cloud security alliance®

algosec

Survey Report

The State of **Hybrid and Multi-Cloud Security Policy Management**

Acknowledgments

Lead Authors

Hillary Baron

Contributors

Marina Bregkou
Josh Buker
Ryan Gifford

Graphic Design

AnnMarie Ulskey

About the Sponsor

AlgoSec, a global cybersecurity leader, empowers organizations to secure application connectivity by automating connectivity flows and security policy, anywhere. The AlgoSec platform enables the world's most complex organizations to gain visibility, reduce risk, achieve compliance at the application-level and process changes at zero-touch across the hybrid network. Over 1,800 of the world's leading organizations trust AlgoSec to help secure their most critical workloads across public cloud, private cloud, containers, and on-premises networks.

Table of Contents

Executive Summary

Hybrid and multi-cloud have become the operating reality for most enterprise applications. Business-critical workloads now run simultaneously across public cloud, private cloud, and on-premises infrastructure, while regulatory frameworks like DORA and NIS2 raise the bar on continuous compliance and operational resilience. The security policies that govern how those applications connect—across environments, across teams, across tools—have not kept pace. What was manageable as a network-configuration task in a simpler environment has become an application-connectivity problem in a fragmented one, and the friction is showing up in production.

1. Production Is Paying the Price for Manual Policy Management

Nearly half of organizations describe their security policy changes as mostly or fully manual (48%), and manual configuration errors emerge as the highest-impact bottleneck to deploying new applications. The result shows up in production: **65% of organizations experienced at least one business-critical application outage caused by a misconfigured security policy in the past 12 months**, and 46% experienced two or more. Delayed deployments, emergency rollbacks, and audit findings compound the operational cost.

2. Fragmented Ownership Creates the Conditions for Misconfiguration

Responsibility for security policy is now distributed across at least four teams—security operations, network operations, cloud architecture, and DevOps—none of which owns the function alone. Two-thirds of teams operate across three or more security consoles daily, and **92% of respondents report at least some difficulty getting a single, accurate view of policies across their environments.** Misconfiguration becomes an emergent property of the operating model, not a failure of individual attention.

3. Manual Compliance Effort Is Not Producing Compliance

Manual review is the most common compliance posture, cited by 40% of organizations. Yet one in four respondents reported a failed compliance audit or audit finding tied to policy gaps in the past year, and **38% take more than 3 days to remediate a production policy issue.** Manual compliance may still be necessary, but it is no longer sufficient on its own in environments where policy changes continuously.

4. Organizations Are Requesting the Application Before the Foundation Exists

The single most-requested capability is risk or impact analysis before a policy change is committed, cited by 32% of respondents—twice the rate of any other capability. But that analysis depends on a unified view of policy state and application connectivity that most organizations say they struggle to assemble. **Organizations are asking for the analytical capability before the visibility foundation it requires exists.**

5. The Path Forward Is Operational Redesign, Not Incremental Spending

Only 44% of organizations expect any budget increase for 2026, and only 7% expect a significant one. Yet **only 9% describe their security policy management as fully integrated into development and deployment workflows; 61% remain in manual or reactive postures.** Closing the gap between preventive intent and reactive practice is a structural problem—one that requires consolidating visibility, clarifying ownership, and embedding policy into the deployment workflow, not another line item.

Takeaway

Security policy management has quietly become central to how business-critical applications connect across hybrid and multi-cloud environments. The findings in this report show that the current operating model—fragmented across teams, tools, and environments, and still largely dependent on manual work—is producing measurable costs in production uptime, deployment velocity, and compliance readiness. The organizations that will close the gap will treat this as an operational redesign centered on the applications they run, not as a matter of adding more tools or effort to a model that is straining under conditions for which it was not built.

Key Findings

Business-critical applications no longer live in one place. They run across public cloud, private cloud, and on-premises infrastructure at the same time, and the security policy that governs how they communicate has to move with them—rewritten constantly, spread across more environments than any single team can see, managed through more consoles than any one team can reconcile, changed by hand by teams working from different views of the whole. What was once a network-configuration problem has become an application-connectivity problem, and the older, infrastructure-centered lens—defining policy device by device and rule by rule—is a poor fit for the newer reality. For a long time that mismatch was merely inefficient. It is now becoming a production risk—less because people are careless than because an operating model built for a simpler environment is straining under a fragmented one.

Key Finding 1:

Two in Three Organizations Took a Production Hit from a Policy Misconfiguration

Security policy misconfiguration is often treated as a hygiene issue—something to be caught in review, corrected, and forgotten. The data points to something closer to a recurring operational failure. For most organizations, a misconfigured security policy is not a rare event that occasionally slips through; it is a routine cause of production disruption, and it sits on top of a policy management model that remains predominantly manual. That disruption is no longer confined to the security backlog: it affects application uptime, release velocity, emergency change volume, and audit readiness.

That manual foundation is still the norm. **Nearly half of organizations describe their security policy changes as mostly or fully manual (48%), while only 22% describe them as mostly or fully automated**, and the remaining 31% report a balanced mix. When teams ranked the biggest bottlenecks in

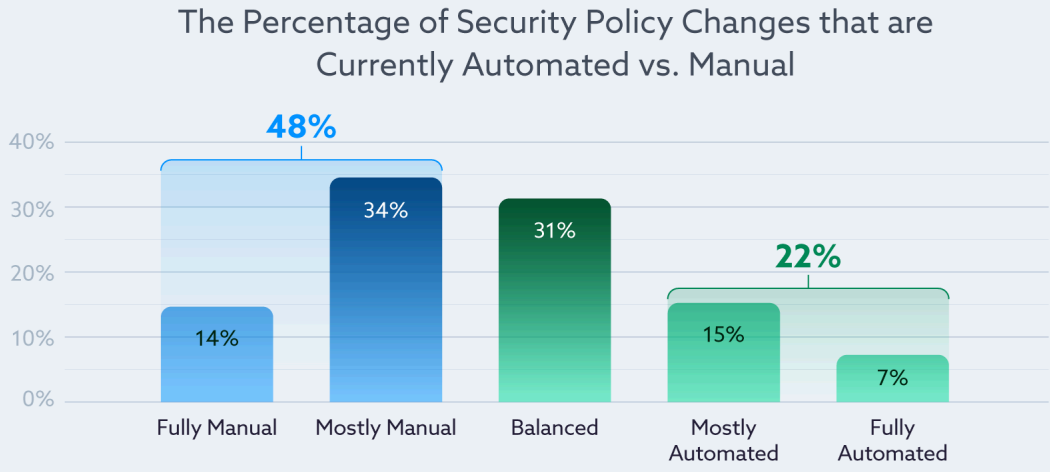

deploying a new application, manual configuration errors leading to rework carried the heaviest weighted impact (3.24). The next two bottlenecks—cross-team coordination (3.37) and getting security policy approvals (3.42)—fall within 0.18 points of the top, so the three are best read as a tight cluster rather than a single clear leader. Even read that way, the most disruptive bottlenecks are concentrated in human-mediated policy work rather than in the technology itself. That concentration reflects a broader pattern in how policy has traditionally been organized: at the infrastructure level, rule by rule and device by device, rather than starting from what applications actually need to connect to. In hybrid and multi-cloud environments, that mismatch is what shows up as the coordination, approval, and configuration friction the data describes.

The production consequences are measurable. **In the last 12 months, 65% of organizations experienced at least one business-critical application outage that respondents attributed to a misconfigured security policy, and 46% experienced two or more.** A quarter reported no such outage (25%) and 11% were unsure. Outages are also only the most visible cost. Within the same period, respondents reported that misconfiguration or policy gaps were associated with delayed application deployments (40%), near-miss security incidents (34%), unplanned rollbacks or emergency changes (31%), failed compliance audits or audit findings (25%), and actual security incidents or breaches (18%). The disruption is spread across delivery speed, operational stability, audit standing, and security posture at once.

When something does go wrong, the exposure does not close quickly. **Only 48% of organizations can remediate a production policy issue within 3 days** (less than 24 hours, 23%; 1–3 days, 25%). For 38%, remediation takes 4 days or longer (4–7 days, 27%; 1–2 weeks, 5%; more than 2 weeks, 6%), and another 14% cannot reliably estimate how long it takes at all (varies too widely to estimate, 9%; unsure, 5%). That one in seven teams cannot even characterize their own remediation time is itself a signal of limited operational visibility into the policy layer.

48%

of organizations can remediate a production policy issue within 3 days

The pattern is a production reliability problem with a policy root cause. The failures organizations are absorbing—outages, rollbacks, delayed releases, audit findings—are the downstream cost of managing a high-volume, high-consequence control surface largely by hand. This is the first movement of the story the data tells: manual policy management has crossed from inefficiency into operational risk, and that risk is materializing where the business can least afford it.

What this means in practice

For the CISO, this reframes policy misconfiguration from a control-hygiene metric into an enterprise risk with a recurring, quantifiable production footprint—one that belongs in board-level risk reporting rather than an operational backlog. For the CIO, the cost lands where the business feels it: application availability, delayed releases, and emergency change activity that disrupts delivery. For cloud and network teams, the multi-day remediation window means manual effort is not only error-prone at the point of change but slow to recover afterward, extending every exposure. The business implication is direct: policy work organized around application connectivity—rather than around individual rules or devices—is what would let security keep pace with application delivery rather than slow it down. The reasons these windows stay open have less to do with individual errors than with how the work is structured across the organization—which is where the next finding begins.

Key Finding 2:

Everyone Owns Policy. No One Has the View.

A misconfiguration is usually described as a mistake, but mistakes happen inside conditions that make them likely. The conditions for security policy management in hybrid and multi-cloud environments are fragmented by default: responsibility is divided across teams, daily work is divided across consoles, and the environments themselves are divided across clouds and data centers. There is no single source of truth for policy across those environments, and every team is making decisions from a partial map—the second movement of the story, and the structural setup in which misconfiguration becomes likely rather than exceptional.

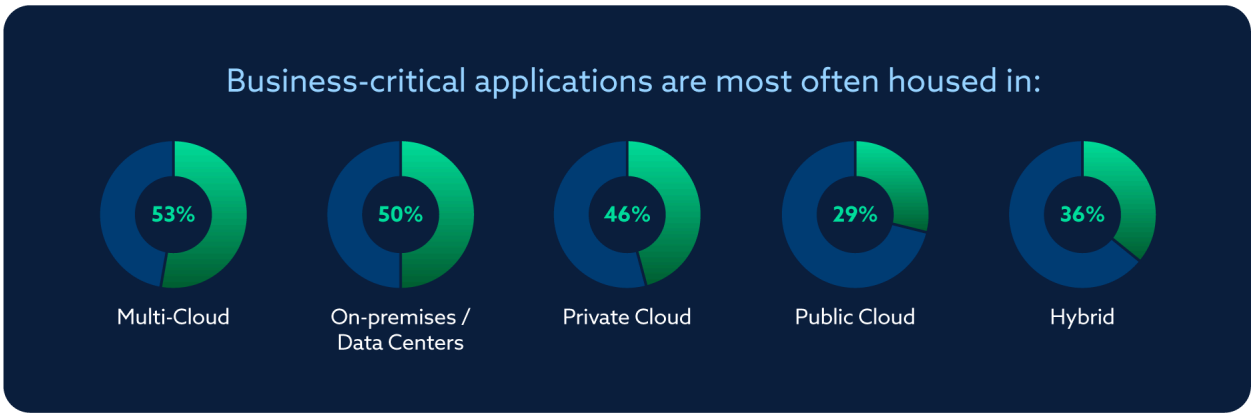

The environment itself resists consolidation. **Business-critical applications are most often housed in multi-cloud deployments (53%) and on-premises data centers (50%)**, with private cloud (46%), hybrid (36%), and single-provider public cloud (29%) close behind. Because organizations selected all that apply, the picture is not one environment giving way to another but several operating in parallel; no single environment dominates, and most organizations manage policy across more than one at the same time.

Ownership is distributed in the same way. **Four distinct teams are most commonly responsible for defining application connectivity policy within organizations: Security Operations (51%), Network Operations (46%), Cloud Architects (46%), and DevOps or application owners (41%)**, with CISO and security leadership (35%) and GRC and compliance (23%) also involved. No team owns the function outright. Connectivity policy is a shared responsibility, which means consistency depends on coordination rather than on a single accountable owner.

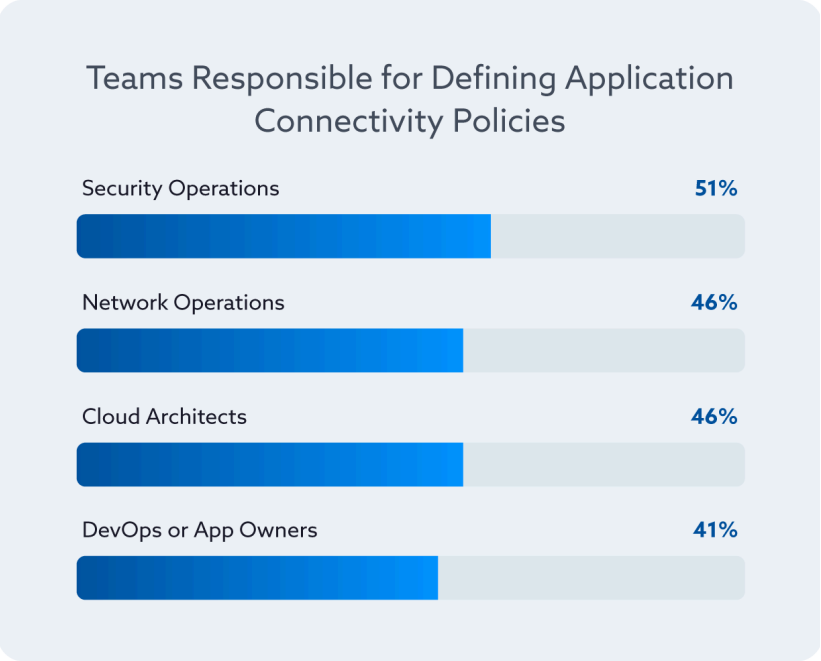

That coordination plays out across a sprawl of tools. **Two-thirds of teams use three or more security management consoles daily to manage connectivity (67%)**, and 16% use six or more. Each console is a separate place where policy can be defined, changed, and viewed differently from the others. Confidence in cross-environment visibility is correspondingly limited: fewer than half of respondents rate their visibility across hybrid and multi-cloud environments as strong, a 4 or 5 on a five-point scale (49%), while a plurality sit in the middle at 3 (38%). Asked more directly, 92% report at least some difficulty getting a single, accurate view of security policies across all their environments (slightly difficult, 33%; moderately difficult, 40%; highly difficult, 19%), and 59% rate that difficulty as moderate or high. Only 7% find it not difficult.

Two-thirds of teams use three or more security management consoles daily to manage connectivity.

The pattern is a visibility deficit produced by structure, not by negligence. When responsibility is split across four teams, daily work is split across three or more consoles, and applications are split across several environments, the unified view that would let anyone catch an inconsistency before it reaches production does not exist by default. The map of how business applications actually connect across the hybrid infrastructure stays partial for everyone and would have to be deliberately assembled to become whole. This is the substrate beneath the production failures in the previous finding, and it explains why asking teams to simply be more careful does not work: the errors are emergent properties of a fragmented operating model.

What this means in practice

For the CISO, fragmented ownership produces accountability without control: the executive answerable for policy risk depends on the coordinated behavior of four teams and a half-dozen consoles, none of which provides an authoritative view of the whole. For the CIO, console sprawl is a measurable efficiency and consistency tax on the teams running hybrid infrastructure. For cloud and network teams, switching among consoles is where drift between environments originates. For application owners—who think in terms of whether an application can safely connect to the services it needs, not in terms of firewall rules or security groups—the absence of a unified policy view means the impact of a connectivity change is hard to predict before it is made. That last point—predicting the impact of a change before committing it—turns out to be exactly what organizations most want to do, and a finding we'll return to later.

Key Finding 3:

Manual Compliance Means High Effort, Mixed Outcomes

Manual compliance is usually framed as the careful, conservative choice—slower than automation, but safer because a human is checking the work. The data complicates that framing. Manual review is the most common compliance posture organizations report, yet it coincides with a non-trivial rate of audit findings and a remediation cadence that sits uneasily with the continuous compliance organizations say they want. The effort is real; the outcomes are mixed. Manual compliance may still be necessary, but it is no longer sufficient on its own for hybrid and multi-cloud environments that change continuously.

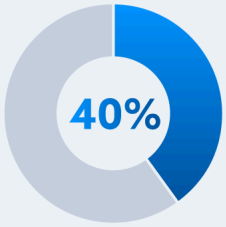

of compliance is maintained primarily through manual processes and reviews

Manual review is the prevailing model. **Asked which phrase best describes their current security posture, 40% said compliance is maintained primarily through manual processes and reviews**—the single most common answer, outnumbering the 9% whose policy management is fully integrated into development and deployment workflows by more than four to one. Separately, one in four organizations reported a failed compliance audit or audit finding in the past 12 months (25%). These are two independent survey measures. The data does not establish that the organizations relying on manual compliance are the same ones experiencing adverse audit findings, and it would be an overstatement to read a direct causal link between the two. What the data supports is more measured: manual, periodic review is the dominant model at the same time that audit findings occur at a meaningful rate across the sample.

The remediation lag sharpens the concern. **With 38% of organizations taking 4 or more days to remediate a production policy issue**, the interval between identifying a policy gap and producing evidence that it has been closed is frequently measured in days. That cadence is difficult to reconcile with continuous compliance, which depends on an always-current evidence trail rather than a point-in-time review. Continuous compliance and audit readiness is itself among the outcomes organizations rate most important (35%), placing just behind speed of application delivery and reduced misconfiguration and human error (both 40%). The priority is consistent with the direction of recent regulatory developments such as DORA, NIS2, and PCI DSS v4.x, each of which leans toward continuous evidence over periodic attestation. The appetite for continuous compliance is present; the operating model underneath it is still periodic and manual.

The reasonable inference is not that manual review causes audit failures, but that manual, periodic compliance may be insufficient for the scale and rate of change in hybrid and multi-cloud environments. Effort spent on manual review does not, on its own, produce the always-current assurance that continuous compliance requires, which means organizations can invest real time in compliance work and still carry exposure they cannot readily demonstrate as closed. This is the same fragmentation problem viewed through a compliance lens, and it points toward the same conclusion the rest of the data builds to: the limit is in the model, not in the diligence.

What this means in practice

For the CISO and the GRC function, high manual effort is not the same as defensible assurance; the gap between the two is where audit risk accumulates. For the CIO, the multi-day evidence lag is incompatible with the continuous reporting that regulators and customers increasingly expect. For cloud and network teams, manual review is a recurring time cost that scales with environment complexity rather than with risk reduced. The capability that would most directly address this gap is one organizations are already asking for—the subject of the next finding.

Key Finding 4:

The Visibility Gap Underneath the Capability Wishlist

When organizations are asked what would most improve their security policy management, their answer reveals as much about what they lack as what they want. The most requested capability quietly assumes a foundation the same respondents indicate they do not have, which returns the story to its second movement: the visibility gap that sits underneath nearly every other ambition.

32%

of organizations chose risk or impact analysis performed before a policy change is committed

The demand signal is clear. **When asked which capability would most improve their security policy management over the next 12 months, organizations most often chose risk or impact analysis performed before a policy change is committed (32%)**—selected at twice the rate of any other option. The next three capabilities cluster well behind at 16% each: AI-assisted anomaly detection on policy traffic, automated policy provisioning end-to-end, and continuous compliance auditing of policies in production. Pre-change risk analysis is not just the leading request; it is the clear priority.

The difficulty is that analyzing the impact of a change before committing it requires first being able to see the full policy state the change would touch—and to understand which application flows and business services depend on that policy state. That foundation is largely absent: **92% of organizations report at least some difficulty getting a single, accurate view of policies across their environments, and fewer than half rate their cross-environment visibility as strong (49%).** A pre-change analysis is only as reliable as the picture of the current state it runs against, and that picture is precisely what most respondents say they struggle to assemble. The most-wanted capability presupposes the unified visibility identified earlier as the central structural gap; organizations are requesting the application before the foundation it depends on exists.

What this means in practice

For the CISO and CIO, the sequencing matters for investment: funding pre-change analysis or automation before unified visibility is in place risks buying a capability that cannot perform reliably. For cloud and network teams, the value of any pre-change check depends on a single source of truth for policy across environments—without it, the analysis runs against an incomplete map. For application owners, reliable pre-change impact analysis is what would let connectivity changes ship with confidence rather than caution. Wanting better decisions at the moment of change is reasonable; getting them depends on first solving the visibility problem underneath, which is less a matter of new capability than of how the operating model is built—the third and final movement of the story.

Key Finding 5:

Prevention Is the Goal. Reaction Is the Practice. Budgets Won't Close the Gap.

Across the survey, organizations express a consistent preference for getting ahead of policy problems rather than cleaning up after them. Pre-change risk analysis is the most requested capability (32%) and reduced misconfiguration and human error is tied for the most important outcome (40%). The aspiration is preventive. The practice is not, and the gap between the two will have to be closed without much new money—the third movement of the story, and the one that determines whether the first two can be resolved.

A useful way to read the survey is as a maturity progression. The four postures respondents could choose from in describing their current security posture map onto four stages of policy management maturity:

- **Stage 1 — Manual:** Compliance is maintained primarily through manual processes and reviews (40%)
- **Stage 2 — Reactive:** Security policy issues are addressed as they are identified or reported (21%)
- **Stage 3 — Proactive:** Automated tools are used to detect and flag potential risks before incidents occur (30%)
- **Stage 4 — Integrated:** Security policy management is fully integrated into development and deployment workflows (9%)

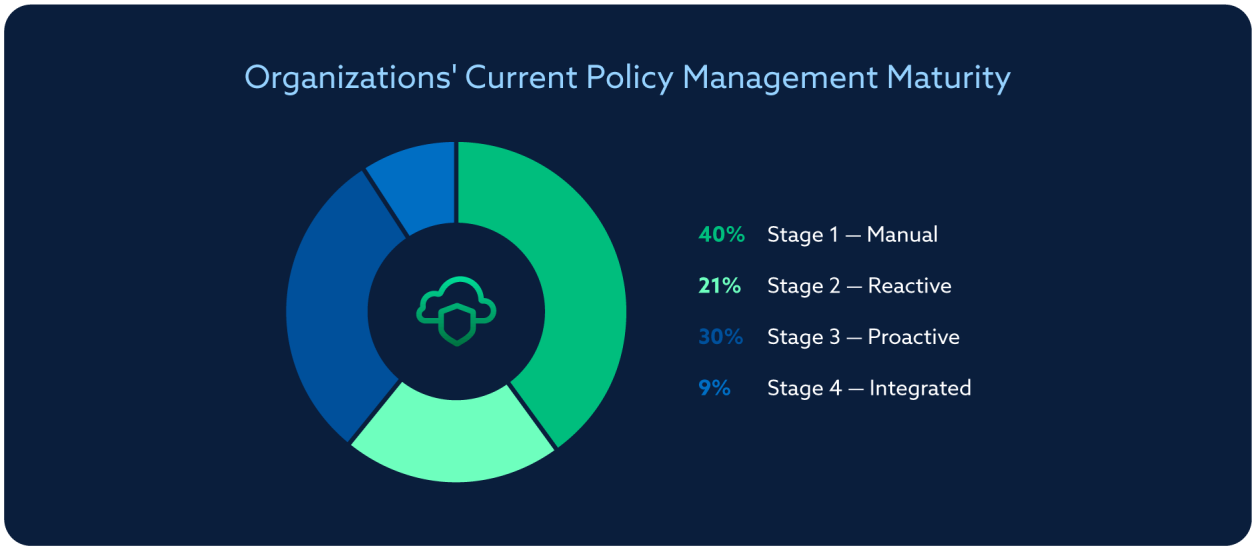

The distribution is bottom-heavy. **A combined 61% of organizations sit in the first two stages**, where policy is governed by manual review or after-the-fact response; 30% have reached the proactive stage, having acquired detection tooling without yet embedding it in how applications are built and shipped; and only 9% operate at the integrated stage that the preventive aspiration actually requires. Most organizations want to operate one or two stages above where they are.

The money that might force the transition is not arriving in volume. **Only 44% of organizations expect any increase in their 2026 security budget, and just 7% expect a significant increase of more than 20%**; most expect budgets to stay the same (39%) or rise only slightly (1–20%, 37%), with 12% unsure. The distance between Stage 1 or 2 and Stage 4 will have to be covered on roughly flat funding.

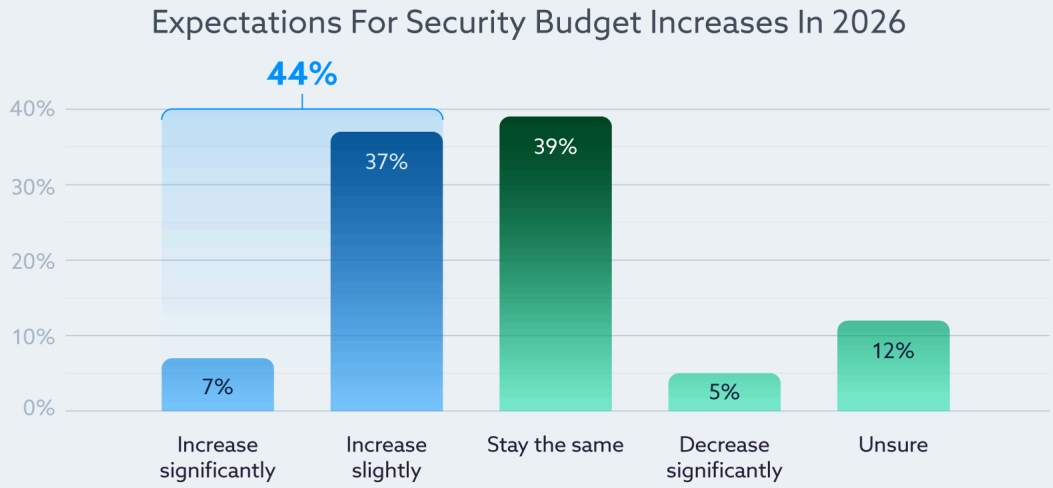

This is the report's central tension, and it points to a structural conclusion rather than a financial one. Advancing through the maturity stages is not something budget alone buys; it depends on consolidating fragmented visibility into a single accurate view, clarifying who owns policy across teams, and integrating policy management into the deployment workflow rather than bolting controls on afterward. Spend layered on top of a fragmented operating model tends to add tools without moving the organization up a stage. The three movements of the story resolve here: manual management became a production risk because fragmented ownership and limited visibility made it one, and the way out is to redesign the operating model rather than to fund the existing one more generously.

What this means in practice

For the CISO, the maturity model offers a way to benchmark the organization's current stage and to frame investment to the board as an operating-model transition rather than a tool purchase. For the CIO, flat budgets mean the path to fewer outages and faster delivery runs through consolidation and integration, not incremental headcount. For cloud and network teams, reaching the proactive and integrated stages means trading manual review and console-switching for automated, pre-change-validated workflows. For application owners and DevOps, the integrated stage is where security policy stops being a deployment bottleneck and becomes part of the pipeline. With most organizations expecting flat or only slightly increased budgets, more manual effort is not a viable path forward. The data points elsewhere: to the unified visibility, automated workflows, and integrated compliance the earlier findings identified as the operating gaps.

Conclusion

The way most organizations manage security policy was built for an environment that no longer exists—fewer environments, a clearer line of ownership, and a slower rate of change. In hybrid and multi-cloud estates, that model is reaching its limit, and the survey shows the strain in concrete terms: production outages, multi-day remediation windows, audit findings, and a preventive ambition most organizations cannot yet act on. These are not separate problems. Manual policy management has become a production risk because fragmented ownership and limited visibility made it one, and the distance between the reactive practice most organizations live in and the preventive posture they want is one that flat 2026 budgets will not close on their own. The constraint is the operating model itself.

That reframes the path forward. The question is not which tool to add, but which capabilities a modern operating model has to have—and the data points to four such capabilities that build on one another in a deliberate order rather than arriving as separate purchases:

- **Unified visibility across environments comes first.** A single, accurate view of policy spanning cloud and on-premises is the foundational gap, reported by the 92% who cannot readily assemble one today. Nothing above it is reliable until it exists
- **Pre-change risk analysis builds on that visibility.** The capability organizations most want (32%) only works once there is a complete picture of the current state to analyze against, letting the impact of a change be evaluated before it is committed rather than discovered in production
- **Automation removes the manual surface where errors originate.** With 48% still mostly or fully manual and manual error the heaviest deployment bottleneck, automating routine policy provisioning and change cuts the dominant source of misconfiguration and shortens the remediation window
- **Continuous compliance replaces periodic review with always-current evidence.** This is what addresses the audit findings and the multi-day evidence lag at once, and it aligns with the direction regulation is already taking

Read together, these four capabilities describe the upper stages of the maturity progression—the move from manual and reactive postures, where 61% of organizations sit today, toward the integrated operating model only 9% have reached. They also map onto a practical sequence any organization can use to locate itself—one where each capability depends on the one before it: establish the unified view, use it to analyze change before it ships, automate the routine work that view exposes, and let the evidence trail run continuously rather than on audit cycles. The organizations that advance will treat this as an operational redesign—consolidating visibility, clarifying ownership, and embedding policy into the deployment workflow—not as another line item layered on top of a fragmented model. In hybrid and multi-cloud environments, the limiting factor is no longer how hard teams work on policy, but how the work itself is designed. And the design that fits this environment starts from the applications the business runs on: from how they connect, how they change, and what they depend on to keep working.

Full Results

The Modern Environment

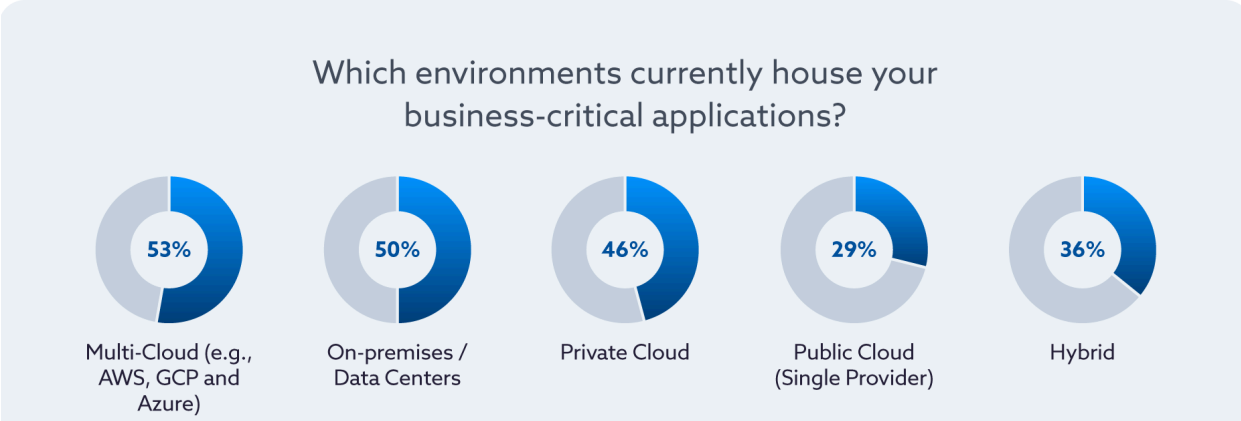

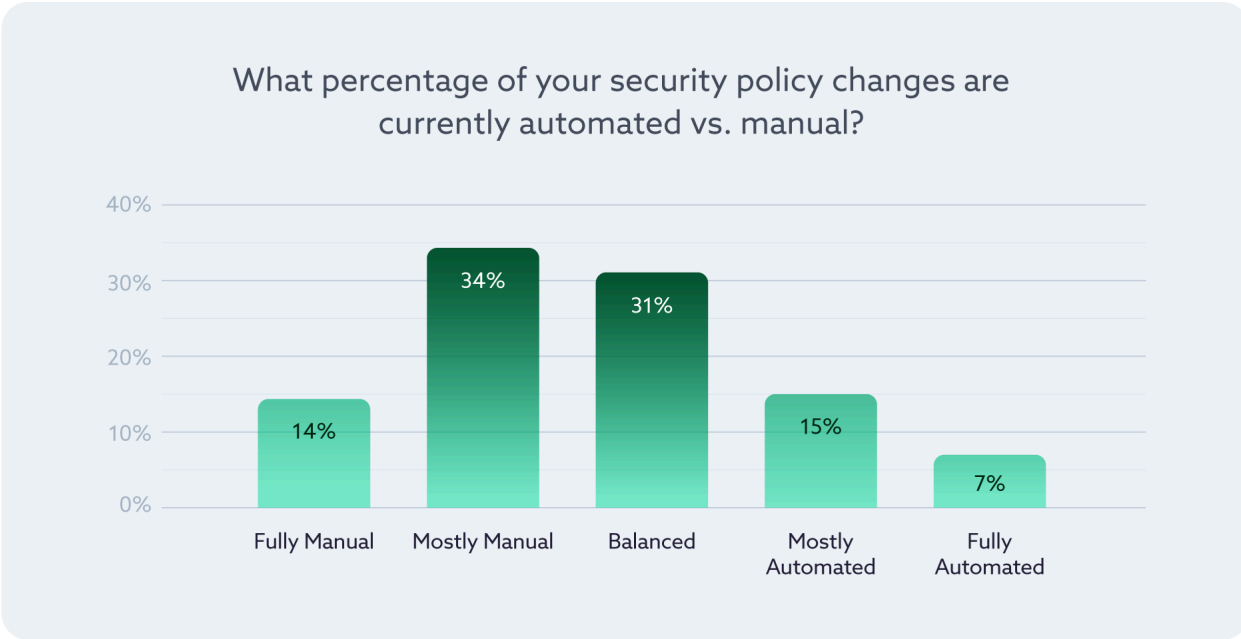

Which team(s) are responsible for defining application connectivity policies in your organization?

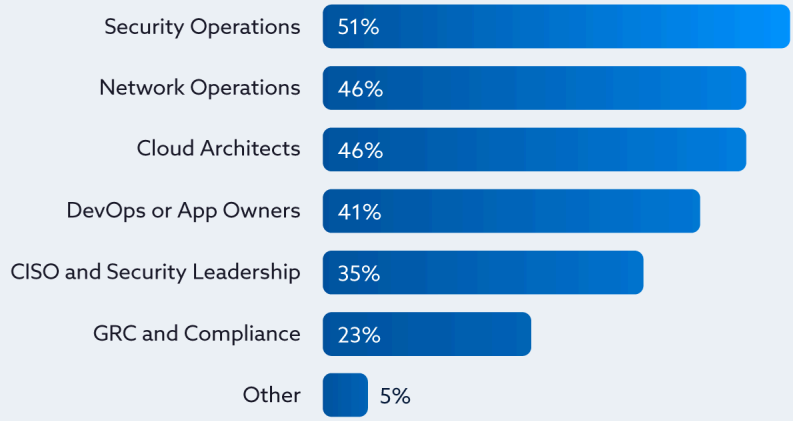

How many different security management consoles does your team use daily to manage connectivity?

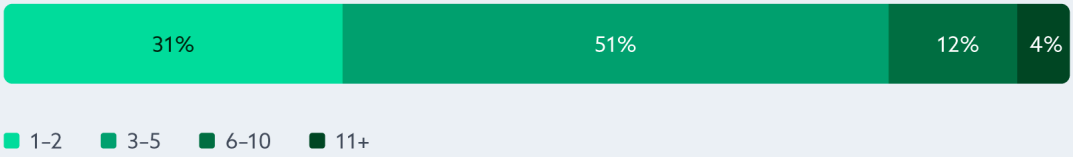

Pain Points

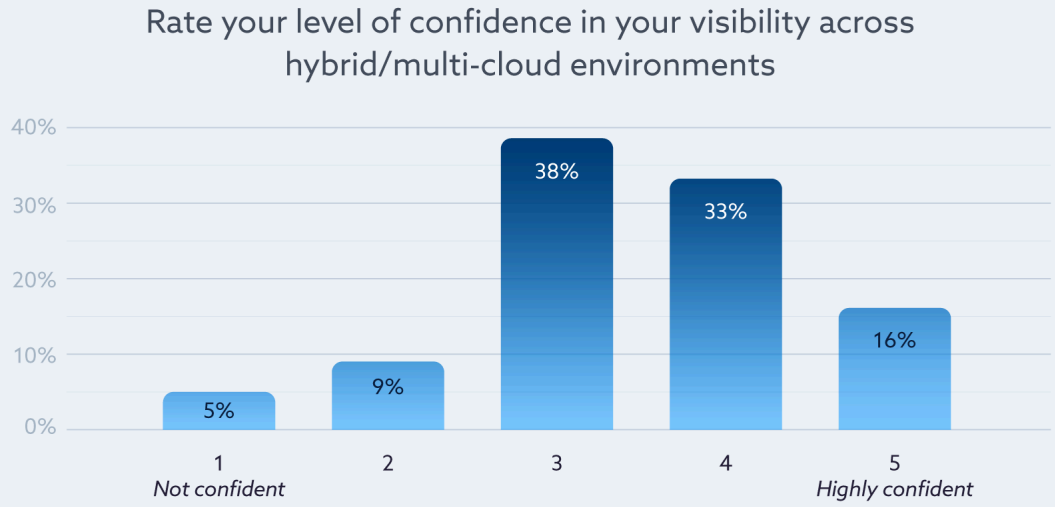

In the last 12 months, has a misconfigured security policy caused a business-critical application outage?

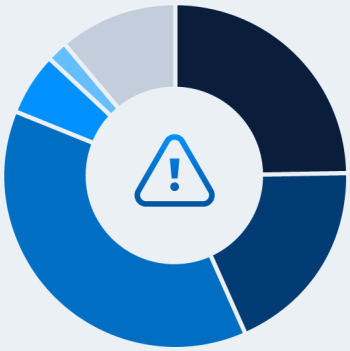

25%	Never	6%	6–10 times
19%	Once	2%	More than 10 times
38%	2–5 times	11%	Unsure

In the past 12 months, has a misconfigured security policy or policy gap caused any of the following?

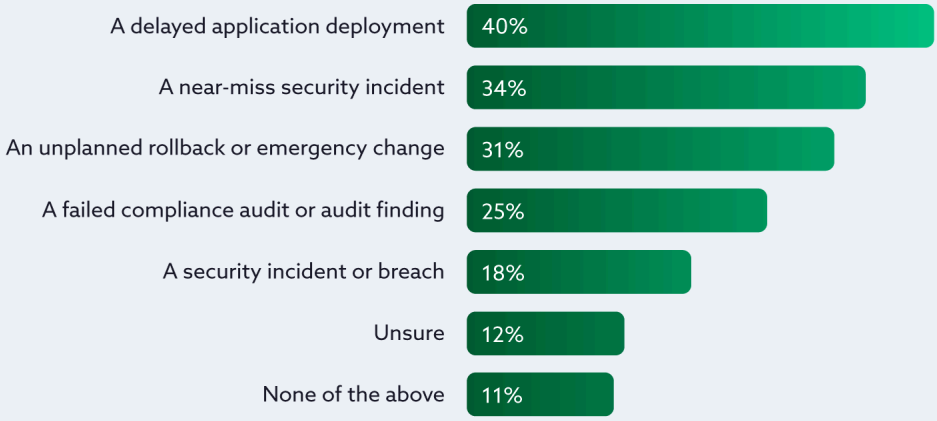

When a security policy issue is identified in production, how long does it typically take to remediate?

Strategic Priorities

Which of the following capabilities would most improve your organization's security policy management over the next 12 months?

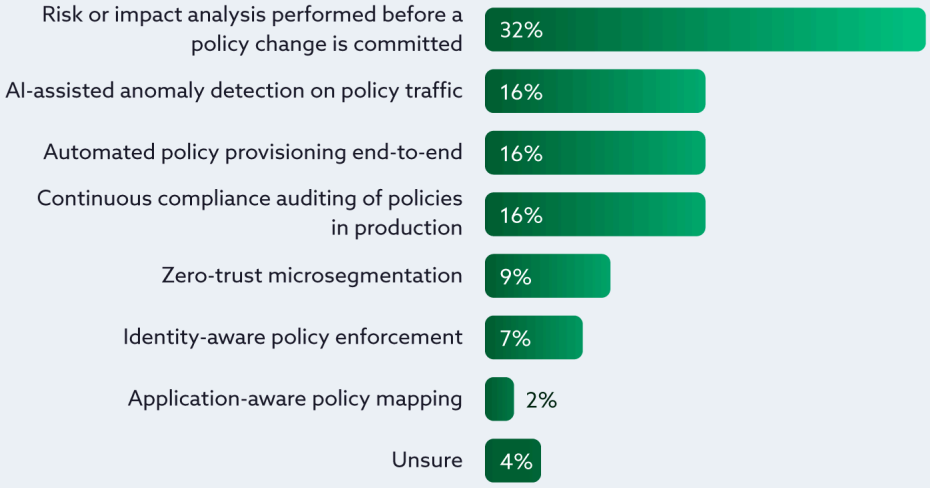

Which of these phrases best describes your current security posture?

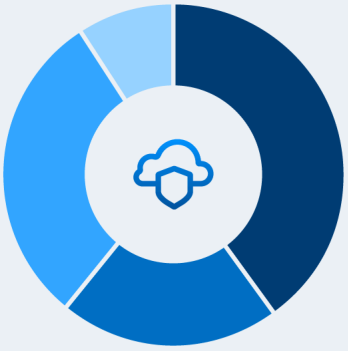

40% Stage 1 — Manual: Compliance is maintained primarily through manual processes and reviews

21% Stage 2 — Reactive: Security policy issues are addressed as they are identified or reported

30% Stage 3 — Proactive: Automated tools are used to detect and flag potential risks before incidents occur

9% Stage 4 — Integrated: Security policy management is fully integrated into development and deployment workflows

Demographics

Location

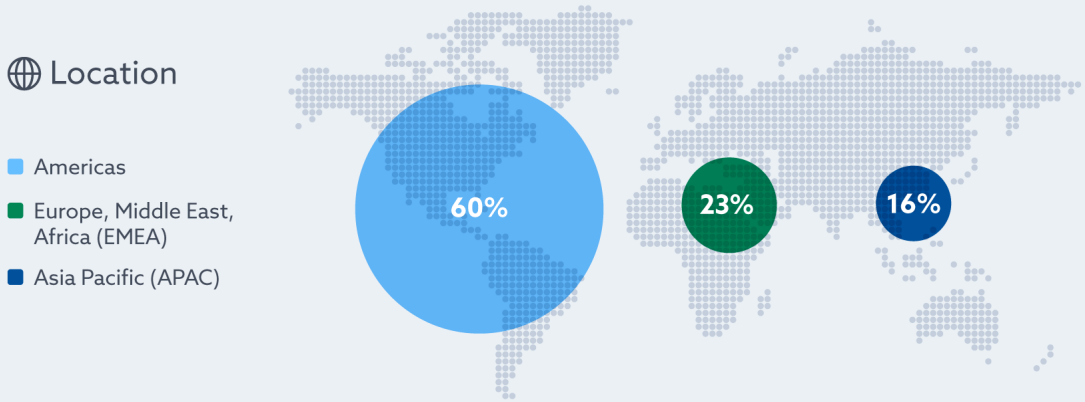

Organization Size

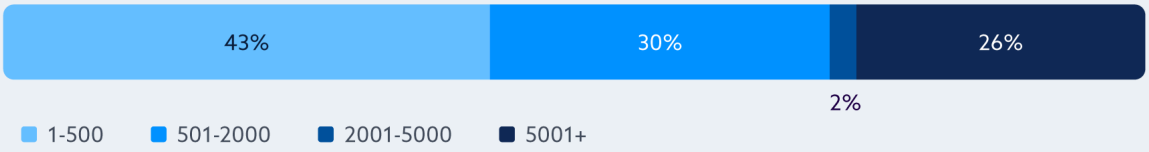

Job Level

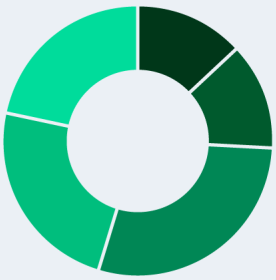

13% C-level or Executive
13% Director
29% Manager
24% Team Lead
22% Individual Contributor

Organization Industry

- 29% Telecommunications, Technology, Internet & Electronics
- 14% Finance & Financial Services
- 10% Education
- 8% Business Support & Logistics
- 6% Insurance
- 6% Entertainment & Leisure
- 6% Prefer not to answer
- 4% Healthcare & Pharmaceuticals
- 4% Manufacturing
- 4% I am currently not employed
- 2% Airlines & Aerospace
- 2% Utilities, Energy, and Extraction
- 2% Government
- 2% Automotive
- 2% Construction, Machinery, and Homes

Primary Functional Area of Role

- 18% Security Operations
- 12% IT Management
- 11% Security Architecture
- 11% GRC / Compliance
- 11% Cloud Architecture
- 9% Network Operations
- 7% Cloud Engineering
- 5% DevOps / Application Development
- 16% Other

Survey Creation and 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.

AlgoSec commissioned CSA to develop a survey and report to better understand the industry's knowledge, attitudes, and opinions regarding security policy management across hybrid and multi-cloud environments. AlgoSec financed the project and co-developed the questionnaire with CSA research analysts. CSA conducted the survey online in May 2026 and received 515 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 Survey

The goal of this survey is to increase awareness of the operational realities, challenges, and opportunities related to security policy management across hybrid and multi-cloud environments. The findings provide insight into how organizations are structuring policy ownership, the production and compliance costs of current approaches, and the capabilities they identify as most important for closing the gap between reactive practice and preventive intent.

In addition, this survey will help the industry:

- Understand the current operating model for security policy management across hybrid and multi-cloud environments, including team ownership, tool sprawl, and automation levels
- Assess the business and operational impact of policy misconfiguration on production uptime, deployment velocity, and compliance readiness
- Identify the structural and visibility gaps that prevent organizations from moving to more preventive postures
- Explore the capabilities, investments, and operating-model changes needed to align security policy management with modern application delivery