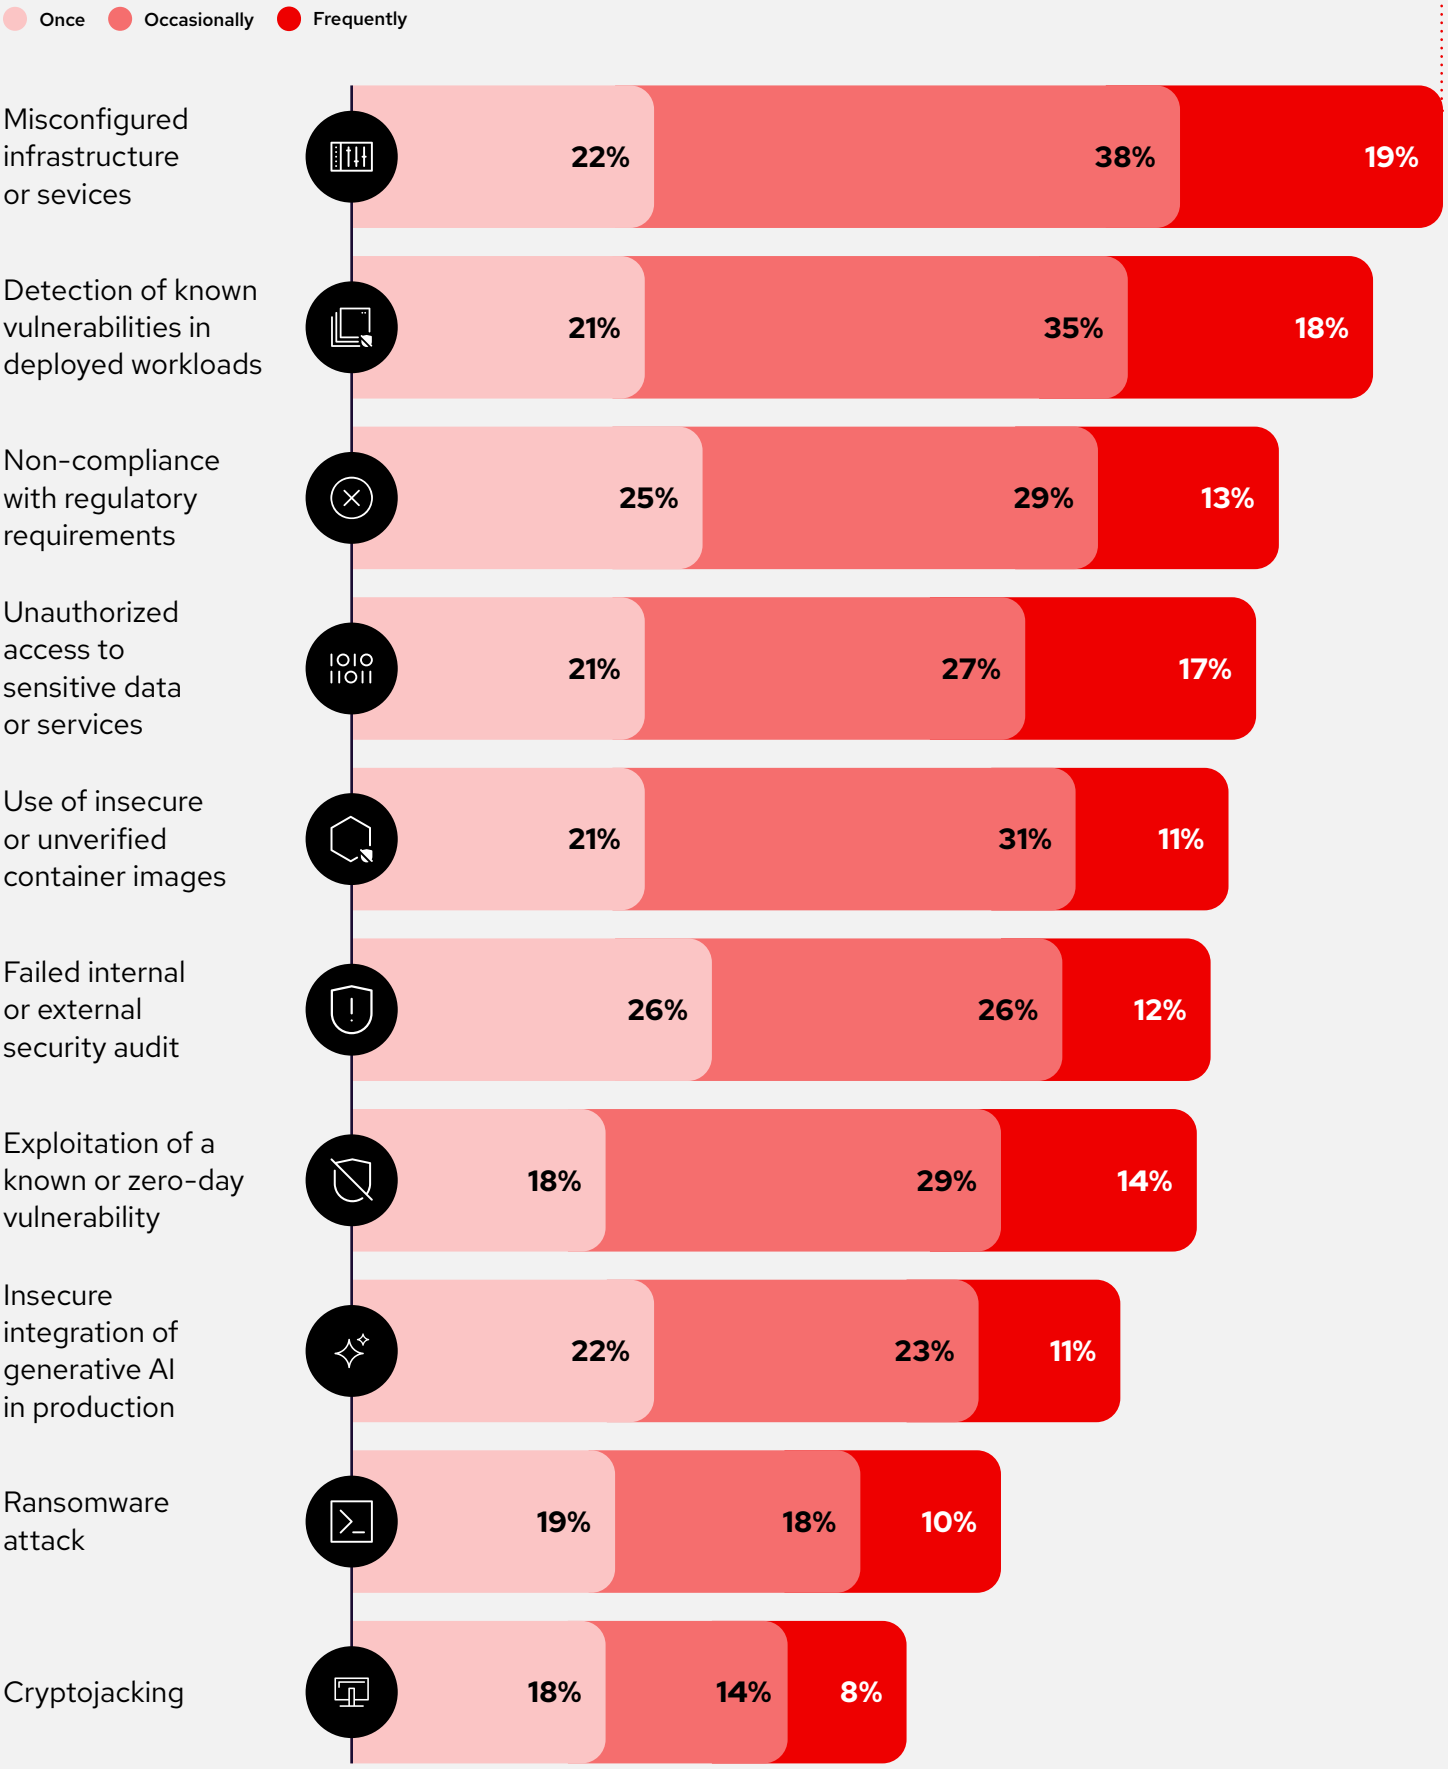


***Figure 1.*** *Misconfigurations and vulnerabilities lead incident types.*

## Security incidents are common and costly

Security issues plague almost all cloud-native teams. **97%** of organizations experienced at least 1 cloud-native security incident in the past year. These incidents carry a tangible business cost, as **74%** of organizations have slowed or delayed application deployments in the last 12 months due to security concerns.

**In short, delays, firefighting, and disruption from security problems are the norm rather than the exception, underscoring the high cost of inadequate security measures.**

**97%**

of organizations experienced at least 1 cloud-native security incident in the past year.

**74%**

of organizations have slowed or delayed application deployments in the last 12 months due to security concerns.

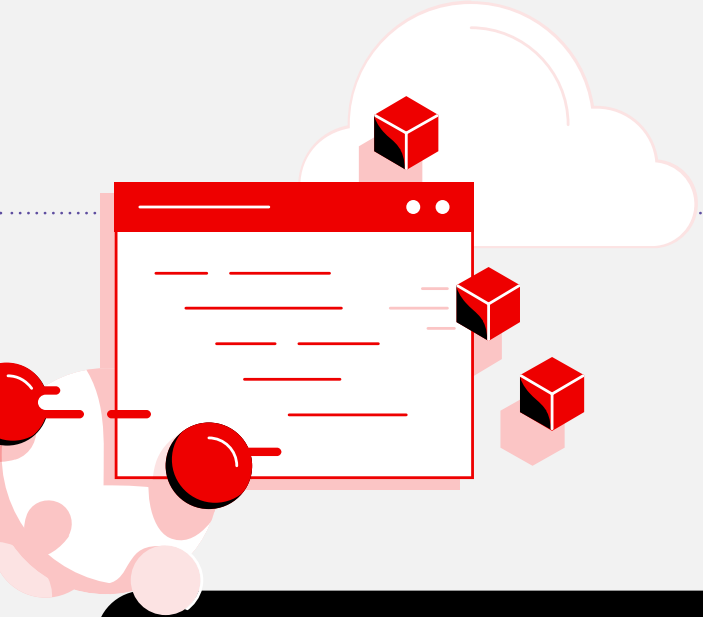

**39%**

of companies report having a well-defined cloud-native security strategy, with over half still developing or evolving a plan.

## Cloud-native security is foundational but uneven

While cloud-native security is widely recognized as critical, maturity levels vary greatly across organizations. Only **39%** of companies report having a well-defined cloud-native security strategy, with over half still eveloping or evolving a plan. At the same time, a majority **(56%)** describe their day-to-day security posture as highly proactive.

**This suggests a confidence that often outpaces actual strategy and execution.**

The gap highlights the need for more structured approaches to cloud security governance and maturity.

## Guardrails define maturity, but adoption is inconsistent

The use of security guardrails (built-in security controls and best practices) is a key indicator of maturity, but implementation remains patchy. For example, basic identity controls are almost universal, and about ¾ of organizations use identity and access management (IAM) tools. However, only roughly half have adopted container image signing and verification for software integrity.

This includes image signing, runtime protection, automated policy enforcement and other measures, leading to an uneven security baseline across the industry. As a security respondent warned: "A major misconception is that cloud-native security is a set-it-and-forget-it solution, ignoring the need for continuous monitoring and adaptation."

**In other words, many teams still overlook critical safeguards.**

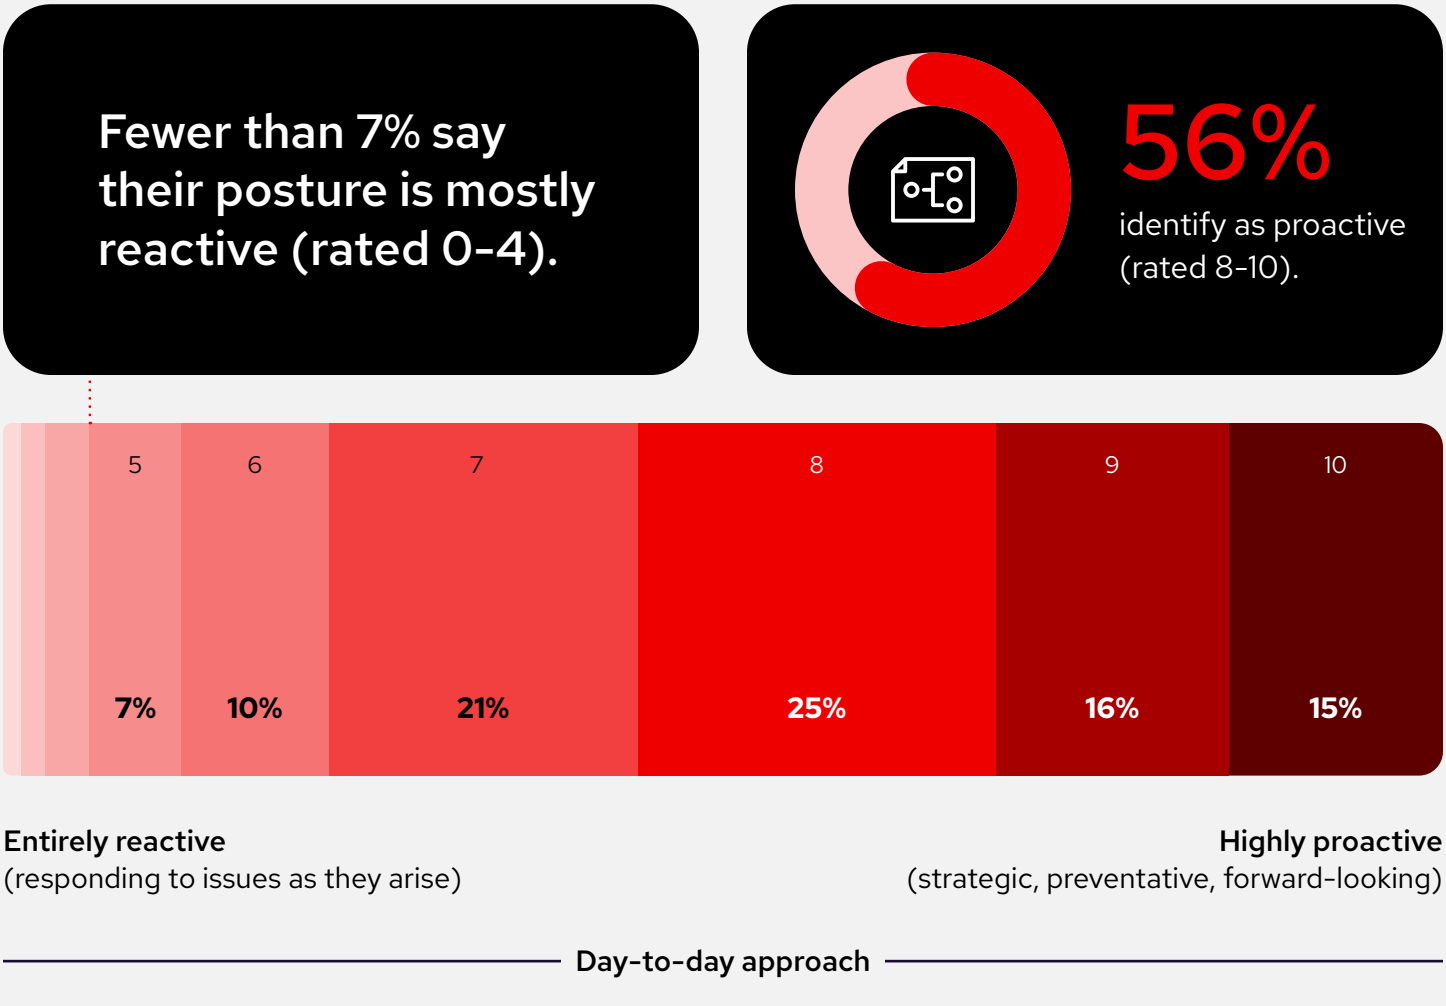


***Figure 2.*** *Confidence outpaces capability in cloud-native security focus.*

## Investments shift toward automation and supply chain security

Looking forward, organizations are rebalancing their security investments to address these maturity gaps. The top priorities for the next 1-2 years are DevSecOps automation and software supply chain security. Over **60%** of surveyed organizations plan to invest in automating security into continuous integration/continuous delivery (CI/CD) pipelines (policy automation, integration, etc.), and **56%** plan to invest in securing the software supply chain (managing integrity from code to runtime). Close behind is an emphasis on expanding runtime protection (**54%** plan to invest) to embed continuous defenses at deployment. This marks a consolidation of efforts around automation and built-in security, aligning investments with the areas that define mature, resilient cloud-native programs.

Over

**60%**

of surveyed organizations plan to invest in automating security into CI/CD pipelines.

**56%**

plan to invest in securing the software supply chain (managing integrity from code to runtime).

**79%**

of respondents agree that gen AI is creating new security challenges in their cloud environments.

**59%**

of organizations lack any documented internal AI usage policies or governance frameworks.

## Governance struggles to keep up with AI risks

Rapid adoption of AI in development and DevOps is introducing new risks faster than governance can respond. **79%** of respondents agree that gen AI is creating new security challenges in their cloud environments. Yet formal policies are lagging, as **59%** of organizations lack any documented internal AI usage policies or governance frameworks.

**This disparity suggests that AI-related risks (from data exposure to insecure AI tools) are growing without corresponding oversight, leaving organizations exposed.**

Chapter 1

# Security incidents and their cost

Security incidents remain a reality in cloud-native environments, and data shows they are relatively common and can disrupt operations.

In the past year, virtually every organization in our survey (**97%**) experienced at least 1 cloud-native security issue. Far from being rare anomalies, incidents such as misconfigurations and known vulnerabilities have become almost routine. In fact, misconfiguration of cloud infrastructure and detection of known vulnerabilities lead the list of incident types. This means that everyday lapses (for example, leaving an S3 bucket open or deploying an unpatched container) are causing more trouble than sophisticated attacks.

## The business impact of these incidents is significant.

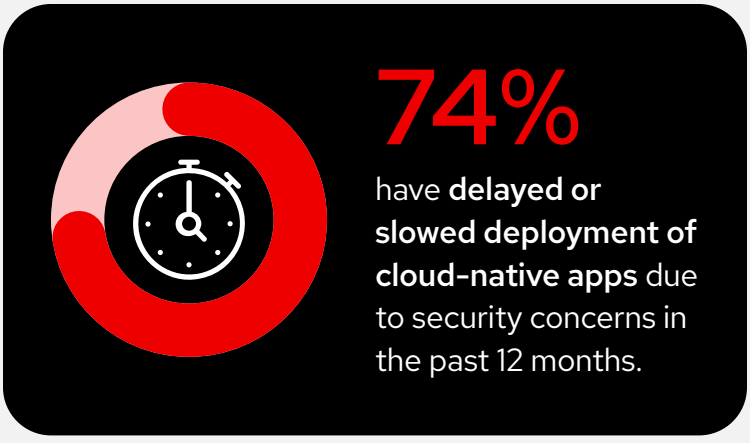


### Impact of cloud-native security incidents


***Figure 3.*** *Security slowdowns are common and costly.*

A striking **74%** of organizations surveyed report that they slowed down or delayed application releases in the last 12 months due to security concerns. In other words, 3 out of 4 teams had to slow deployments because a security issue arose, a direct hit to agility and time-to-market. This kind of delay is not just an inconvenience; it translates into lost revenue opportunity, missed deadlines, and frustrated teams.

Other common consequences of cloud-native security incidents include increased unplanned work and damage to customer trust. According to the survey, **92%** of organizations experienced at least 1 significant impact on their ability to deliver software or meet business goals due to security incidents. The types of impacts reported were numerous, for example:

## Delayed application releases.

Teams must postpone deployments or feature launches while fixing security issues.

## Reduced developer productivity.

Developers lose cycles remediating vulnerabilities or addressing configuration errors.

## Missed internal or customer deadlines.

Security problems cause slip-ups in delivery commitments.

## Loss of stakeholder or customer trust.

High-profile security hiccups erode confidence among leadership and clients.

## Reputational damage or even lost business.

In the worst cases, security failures lead to public scrutiny or customers leaving.

It's clear that weak cloud security directly costs organizations time and money. Delays drain engineering productivity, and there is a ripple effect on the business resulting in missed market opportunities and potential revenue loss. The prevalence of these setbacks, with ¾ of teams slowing down deployments, shows that cloud security is not just a technical concern but a serious business risk.

## Addressing these incident costs will require more proactive measures.

The high rate of misconfigurations (**78%** reported them in the past 12 months) suggests basics such as configuration management and vulnerability patching need improvement. Every organization should assume that without stronger guardrails and processes, they will continue to face frequent incidents and associated delays. The data makes a compelling case for investing in preventative security to avoid the much greater cost of reacting to repeated issues.

Chapter 2

# Governance and maturity in cloud-native security

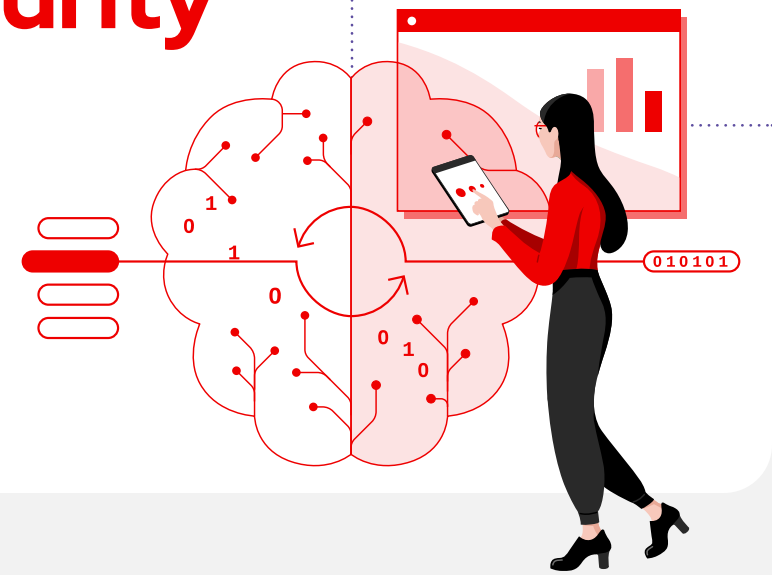


Achieving a strong cloud-native security posture is as much about governance and process maturity as it is about tools. Here, the research reveals a paradox: many organizations believe they are doing well, yet relatively few have actually put in place the formal strategies and controls that define a mature security program.

## Proactive stance

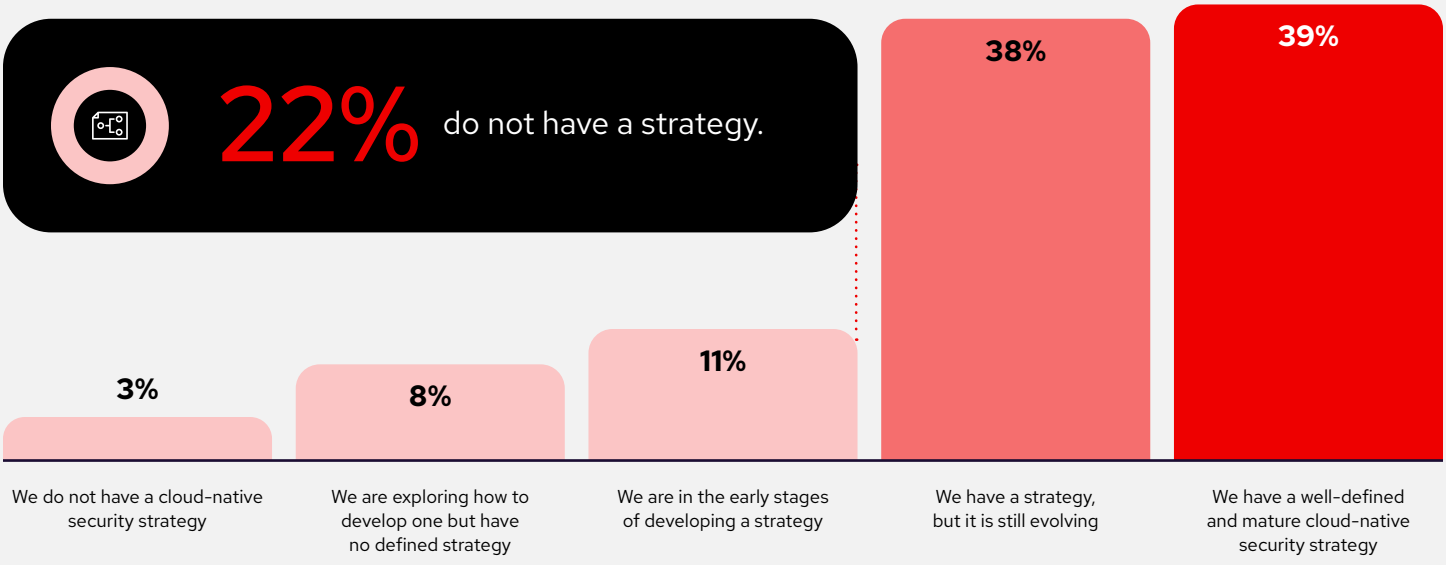


***Figure 4.*** *Without a defined strategy, security remains reactive and fragmented, leaving teams exposed.*

On 1 hand, a majority of teams surveyed claim a proactive stance, as **56%** rate their day-to-day security approach as "highly proactive." Furthermore, fewer than 7% self-identify as mostly reactive, which indicates most organizations aspire to be forward-looking.

## Teams clearly want to be secure. On the other hand, however, far fewer have the foundational governance to back that up.

Only **39%** of companies have a well-defined cloud-native security strategy in place. The rest are improvising, as more than half are still developing, refining, or even just exploring how to create a security strategy. In some cases (about 22% of organizations), there is no cloud security strategy at all yet, which is an obvious maturity gap.

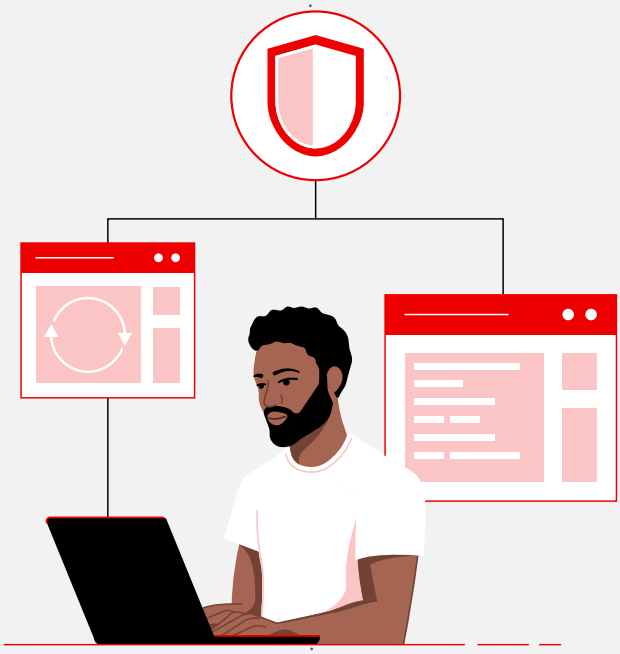

## In practice, this means that many organizations may be overestimating their readiness.

Declaring a proactive posture does not equal readiness if the organizations lack the policies and structure to enforce it. True cloud-native security maturity entails defined objectives, cross-team alignment, and embedded controls. This is where many programs fall short.

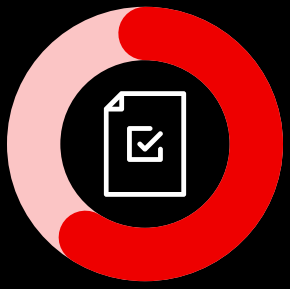

**56%**

say their day-to-day **posture is highly proactive.**

**39%**

have a **well-defined and mature** cloud-native security strategy.

# Security guardrails

A key indicator of maturity is the adoption of security guardrails: the built-in controls and practices (from access management to continuous monitoring) that are designed to keep cloud environments safe. The survey shows guardrail adoption is very uneven across organizations.

Certain basic measures are broadly implemented. For example, around 3 out of 4 of respondents have IAM solutions in place, reflecting that most understand the need for strong identity and authorization controls. However, more advanced or emerging best practices see much lower uptake.

Only about **half** of these organizations have implemented container image signing and verification to manage code integrity, and similarly, many have not yet deployed things like runtime protection or automated policy enforcement.

## In other words, the breadth and consistency of guardrails isn't where it should be.

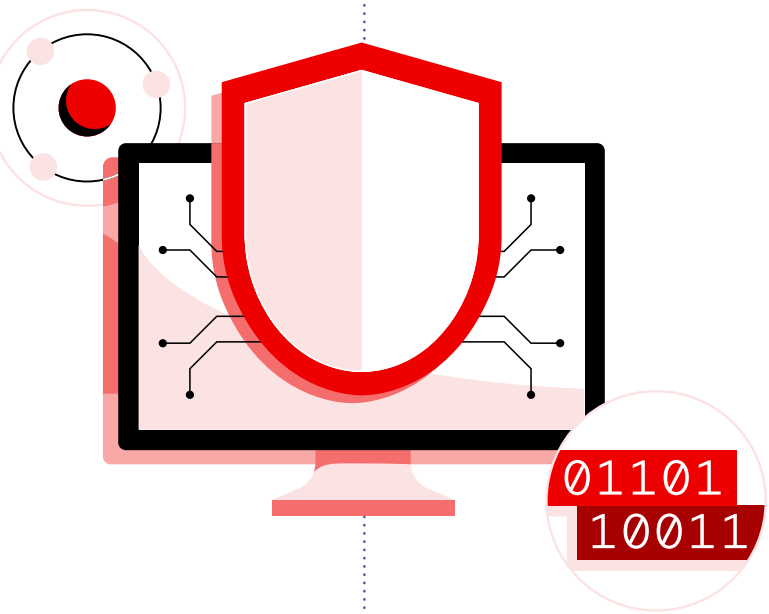


**61%**

of mature organizations are very confident in securing their software supply chain, versus far lower confidence among less mature peers.

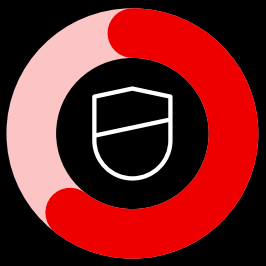

A lot of teams plug certain gaps while leaving others wide open. Without comprehensive, intentional governance, teams can be lulled into a false sense of security by default settings or ad-hoc efforts.

Notably, the research found that organizations with a well-defined security strategy consistently demonstrate higher adoption of such guardrails and greater confidence in their security. Mature programs treat security as part of the platform and pipeline, not an afterthought.

For example, teams that have a clear strategy are far more likely to be using controls like software supply-chain security tooling and automated policy enforcement, compared to those still developing a strategy.

They also report substantially higher confidence in areas like supply chain protection, as **61%** of mature organizations are very confident in securing their software supply chain, versus far lower confidence among less mature peers.

In short, maturity yields tangible security advantages: more consistent guardrails, better visibility, and a stronger security posture overall.

# Regulatory alignment

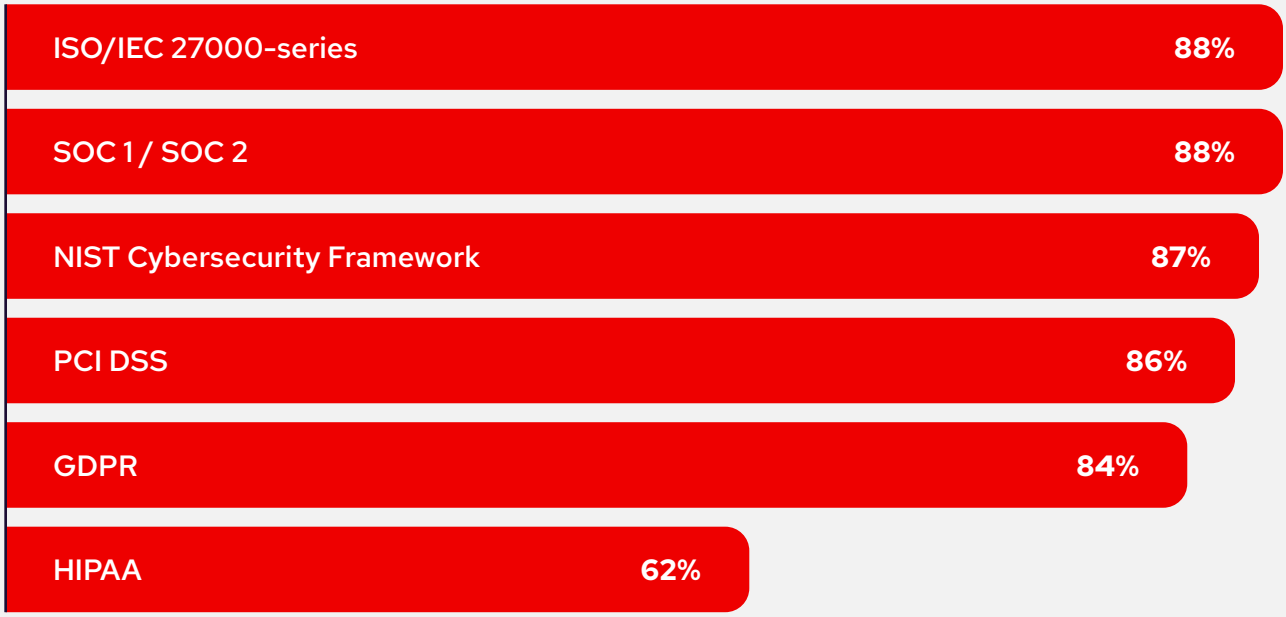


***Figure 5.*** *The response when the surveyed organizations were asked "To what extent do you expect each of the following to impact your organization's cloud-native security strategy over the next 12 months? (reporting some or strong influence)."*

Another critical aspect of cloud governance today is compliance and regulatory alignment. Companies are feeling pressure to formalize security not just from within, but from external requirements. The report indicates that emerging regulations are "turning governance into a requirement." For instance, **64%** of respondents expect the new EU Cyber Resilience Act (CRA) to impact their cloud-native security investments in the next year. Likewise, industry frameworks and standards are exerting broad influence. Whether it's ISO/IEC 27000-series, SOC 2, the National Institute of Standards and Technology (NIST) Cybersecurity Framework, payment card industry data security standard (PCI-DSS), or general data protection regulation (GDPR), a large majority of organizations across regions report that these standards have a strong influence on their security strategy and tooling decisions.

## The implication is clear: aligning early with common security frameworks can pay off.

As the report notes, organizations that embrace shared standards sooner will likely "reduce future cost and complexity" in compliance. In practice, this means governance is no longer optional or an ignorable item; it's rapidly becoming a baseline expectation for doing business in the cloud.

Many teams have the right mindset and recognize the importance of security, but fewer have translated that into structured strategies and full-spectrum controls. The data suggests a need for more organizations to formalize their cloud security programs—defining a clear strategy, implementing uniform guardrails, and embracing frameworks—so that their proactive intentions are backed by real preparedness. When done right, such governance pays off in resilience.

Chapter 3

# Emerging investment trends: Focus on automation and supply chain

Given the challenges outlined in the first 2 chapters, it's no surprise that organizations are adjusting their security investments to address those gaps. The survey points to a clear trend: security efforts in the cloud-native space are entering a phase of consolidation and automation. Rather than spreading resources thinly across too many disparate tools or piecemeal fixes, organizations are concentrating on a few critical priority areas that will harden their cloud-native environments most effectively.

## The top investment areas for 2024-2025 all center on building security into the software lifecycle and infrastructure.

## In particular, 3 themes stand out (each cited by over half of organizations as a planned investment):

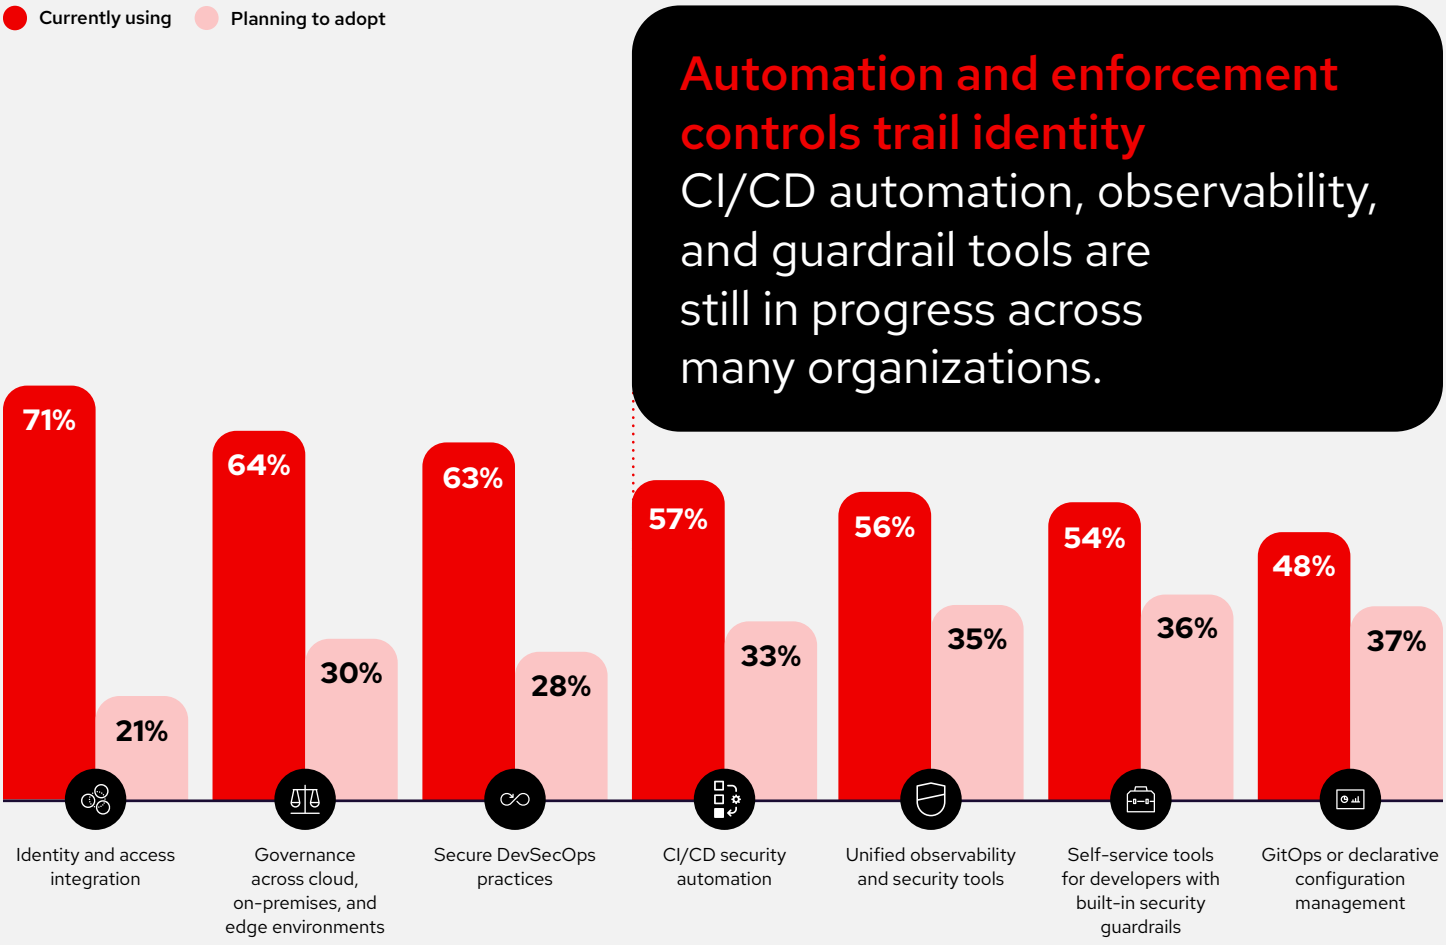


***Figure 6.*** *Automation and policy enforcement lag behind core controls. This gap limits visibility and prevents organizations from fully realizing the benefits of DevSecOps maturity.*

# Automating DevSecOps pipelines

Automation will be a major focus, as **6 in 10** of the organizations interviewed are looking to integrate security into CI/CD pipelines and development workflows. This includes automating policy enforcement and security checks across environments, so that security is not a manual gate at the end but an embedded part of the deployment process. The goal is to catch issues early and consistently (for example, automated code scans, configuration checks, and guardrails in every build/deploy).

**By shifting to "security as code" and automated controls, teams aim to reduce human error and accelerate safe software delivery.**

# Securing the software supply chain

**More than half** of organizations will invest heavily in software supply chain security. This reflects rising awareness that the code and components flowing into applications (open-source libraries, container images, build artifacts, etc.) must be verified and protected. Supply chain attacks—such as tampering with dependencies or injecting malicious code into upstream components—are a growing threat. Managing integrity from code to runtime is the objective here. Initiatives include using tools for software composition analysis, dependency scanning, artifact signing to verify provenance, and security-focused systems.

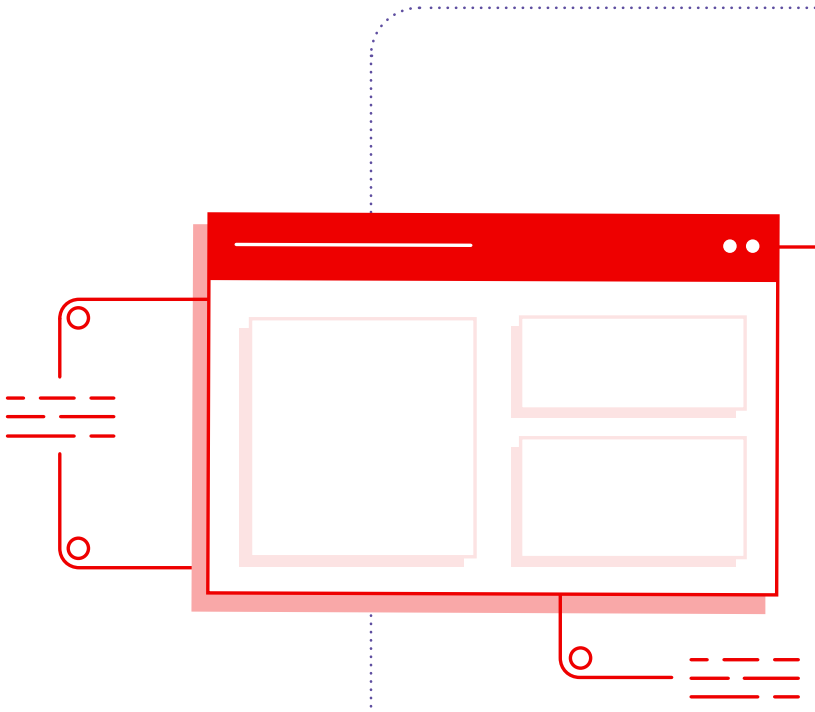

> "Supply chain attacks are soaring because everyone uses open source, but hardly anyone scans or signs their dependencies."
>
> Software engineer (UK)

# Expanding runtime protection

**Just over half** of respondents also prioritized strengthening runtime security in their production environments. This means deploying solutions like container runtime protection, real-time threat detection, and automated response capabilities in clusters and cloud workloads. Many teams have already invested in detection (finding issues) and are now moving toward more integrated, active defense—for example, continuous monitoring of workloads, anomaly detection, and self-healing or blocking of attacks at runtime.

**By embedding a continuous defense within the platform, organizations aim to catch incidents that slip past earlier gates and to limit damage.**

For example, detecting a rogue container behavior or a crypto-mining process and shutting it down immediately.

# Automation and guardrails

Underpinning these specific areas is a broader strategy: invest in automation and guardrails that make security continuous and scalable.

## Investment choices are mirroring the gaps identified in maturity assessments.

Teams are directing resources toward the very capabilities that distinguish mature security programs. In other words, organizations are learning from the data. Since lack of automation and inconsistent guardrails are holding security back (as shown in Chapter 2), budgets are now shifting to fix those issues. Instead of adding more point security tools, there's a push to bake security deeply into development and operations.

This investment shift is also influenced by the external factors discussed earlier (compliance and fear of breaches). With regulations such as the CRA on the horizon, companies want to be ahead of the curve by automating compliance and securing their supply chains now, rather than scrambling later. As well, high-profile supply chain attacks (e.g. dependency hacks) have been a wakeup call, hence the surge in focus there. Companies can use this opportunity to not just to view compliance as a checklist, but to adopt software bills of materials (SBOMs) to do more than just meet regulatory requirements and also prevent tampering, provide transparency, and streamline incident response.

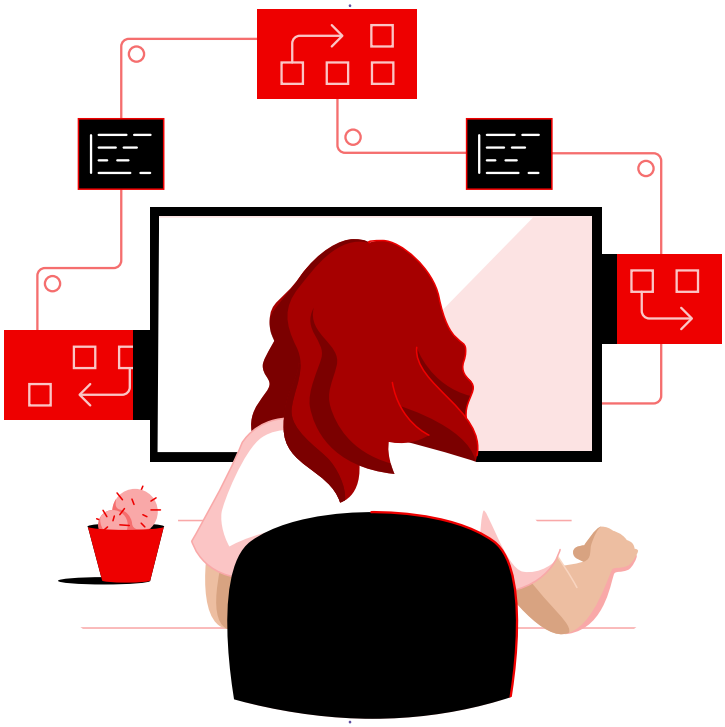

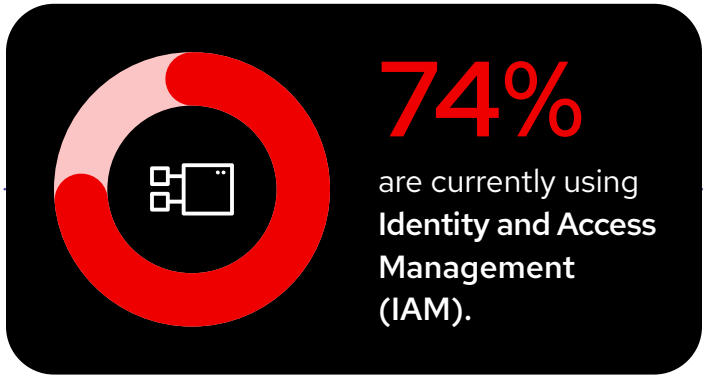


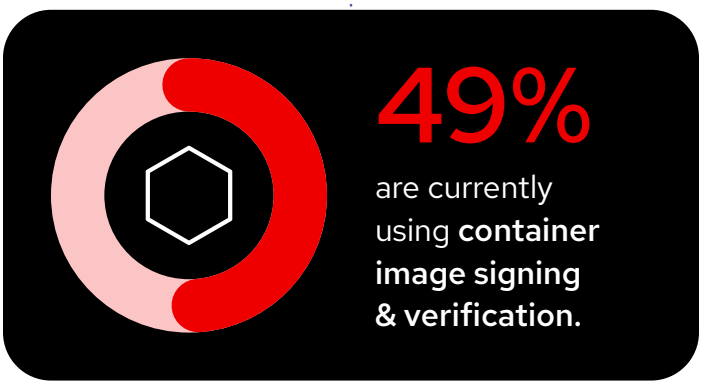


***Figure 7.*** *Guardrails define maturity, but adoption remains inconsistent.*

# Platform consolidation

Another trend is platform consolidation, as many organizations are looking for platforms or unified solutions that cover multiple security needs, rather than a separate tool for each issue. By consolidating, they hope to get that end-to-end visibility and control (from code commit to runtime) in a more streamlined way. Yet only **42%** of organizations surveyed reported investment in adopting cloud-native-application protection platforms (CNAPP).

## The investment patterns in cloud-native security for 2026 show a maturing market.

Companies are putting their money into the fundamentals: working to put a security focus on what they build (supply chain), automating how they put a security focus on it (DevSecOps pipelines), and protecting where it runs (runtime defenses). These are proactive, architecture-level improvements, not just reactive add-ons. Over the next year or two, we can expect the average organization's security toolkit to become more automated and more integrated. The outcome, if these investments are executed well, should be fewer last-minute surprises (as security checks become part of the assembly line) and fewer breach opportunities (as code integrity and runtime monitoring improve).

**42%**

of organizations surveyed reported investment in adopting cloud-native-application protection platforms (CNAPP).

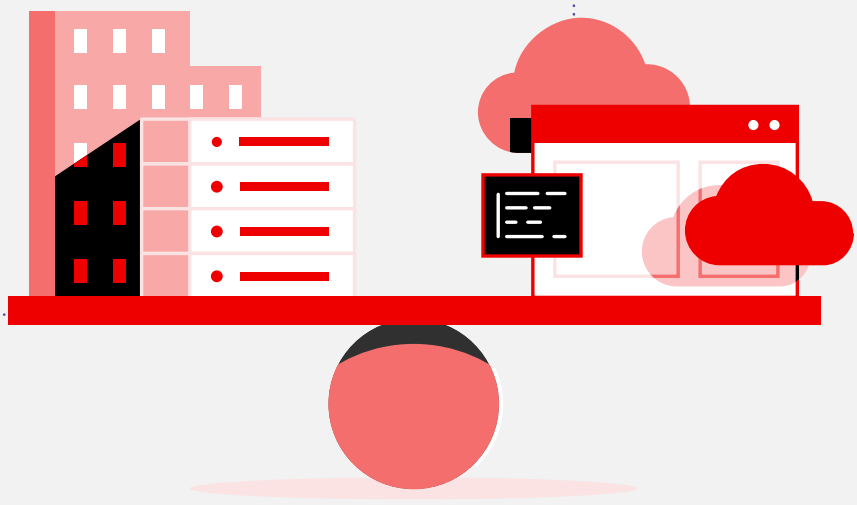

Chapter 4

# The emerging risk frontier: AI and cloud security

No report on technology in 2026 would be complete without addressing AI, and indeed, AI has rapidly emerged as a double-edged sword in cloud-native environments. On 1 side, AI and machine learning (ML) offer powerful capabilities and efficiencies. On the other side, they introduce new security concerns that many organizations are still grappling with.

## The survey results make it clear that while enthusiasm for AI is high, security governance around AI is lagging dangerously behind.

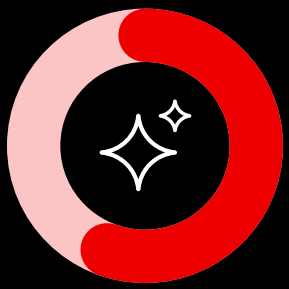

**58%**

say AI adoption significantly shapes their security planning.

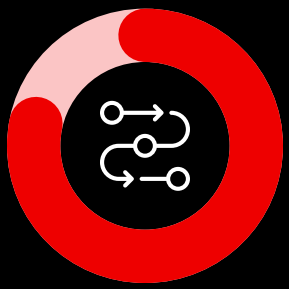

**79%**

agree AI is creating new challenges in their environments.

***Figure 8.*** *AI expands innovation and attack surfaces alike.*

# Everyone is concerned

To start with, virtually everyone is concerned about AI-related risks. An overwhelming **96%** of respondents said they have worries about the use of gen AI in their cloud environments. These concerns are not abstract; they stem from real observed issues and uncertainties. The top AI-related security concerns reported include things like exposure of sensitive data, the presence of shadow AI tools (employees or teams using AI SaaS tools or APIs without approval), and the integration of third-party AI services expanding the attack surface.

## In essence, organizations fear the unseen risks that AI integration might bring.

For example, an engineer might inadvertently feed proprietary code or data into a gen AI service, creating a data leakage risk. Or a team might deploy an AI-based application that has hidden vulnerabilities or that makes security-impacting decisions without proper oversight. There's also concern about how AI systems could be abused, generating convincing phishing content, or introducing logic that traditional security tools don't catch.

> "AI is an overlooked risk. It has to be regulated and made secure so there is confidence in adopting it."
>
> IT operations lead (UK)


**96%**

of respondents said they have worries about the use of gen AI in their cloud environments.

# Yet strong governance is elusive

Despite these widespread fears, most companies have not yet implemented strong governance for AI usage. According to the data, 59% of organizations we spoke to do not have any documented AI-related security policies or guidelines in place.

**In other words, fewer than half have established rules for how developers and employees should safely use AI tools, how AI models should be vetted, how data should be handled.**

A portion are in early stages of drafting some guidelines, and some rely on ad-hoc team-specific rules, but the overall picture is that AI governance is frontier territory right now. This lack of formal policy is a major governance gap, especially considering the speed at which AI adoption is happening.

The mismatch between AI adoption and AI oversight can lead to serious issues. Businesses see the upside of AI, but if they don't put guardrails around its use, they may inadvertently create new vulnerabilities or compliance headaches. For instance, who is accountable if an AI service introduces a security bug? How do you monitor AI-generated decisions or outputs for security implications? These questions often remain unanswered in organizations that lack AI policies.

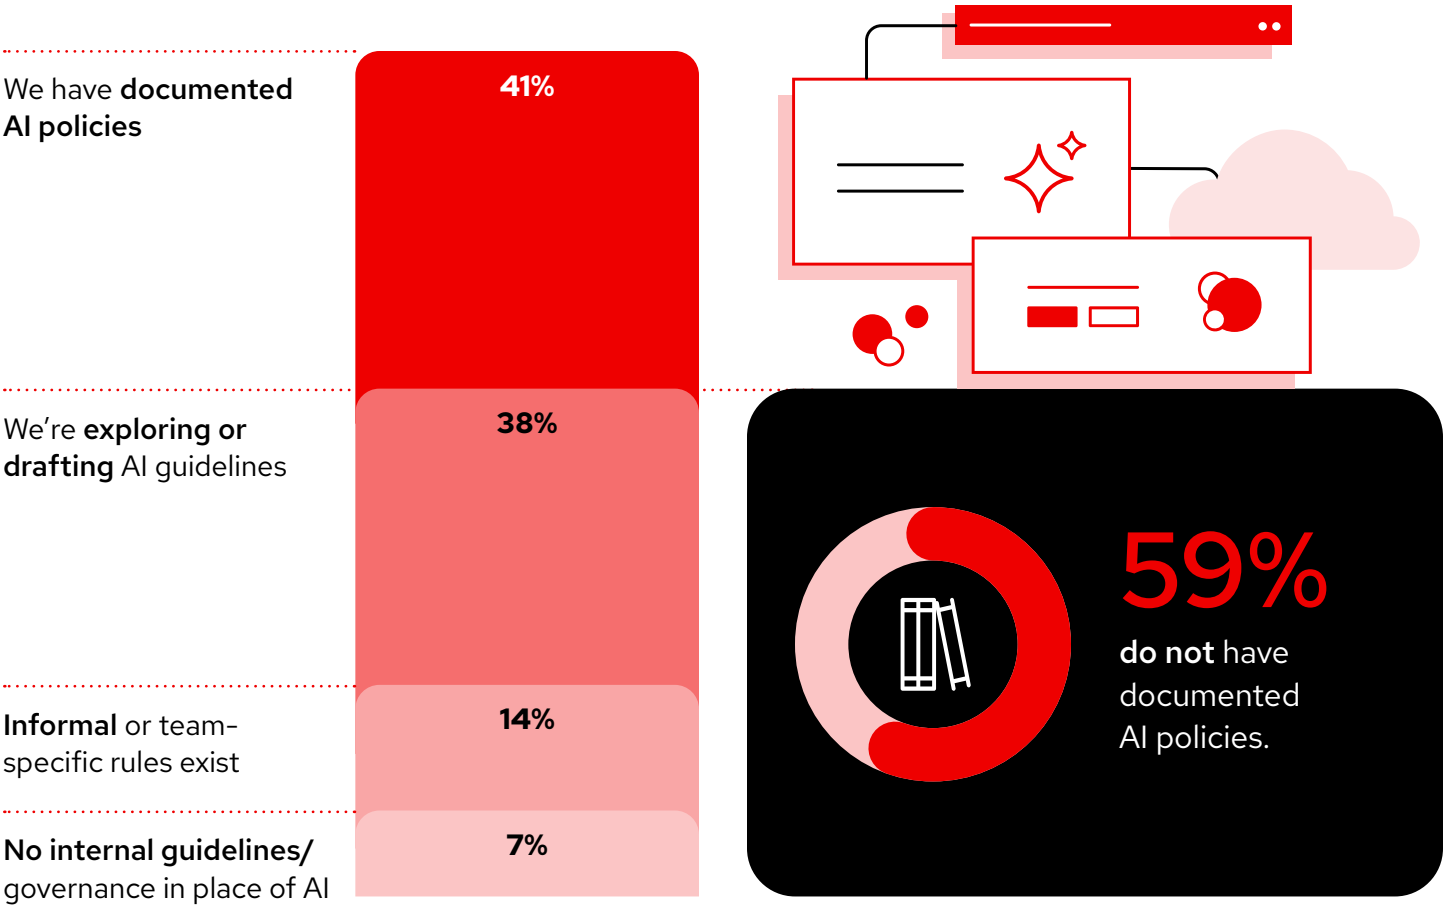


***Figure 9.*** *AI policies lag adoption, with fewer than half of organizations enforcing standardized AI guidelines.*

# AI multiplies risk

Another aspect is that AI can amplify existing security issues. Without proper controls, AI tools might exacerbate identity and access risks, CI/CD pipeline risks, and supply chain issues—the very areas we discussed earlier. An uncontrolled AI script might, for example, spin up cloud resources or alter configurations outside of normal processes (a shadow IT scenario).

## The top AI fears (data leakage, shadow tools, third-party AI) stem from weak visibility and fragmented controls.

Without shared ownership and policies that travel with applications, these risks will scale with adoption. In plain terms, if you don't extend your governance framework to include AI, the more AI usage grows, the greater the potential chaos or exposures.

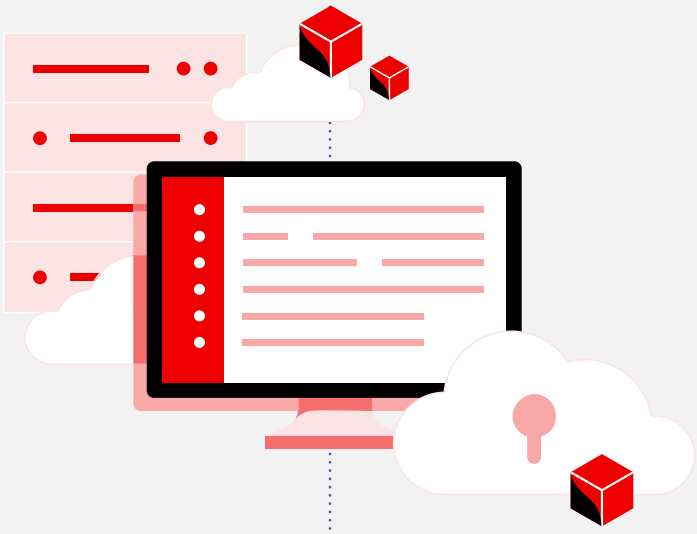

# There are bright spots

It's not all doom and gloom, some organizations are taking steps. A number of respondents indicated they are exploring or drafting AI guidelines now, and a few have internal committees or oversight for AI. Awareness is the first step, and the near-universal concern is forcing leadership conversations about how to tame rogue AI usage. There are also calls for external guidance. Governments and industry bodies are starting to discuss AI regulations (e.g., the EU's proposed AI Act), which may eventually impose requirements similar to data protection laws. But companies can't wait for that. They need to be proactive.

**AI represents a new frontier of risk in cloud-native security, 1 that most organizations are only beginning to get a handle on.** The year 2026 will likely see rapid evolution in this space. On 1 hand, more AI-powered tools and features in DevOps, and on the other hand, a scramble to establish governance around them.

## The key takeaway is that the speed of governance must catch up to the speed of innovation.

Organizations should treat AI in the cloud with the same rigor as they treat any other powerful technology—with clear policies, monitoring, and controls. Those that fail to do so may find that AI, meant to accelerate their business, could instead become the source of the next big security incident.

> "I think AI governance is critical, and we're working on putting clear rules in place."
>
> **Software engineering lead, New Zealand**

# Conclusion

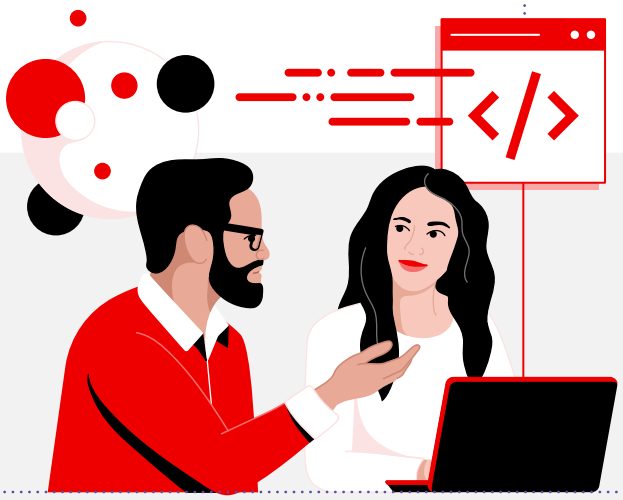

## Data-based recommendations for 2026

The findings of the study highlight several pressing areas where organizations should take action. Below are key recommendations to improve cloud-native security outcomes, based on the report's insights.

**61%**

of the companies in our research lack a defined cloud-native security strategy.

### Establish a formal cloud security strategy and maturity roadmap

If your organization lacks a defined cloud-native security strategy (as is the case for **61%** of the companies in our research), make it a priority to create it.

**A clear strategy, possibly using a cloud security maturity model, will provide a structured path from reactive to proactive security posture.**

It helps keep security efforts aligned with business objectives and keeps all teams informed of the plan. The data shows that companies with well-defined strategies have far greater confidence and consistency in their security programs. Investing time in strategy and architecture now will prevent ad-hoc firefighting later.

## Embed security guardrails and automation into the development lifecycle

The organizations that succeed in cloud security treat it as an integral part of their platform and pipeline, not a bolted-on extra.

**Teams should implement platform guardrails at every stage, from security focused coding standards and pre-commit checks, to Infrastructure as Code (IaC) scanning in the command line, to continuous runtime monitoring in production.**

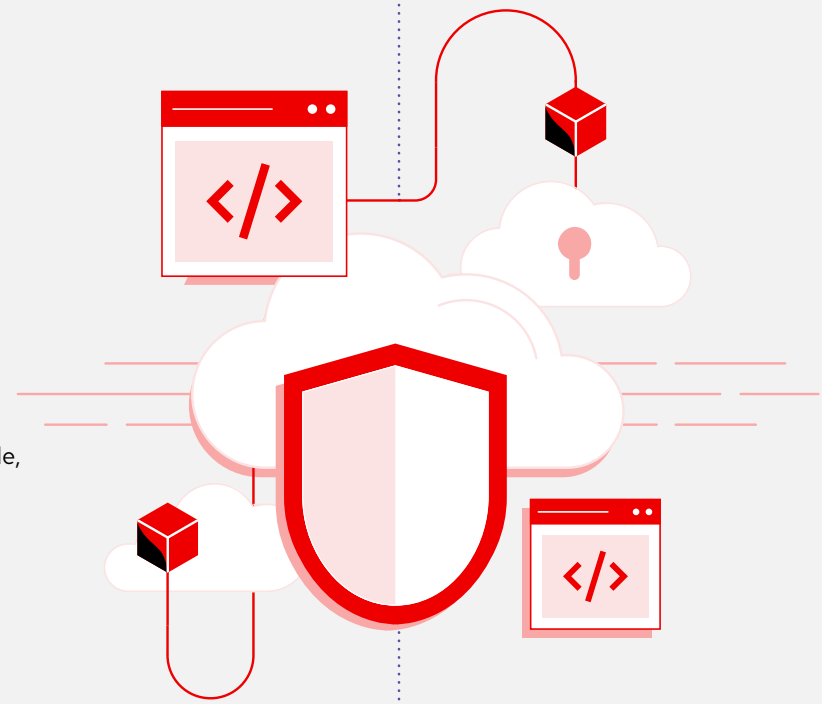

Aim to automate these controls wherever possible (policy-as-code, automatic vulnerability scans, etc.). Automation not only catches issues early but also maintains consistency at scale. This directly addresses the frequent misconfigurations and human errors that cause most incidents. Essentially, make "secure by default" the norm. As the data puts it, organizations need to move from security aspiration to security execution by building controls into their workflows. An automated, unified security platform, such as the adoption of a CNAPP (a priority for **42%** of organizations) will reduce the chance for things to fall through the cracks. This shift requires defining an organizational mandate, often executed by DevOps or platform engineering teams, to scale security without imposing friction on developers.

## Prioritize software supply chain integrity

**To protect your supply chain, implement measures such as dependency scanning, SBOMs, tracking, and image signing for all container images and packages.**

Currently, only about half of organizations are doing image signing and other supply chain security practices, which means many are exposed. Every team consumes open source components; make sure you trust and verify them. Enforce provenance checks (e.g., require signed artifacts) and use tools to detect vulnerable or malicious components before they hit production. A survey respondent noted that it's common to use open source but "hardly anyone scans or signs their dependencies"—make sure your organization is the exception. By shoring up the software supply chain, you cut off a growing avenue of attack and check that what you build and deploy hasn't been tampered with upstream.

## Close the maturity gap through unified visibility and the full-lifecycle feedback loop

**Unify observability and security data across teams, rather than operating in isolated team structures.**

This unification is critical to establishing a full-lifecycle security feedback loop. While organizations are investing heavily in both DevSecOps automation (**60%**) and expanding runtime protection (**54%**), these efforts must be connected. Mature security requires using insights derived from runtime threat detection to prioritize and fix the most critical vulnerabilities earlier in the development and build processes. By extending security across the full life cycle, from build/deploy to runtime, and feeding back intelligence, teams can ensure consistent guardrails and accelerate safe software delivery, turning the DevOps user into a security user.

## Hybrid cloud security posture through sovereign cloud and edge deployments

### Adopt security controls that function across clouds, on-premise data centers, and the edge.

Cloud-native security tooling must be able to operate consistently in disconnected, sovereign, or data-resident environments. Implement tools that enforce policies uniformly, regardless of whether the workload is in a public cloud, a private cloud, or on-premise. Edge security for distributed workloads is an investment priority for **38%** of organizations. Security tooling must be lightweight and autonomous enough to protect these remote environments without continuous, high-bandwidth connections back to a central security team.

## Align with security frameworks and compliance early

### Don't wait for an audit or regulation to force your hand—proactively adopt industry security frameworks that are relevant to your business.

The research indicates that companies who align with shared standards can reduce long-term complexity. Embracing best-practice frameworks provides a structured checklist to harden your cloud environment, covering areas like identity, access, monitoring, incident response. It also prepares you for emerging regulations. For example, if you might be impacted by the EU Cyber Resilience Act, start assessing its requirements now. **64%** of organizations in our survey expect this to have influenced their 2026 investments. By building compliance into your strategy, you avoid last-minute scrambles and keep security and governance unified. In summary, treat compliance as a floor, not a ceiling and use it to bolster your security fundamentals.

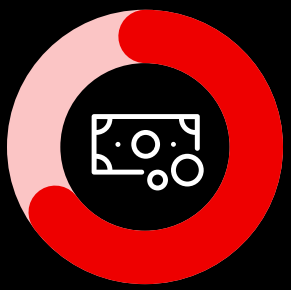

**64%**

of organizations in our survey expect this to have influenced their 2026 investments.

## Implement robust AI governance and policies

**Given the rapid infusion of AI into cloud applications and DevOps,and the associated risks, organizations should put in place clear AI usage policies and oversight as soon as possible.**

The fact that nearly **60%** of companies surveyed have no AI governance today is a huge gap that needs closing. Convene a cross-functional team (security, IT, data science, legal) to develop guidelines on acceptable AI use. This should include how sensitive data can or cannot be used in AI services, what approvals are needed for deploying AI-based solutions, and how to monitor AI outputs for security and ethical issues. Educate your developers and engineers on these policies. Additionally, consider technical controls for AI, such as data tagging to prevent export of confidential data to external AI application programming interfaces (APIs), or monitoring for unusual AI-powered behaviors.

By executing on these recommendations—developing a strategy, building in guardrails/automation, securing the supply chain, aligning to standards, and governing AII, organizations will position themselves to dramatically improve their cloud-native security posture.

The 2026 outlook shows threats continue to evolve, but also that there is more data than ever to inform our defense.

**60%**

of companies surveyed have no AI governance.

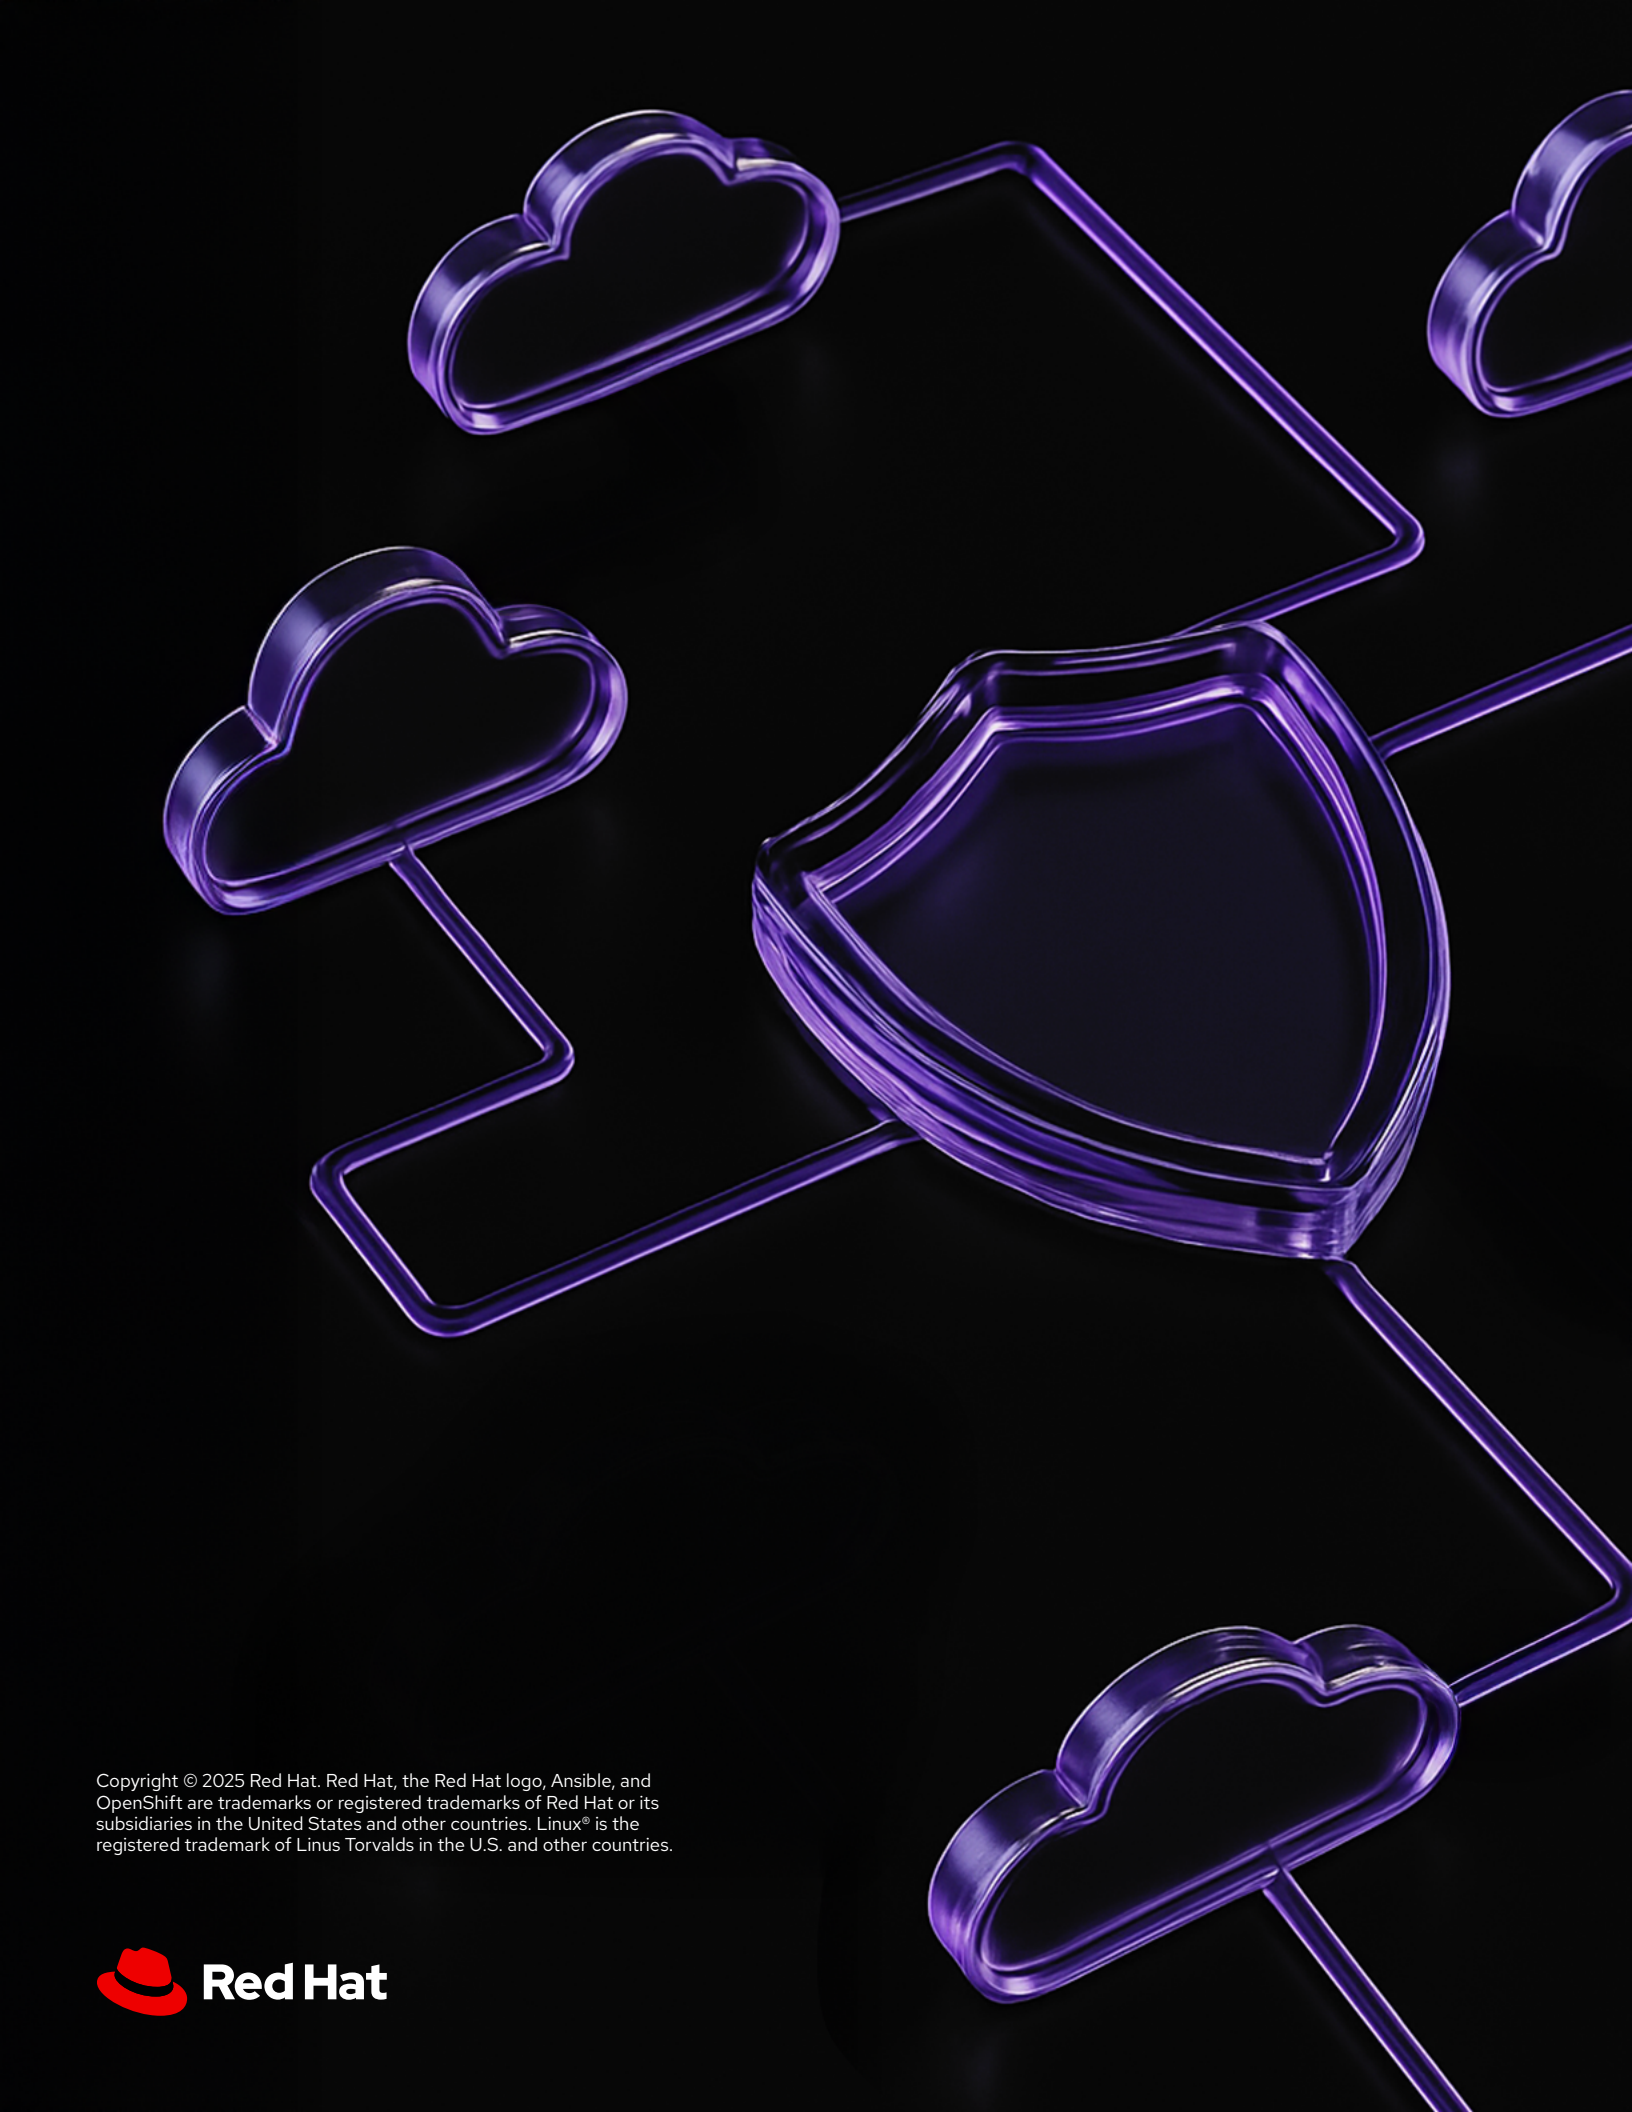



Red Hat