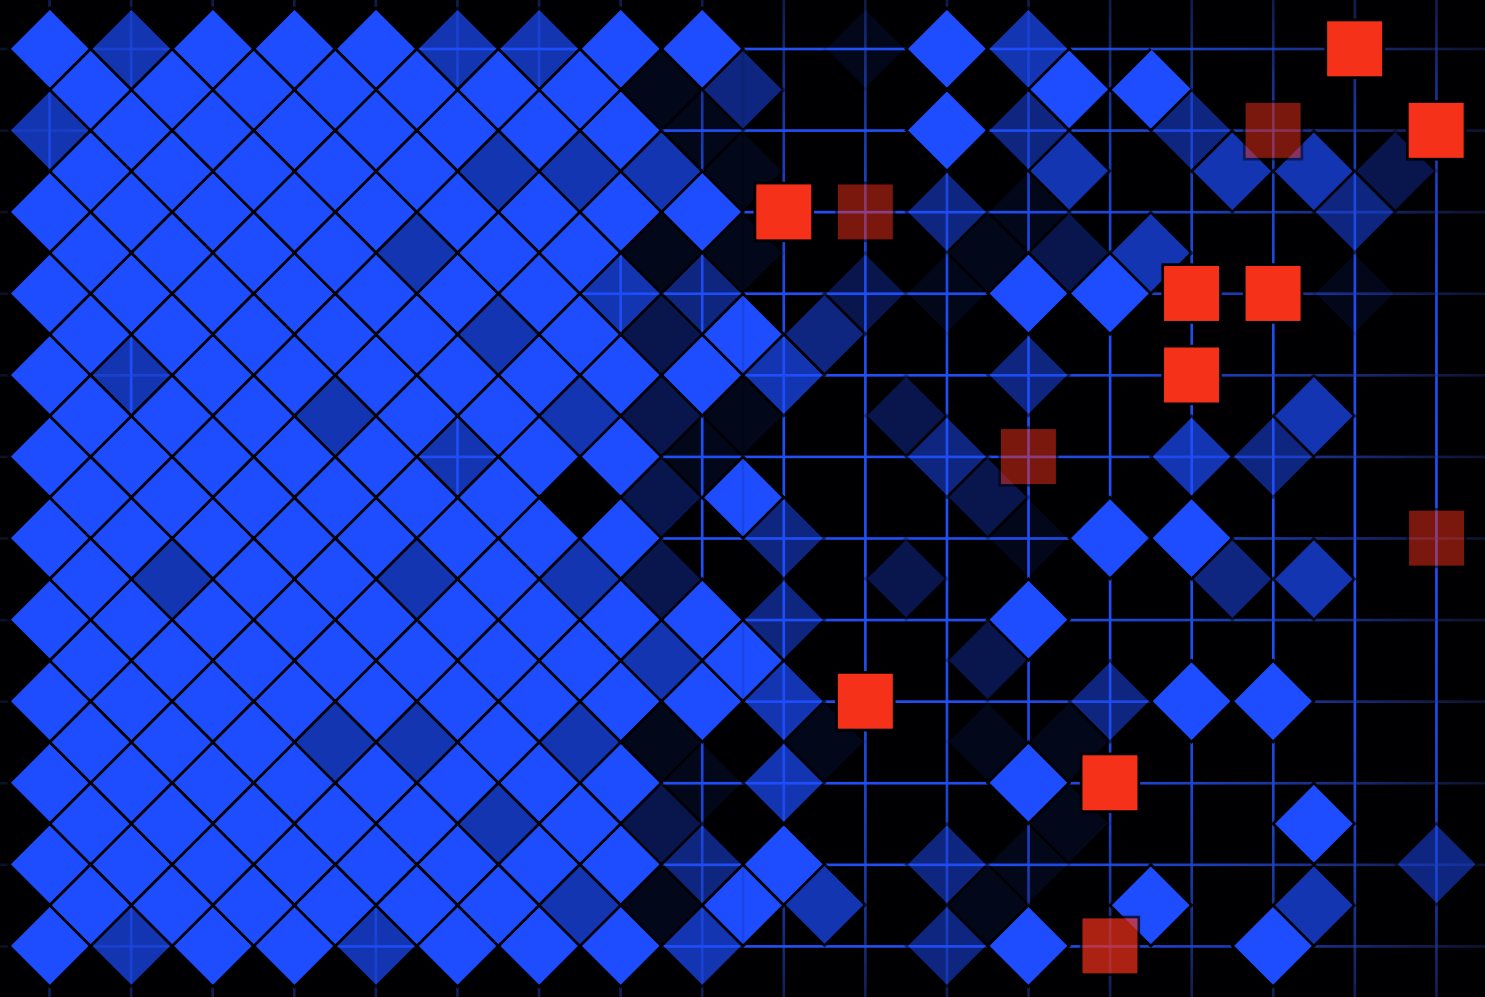

Reach

# Configure → Drift → Breach → Repeat

**Understanding the Cycle of Cybersecurity Control Configuration Risk**

A research report from Reach Security

# Introduction

There's an uncomfortable reality in today's cybersecurity landscape. While organizations battle to defend against malicious attacks by investing heavily in detection and incident response capabilities, they remain exposed to weaknesses that have crept into the first line of defense.

We are talking about blind spots: hidden misconfigurations, configuration drift, and unused capabilities in the organization's existing cybersecurity stack that would prevent, detect, or block attacks if properly optimized.

These blind spots are inevitable in every mature organization. Complex security architecture, business growth, frequent software updates, and system changes create an environment in which even correctly deployed tools rapidly – but silently – drift into a suboptimal state. The resulting exposures are difficult to identify and harder still to remediate, leading to invisible yet pervasive cyber risks that permeate the business.

To uncover the scale, impact, and causes of cybersecurity tool configuration risk, we commissioned a survey among US cybersecurity professionals across a range of sectors. The results are illuminating.

Read our research and analysis to gain deeper insight into the cybersecurity configuration drift cycle, its causes, management, and recommendations for tackling it in your organization.

# Executive Summary: Tool Sprawl, Breaches, and Slow Remediation

Organizations are struggling under the weight of sprawling cybersecurity environments, using an average of 35 separate tools. This complexity is directly fueling cyber breach risk: **97% of respondents reported either a confirmed breach or a near miss due to a cybersecurity tool misconfiguration in the past year.**

Despite this, our findings show that cybersecurity investment is weighted towards detection and response rather than prevention, reinforcing reactive approaches.

The practice of configuration management is still in its infancy. On average, **organizations review tool configuration 6.5 times per month**, but when issues are identified, **remediation takes an average of 8.3 days**, leaving ample time for security risk to accumulate.

Right now, configuration drift management is more of a compliance-focused exercise than a business facilitator. Most of the organizations we surveyed track adherence to baseline configurations rather than outcome-driven metrics such as reduced incidents or measurable drift reduction. However, today cybersecurity investments must deliver concrete benefits that go beyond a neat compliance checklist or completed audit; organizations need to know the tools they have invested in are performing optimally and actively preventing breaches and disruptions.

The root causes of configuration drift indicate the challenges of effective management: they arise from **operational overload, unclear ownership, and changes made outside formal governance processes**. Compounding this, we found **that manual processes remain a significant component of workflows used for identifying and managing configuration drift**, which is not a sustainable approach at the scale required.

Overall, the findings paint a picture of security teams battling inevitable blind spots in complex systems, with limited resources and a reliance on manual monitoring, resulting in slow remediation cycles and escalating exposure risk. And while organizations recognize the potential of AI-powered solutions to address the challenge, they are cautious about unleashing autonomous solutions on existing fragmented, low visibility environments.

# Contents

# The Configuration Risk Landscape

## 97%

had either a confirmed breach or a near miss as a result of a cybersecurity tool misconfiguration in the past year.

## 74%

have suffered a confirmed incident or breach

## 35

The average number of different cybersecurity tools in use

## 8.3 days

The average time taken to remediate an identified misconfiguration

## 6.5

The number of times per month organizations check for cybersecurity misconfigurations

## Barriers to controlling configuration drift

1. Changes made outside formal change control processes
2. Compliance or audit requirements expose drift faster than operations can respond
3. Too many tools and not enough people to manage them

# Configuration Risk Factors: Tool Sprawl Drives Drift, and Breaches Result

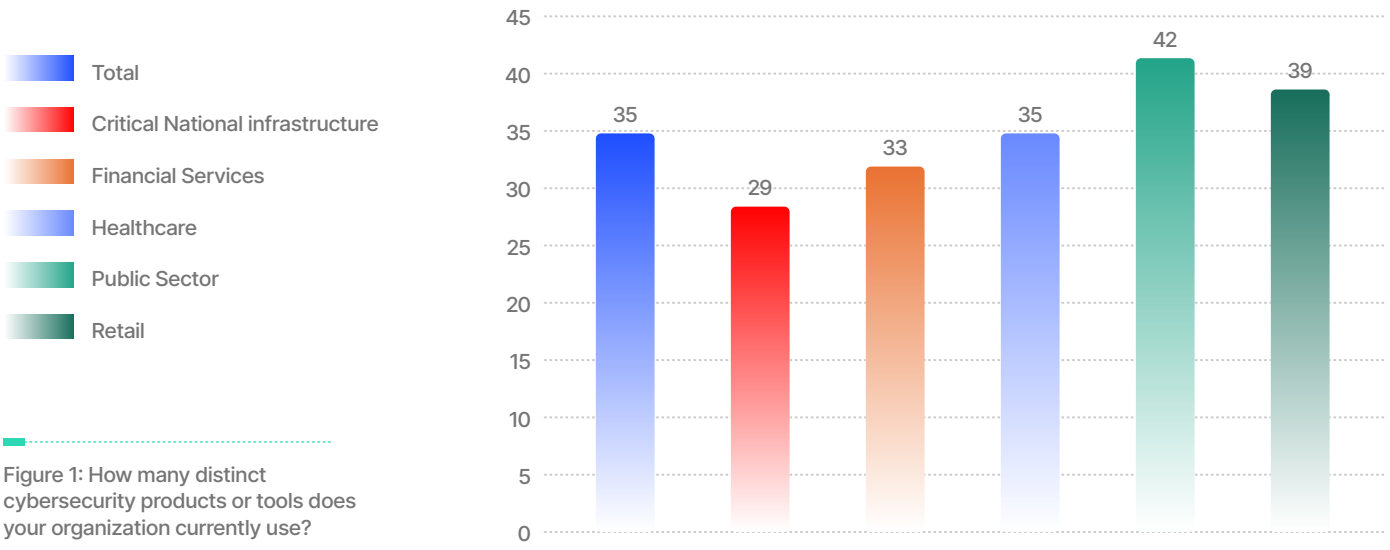


Figure 1: How many distinct cybersecurity products or tools does your organization currently use?

Our research reveals the scale and complexity of a typical cybersecurity tool stack. **Respondents use an average of 35 distinct cybersecurity tools**.

Even if each of those tools is initially deployed correctly, cybersecurity teams must maintain a regular patching and updating cadence to avoid it drifting from its optimal state. That's before you factor in the ripple effect of adding adjacent tools, implementing system changes, and those "temporary" exceptions and workarounds added on the fly to keep teams productive.

Our respondents reported their experience that all tools undergo some form of drift during their lifecycle. It's not surprising, therefore, that organizations cite "**too many tools and not enough people to handle them**" as a leading barrier to effective configuration drift management.

Analysis of the typical number of feature updates issued by vendors in a 12-month cycle validates this challenge. Reach Security has identified that popular security tools are updated on average 20 times every year. Organizations deploying an average of 35 tools—such as those we surveyed—must therefore expect to manage the addition of around 700 features across a diverse tool stack. Teams need to learn, understand, and identify how to effectively deploy and monitor all those features within their baseline security and control plan.

The result of this ongoing complexity and high administrative load is sobering: **97% of the organizations we surveyed have experienced either a breach or a "near miss" due to a misconfiguration of their existing security tools in the past year**. Seventy-four percent had suffered a confirmed breach.

The most common source of a configuration-related breach or near miss was the firewall, followed by endpoint detection and response tools and failures in identity/access policy controls.

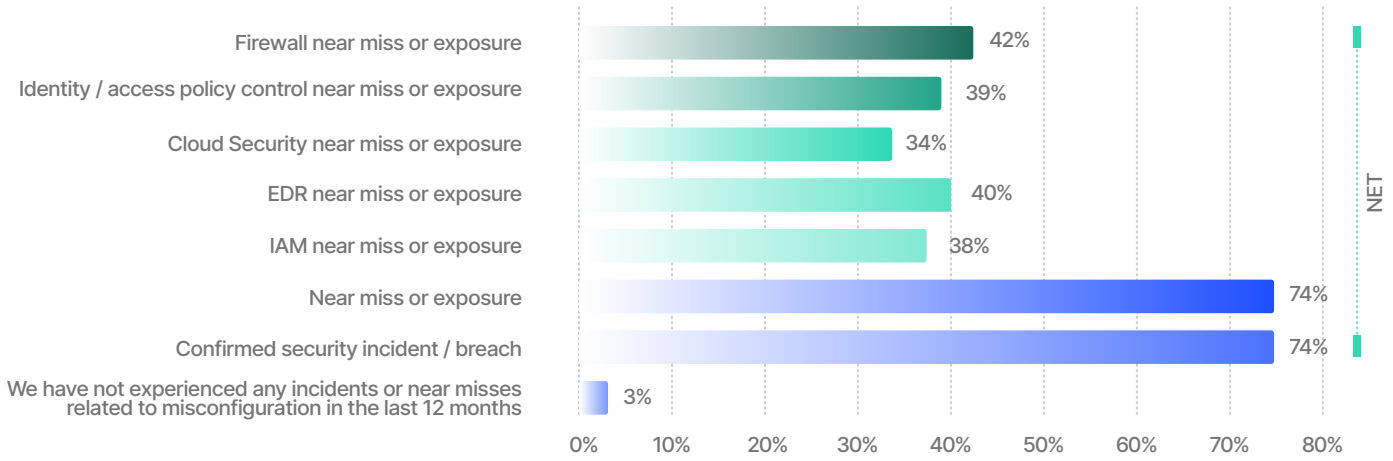


Figure 2: What, if any, security incidents or near misses, resulting from a misconfiguration of your existing security tools, has your organization experienced in the last 12 months? (Select all that apply)

> "We often say it only takes one bad click to create an incident, but with drift risk, it takes no clicks! Not having a program in place to identify and remediate configuration related risk is equivalent to leaving the front door unlocked and just hoping the bad guys don't open it."
>
> **Jay Wilson, CIO & CISO, Insurity**

# Cybersecurity Investment is Predominantly Reactive

To understand the impact of frequent breaches on cybersecurity investment, we asked our survey cohort how they typically spend their budget – left of boom, on tools that prevent attacks; center of boom, on tools that provide attack detection and response; or right of boom, on tools that assist in post-attack recovery. The answer, perhaps unsurprisingly given the acknowledged fragility of cybersecurity infrastructure, is that the lion's share of expenditure is devoted to tools at the center and right of boom. These account for 72% of expenditure, compared to just 28% that is allocated to preventive, left-of-boom tools.[1]

There's a strong economic argument here. If attacks can be stopped in their tracks by optimizing existing tools and maintaining them in a high-performing state, the downstream financial benefit is amplified. The organization suffers fewer breaches, avoiding all the financial, operational, regulatory, and reputational costs they entail.

By prioritizing prevention through more effective configuration management, long-term risk can be reduced. However, there's a risk window in shifting finite resources towards a preventive approach; the transition must be rapid and highly effective to avoid exposures being exploited. Our findings show that current approaches to configuration management are not robust enough to provide a solid foundation for a "prevention-first" strategy.

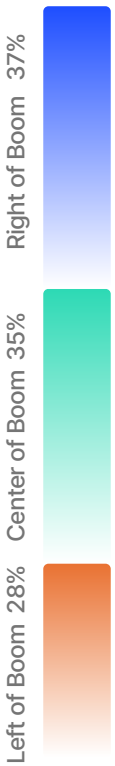


[1] Combines 'Right of boom (i.e. tools / products that help us recover from attacks)' and 'Center of boom (i.e. tools / products that help us detect and respond to attacks)'

# Configuration Management Today: a Nascent Discipline

## A window on review and remediation

Most organizations review cybersecurity tool configurations regularly, but not continuously, resulting in an average cadence of **6.5 reviews per month**. Only one in twenty organizations reviews tool configurations every day.

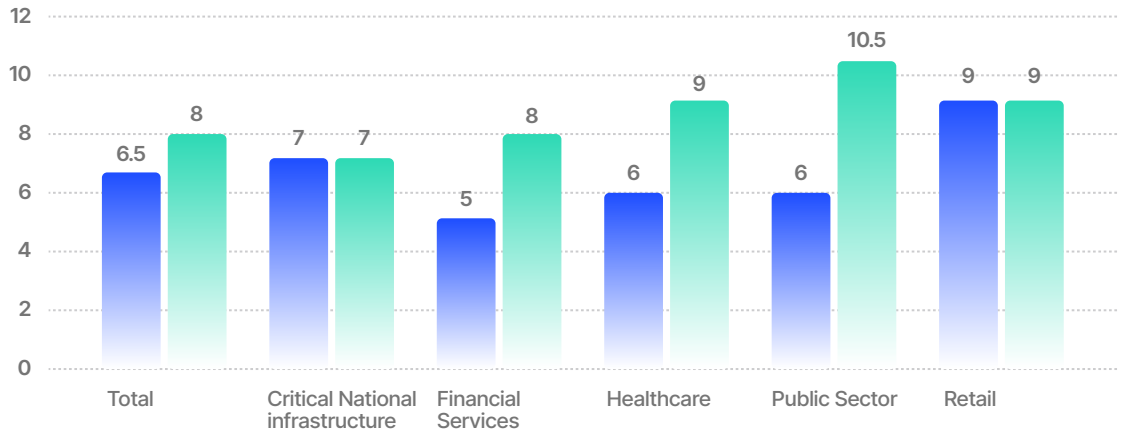


Figure 3: How often, if at all, do you review your cybersecurity tool configurations to ensure they are optimized to prevent attacks and remain aligned with policy and compliance requirements? In your organization how long, on average, does it take to remediate a misconfiguration or incident of drift once it has been identified?

Of equal importance is how quickly teams can remediate misconfigurations and drift incidents once identified. Our research reveals a significant exposure gap, **with organizations taking an average of 8 days to resolve issues**. Just 2% of respondents say they can fix issues in less than a day.[2]

Concerningly, our study uncovered some evidence that even when organizations are faster than average at remediating misconfigurations and vulnerabilities, this doesn't automatically translate to a reduction in breaches.

This points to a prioritization problem: teams are responding to the “noisiest” or most frequent alerts, but not necessarily the most material and high-risk issues. Identification without prioritization diverts resources away from where they are needed. There’s also a performance measurement component here: if workers are monitored on activity volume, rather than their actual impact on risk reduction, they may prioritize quicker, easier fixes over tackling more complex problems.

[2] Mean: (No. of days excl. "My organization doesn't have a consistent remediation timeframe", "My organization hasn't remediated a misconfiguration or incident of drift")

# Drift detection: manual point-in-time approaches prevail

Organizations reported using a range of tools and processes to assess whether cybersecurity tools have drifted from their optimal intended or compliant configuration and created exposure. The most common approach is having a dedicated team or person responsible for configuration management, mentioned by more than one-quarter of respondents.

Compliance-driven approaches are also common: 26% of respondents rely on external audits or certification processes (such as PCI-DSS, SOC2, and ISO) to identify drift, 24% use penetration tests to identify configuration weakness, and 23% say internal audit or internal control testing identifies drift. Less frequently, organizations use spreadsheets or checklists (18%), and manual configuration reviews (14%).

Overall, the balance is weighted toward point-in-time manual review and testing procedures, which align with the review cadence reported earlier. However, there is a definite appetite for automated tools: 24% report using external scanning or posture management tools, 23% rely on automated configuration monitoring within existing security tools, and 21% have developed custom internal automation or scripts.  This points to recognition that manual approaches will become increasingly inadequate as the size and complexity of systems grows. They are high cost in terms of time and financial commitment, and they cannot scale efficiently. But while organizations are exploring automation, this approach is still in its infancy. The use of proprietary customized internal automation indicates that teams are trying to solve the issue independently, but such approaches can become problematic in the long-term if they don't integrate with the security ecosystem.

According to our respondents, the tools most commonly drifting from their optimal configuration are those related to identity, access, cloud, and endpoint management. The pressure points lie where environments are at their most dynamic, where configuration requires regular, fine-grained tuning. Today's distributed workforce is also having an impact, as perimeters are frequently recalibrated, with VPNs/remote access tools cited as problematic by 18% of respondents.

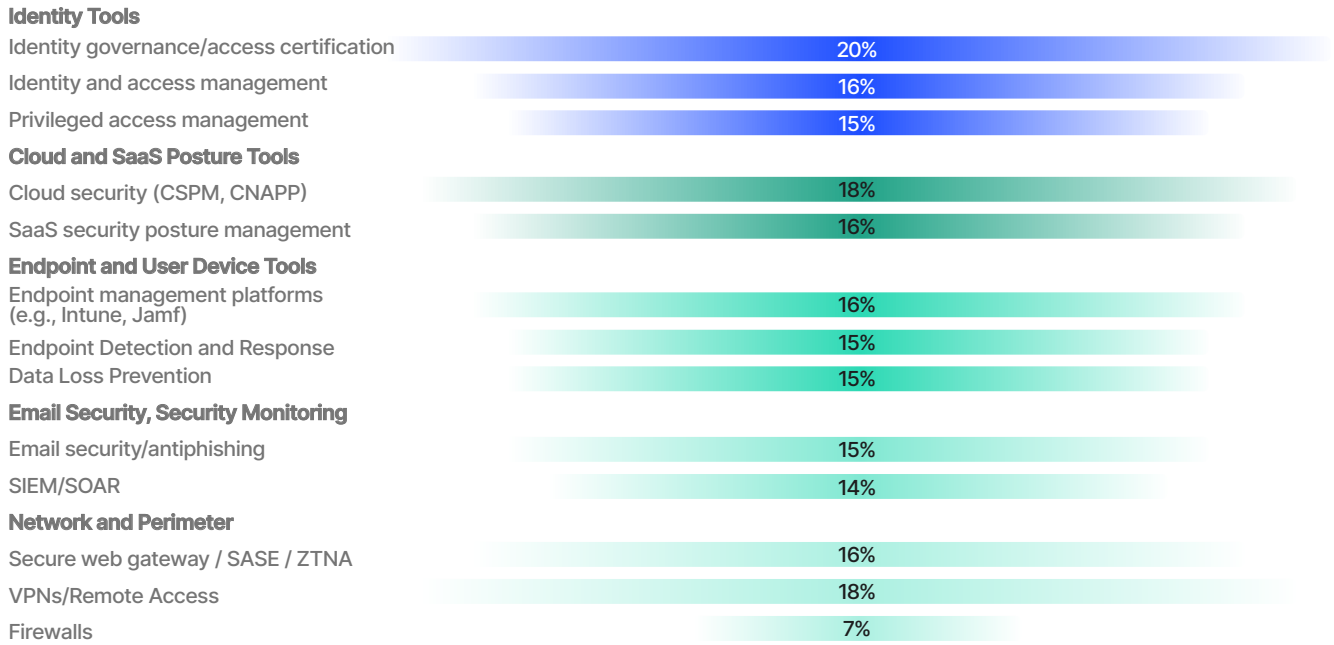


Figure 4: In your experience which, if any, tools suffer the most drift from their optimal configuration? (Select up to 3)

Endpoint and user-device ecosystems are also prone to drift due to device heterogeneity, frequent updates, and users who may circumvent controls for convenience. Cloud and SaaS posture tools drift because the environment often changes faster than controls can keep up.

## Measuring configuration exposure management

Respondents were most likely to say that they measure the success of their exposure management program by assessing the percentage of controls that meet baseline configurations – one in four respondents use this method.

Many of the top metrics focus on visibility, baseline adherence, and coverage, rather than on the reduction in incidents or drift. This suggests organizations are still early in their configuration management maturity curve and see it as more of a hygiene and compliance exercise than a business facilitator. However, in today's commercial landscape, cybersecurity investment must align with business outcomes.  Boards and business leaders expect cybersecurity professionals to tie security initiatives to measurable business objectives – revenue protection, operational resilience, speed to market, and risk-adjusted growth – making cybersecurity a demonstrable business enabler rather than a cost center. Consequently, metrics need to be transparent and relatable for business audiences. Reductions in incidents linked to misconfiguration and cutting the time to detect and remediate drift are easier to link to business resilience.

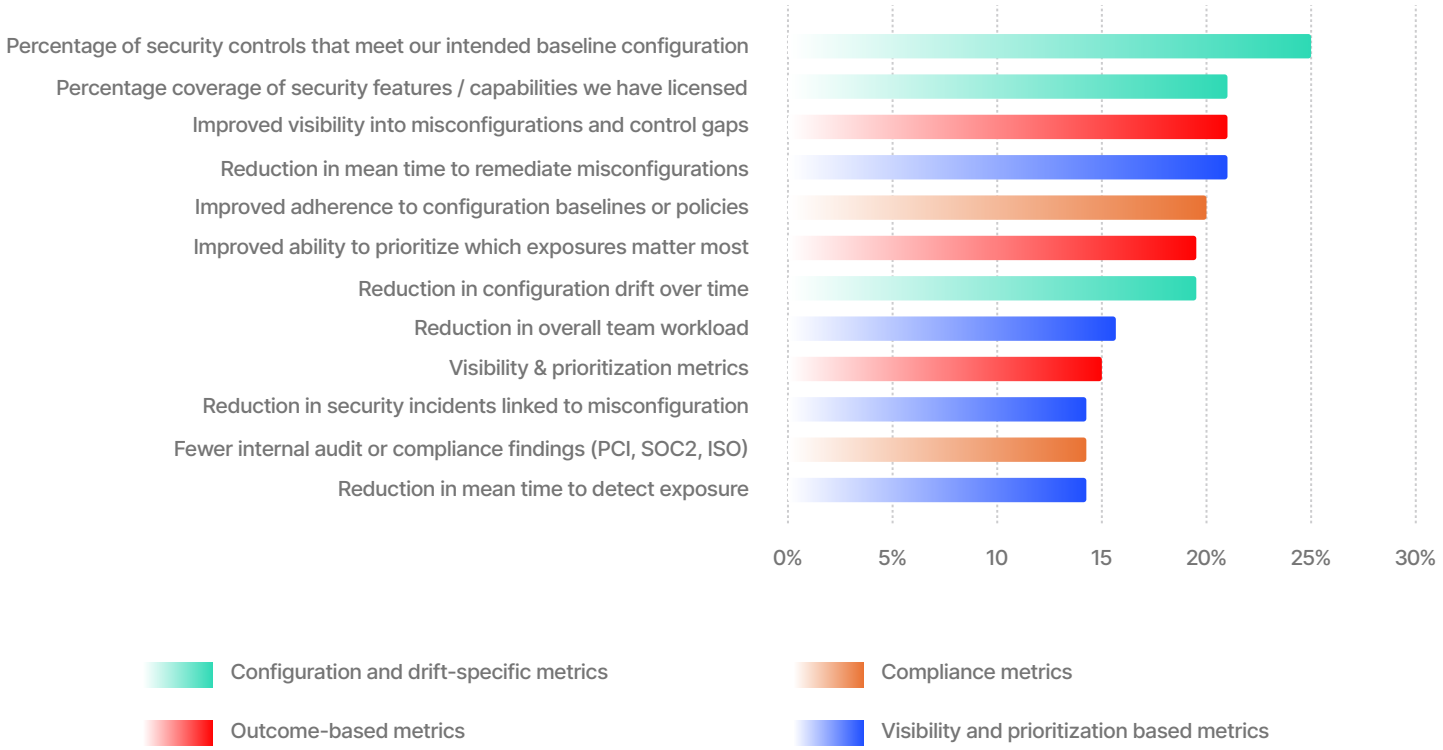


Figure 5: What, if any, key metrics does your organization use to measure the effectiveness of your exposure management program? (Select up to three)

# What's blocking better configuration drift management?

The core challenge of tool sprawl is a significant blocker to configuration drift management, with one in five respondents saying they have "too many tools and not enough people to manage them". However, the leading barrier is a lack of rigorous governance around configuration changes; 21% of respondents say that "configuration changes made outside formal change control processes" is a top-three barrier. A similar proportion says that compliance or audit requirements expose drift faster than operations can respond.

Visibility is another problem. Lack of real-time or accurate configuration data and unclear or outdated configuration baselines for drift measurement are equally challenging, with 18% putting them in the top three problems.

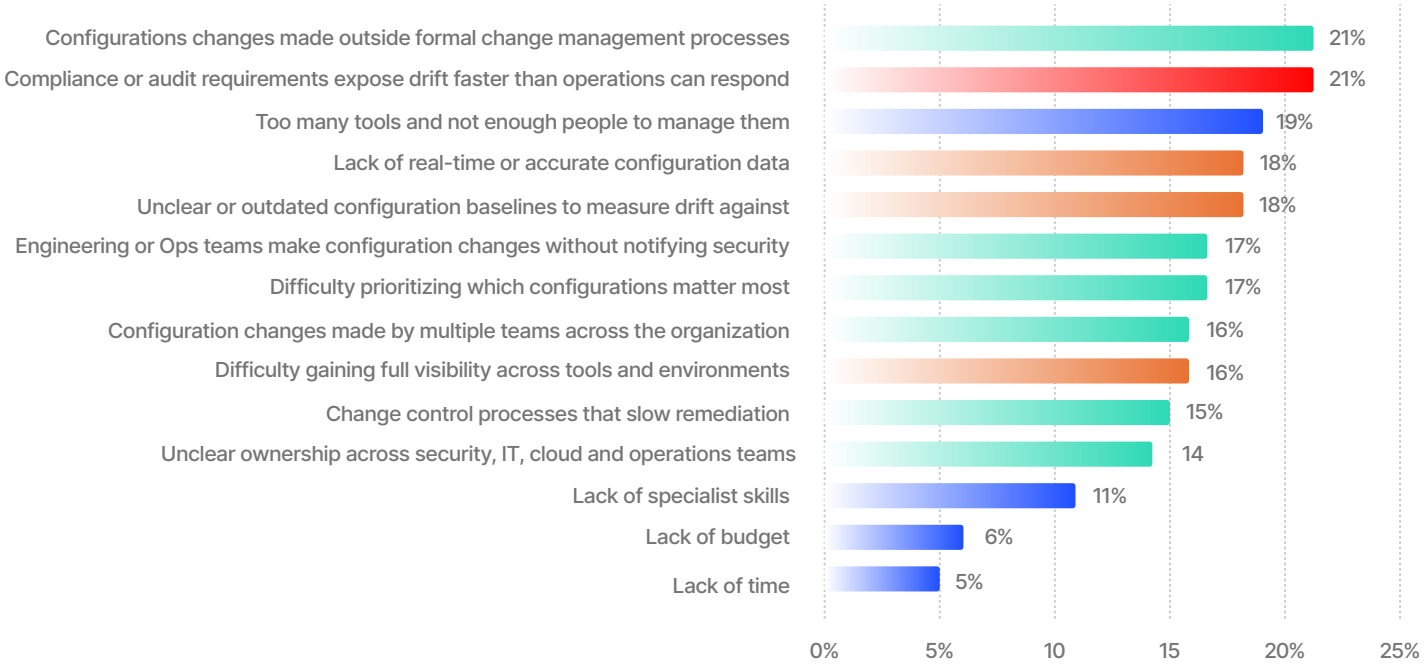


> "Often the challenge in managing configuration drift has been a signal to noise issue. I have lots of tools telling my teams that there are exploits to be remediated, but before Reach, nothing was telling me my own configuration and settings were just as risky as an exploitable CVE."
>
> **Jay Wilson, CIO & CISO, Insurity**

Figure 6: What, if any, are the main obstacles to your organization better controlling configuration drift? (Select up to three)

The wide range of barriers to better managing configuration drift cited by respondents offers further insight into the scale of the issue. Solutions that cost-effectively surface key data and provide better visibility, while also streamlining governance and ownership, will have an immediate impact on risk reduction and ongoing cybersecurity posture.

# AI in Configuration Drift: Help and Hindrance

Introducing AI into a constantly changing, highly configurable security stack increases both capability and volatility. AI is an extraordinary accelerator, but when AI agents can reconfigure controls in real time, they don't just increase operational speed, they increase the rate of configuration drift. What was gradual becomes exponential. That makes drift management foundational. And if AI contributes to the problem, it must also be central to the solution.

However, despite the potential for AI-powered tools to address drift at scale and pace, our survey cohort has reservations about unleashing AI agents to detect and remediate configuration drift. The top concerns reflect their experience to date; they fear that AI will introduce new forms of drift or configuration inconsistencies and that there'll be a lack of transparency and explainability into why the AI recommends or performs an action. They're also concerned about ensuring the appropriate oversight, approvals and guardrails, and the time and resources needed to deploy an AI-based system.

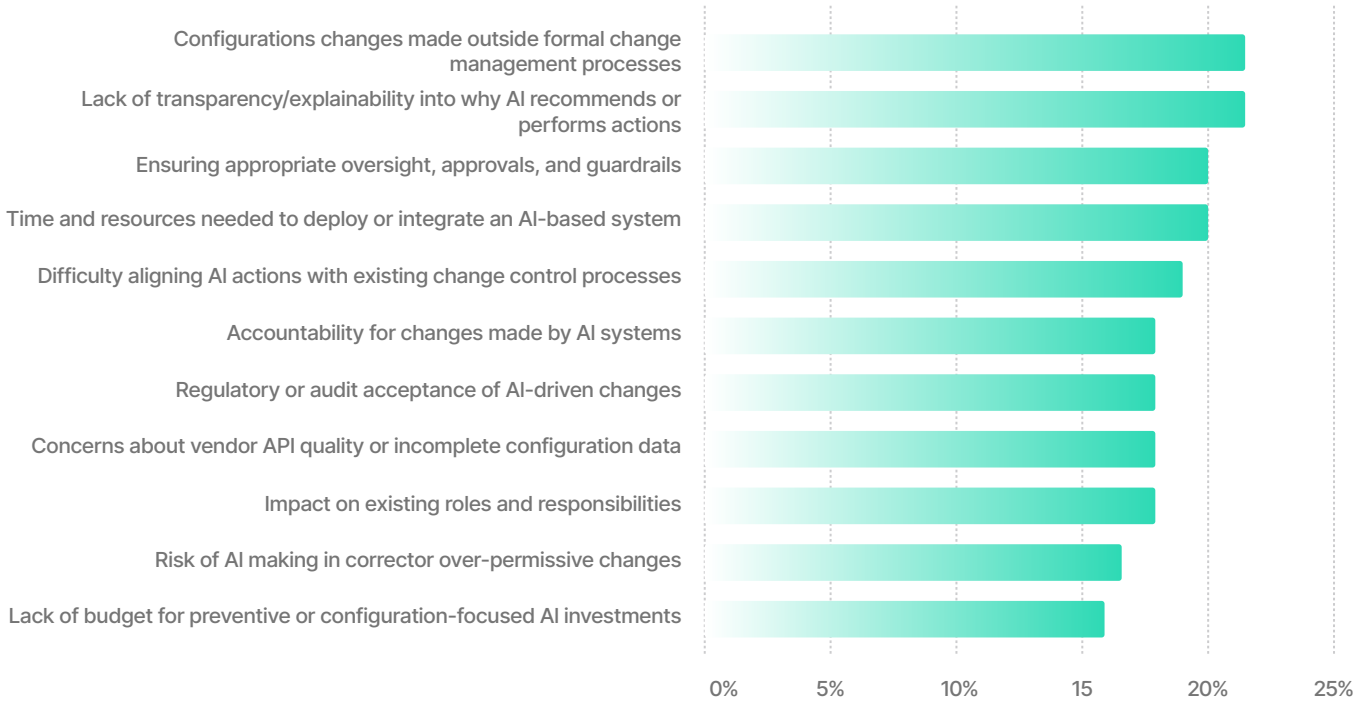


Figure 7: What, if any, are the main concerns you have about using AI agents to detect and remediate configuration drift? (Select up to 3)

The breadth of concerns our respondents cite about introducing AI for configuration drift management underlines the careful approach required when designing a solution that is effective, transparent, straightforward to implement, and - above all - trustworthy.

# Building trust in AI-powered cybersecurity

Trust is only established when users are confident in the quality and accuracy of AI reasoning, and solutions built on generic LLMs don't meet the standard required for cybersecurity use cases. Hallucinations and out-of-context decision-making elevate risk and compromise trust.

Reach Security has addressed this problem by building a domain-specific language model (DSLM) trained on cybersecurity-specific data sources, including vendor APIs, configuration baselines, threat intelligence, and frameworks such as MITRE and NIST. This delivers precise, explainable assessments of control coverage, configuration state and exposure impact across identity, endpoint, network, email and cloud controls.

Control is also crucial when AI tools are deployed. Reach's MastermindAI™ identifies configuration drift blind spots, prioritizes fixes, and - when granted full authorization - autonomously remediates misconfigurations and inactive controls. However, the level of agency delegated to AI can be varied. Instead of full automation, the AI can raise a ticket listing the issues identified for a human to action, or stage an automated change that requires human review prior to implementation. This provides a level of human oversight that the organization is comfortable with, and builds trust over time.

# Conclusion:
# Breaking the Configuration Drift Cycle

Our findings are clear: **97% of organizations have suffered breaches and/or near misses as a result of cybersecurity configuration drift**, making it a dominant risk driver in corporate environments. Tool sprawl, governance challenges, and slow identification and remediation cycles have created an environment where control is significantly compromised. As a result, companies are prioritizing investment in reactive threat detection and incident response solutions and are cautious about shifting focus to proactive, preventive approaches. Adding AI to the mix is accelerating drift risk and reinforcing reactive thinking. The configure, drift, breach, repeat cycle is firmly embedded as a systemic pattern across all organizations.

Breaking this cycle is possible and essential if organizations are to reduce risk to acceptable levels and achieve a level of resilience that enables the business.

## At the same time, our research highlights two deeper structural issues that must be acknowledged:

### Configuration change is creating risk faster than security teams can keep up

Most organizations are still using processes and assumptions built for a slower, more predictable technology environment. Today's hybrid, multi-cloud, multi-vendor reality produces configuration changes at a scale and velocity that manual processes and legacy operating models simply cannot absorb. This mismatch creates a systemic vulnerability: even well-resourced teams are unintentionally operating with outdated models of control and visibility. Until organizations modernize the way they manage configuration integrity, they will continue to fall behind.

### Configuration drift is eroding trust in security controls and increasing costs

As drift undermines the reliability of controls, organizations compensate with layers of redundant tooling, manual verification, and reactive monitoring. This increases cost and complexity while widening the "assurance gap" between perceived and actual control effectiveness. Without continuous, automated validation, leaders cannot reliably assess whether their security posture is holding - and teams remain trapped in a cycle of overspending on reactive measures without achieving meaningful risk reduction.

# Key recommendations

**Pursue continuous visibility and validation:** Drift is inevitable and continuous, and blind spots are no one's fault. Even the most conscientious and well-resourced teams can't keep pace manually across the breadth and depth of tools in place - nor is it cost viable to do so. Look for solutions that offer continuous posture assessment and drift identification to deliver visibility at the level and cost required in fast-changing organizations.

**Prioritize remediation to achieve material risk reduction:** Our survey shows that even organizations that are quick to address drift don't necessarily see material risk reduction. The volume of alerts from multiple tools can funnel teams down the path of least resistance, leading them to opt for quick, easy wins rather than fixing the highest-risk issues. By implementing a solution that surfaces the most important problems, the organization is better protected.

**Close the gap between drift identification and fix action:** Knowing there is an exposure is not the same as fixing it. Look for a solution that not only surfaces drift but also helps remediate it with specific fix guidance that cuts the time between identification and resolution. Choose a solution that integrates with ticketing systems and can automatically stage a fix for human review, helping teams work faster, even when they lack expertise in the tool in question.

**Choose AI-powered automation solutions with strong subject-matter expertise and human-in-the-loop facility:** Look for solutions that draw on domain expertise and provide automated workflows to resolve issues with an appropriate degree of human oversight where required.

**Quantify drift risk in terms of business outcomes:** Cybersecurity is a business enabler and therefore a board-level issue. Configuration management should be framed as a pillar of business resilience underpinning the ability to operate. By framing board reporting in terms of incident reductions and costs avoided, cybersecurity leaders will resonate with corporate leaders and help justify a preventive investment approach.

# Methodology

The research was conducted by Opinion Matters, among a sample of 250 Cybersecurity Professionals (aged 25+) in the US working in companies employing 2000+ people in Financial Services, Retail, Public Sector, Healthcare and Critical National Infrastructure sectors. Natural fallout within those sectors. The data was collected between Dec 12, 2025 - Dec 22, 2025. Opinion Matters abides by and employs members of the Market Research Society and follows the MRS code of conduct and ESOMAR principles. Opinion Matters is also a member of the British Polling Council.

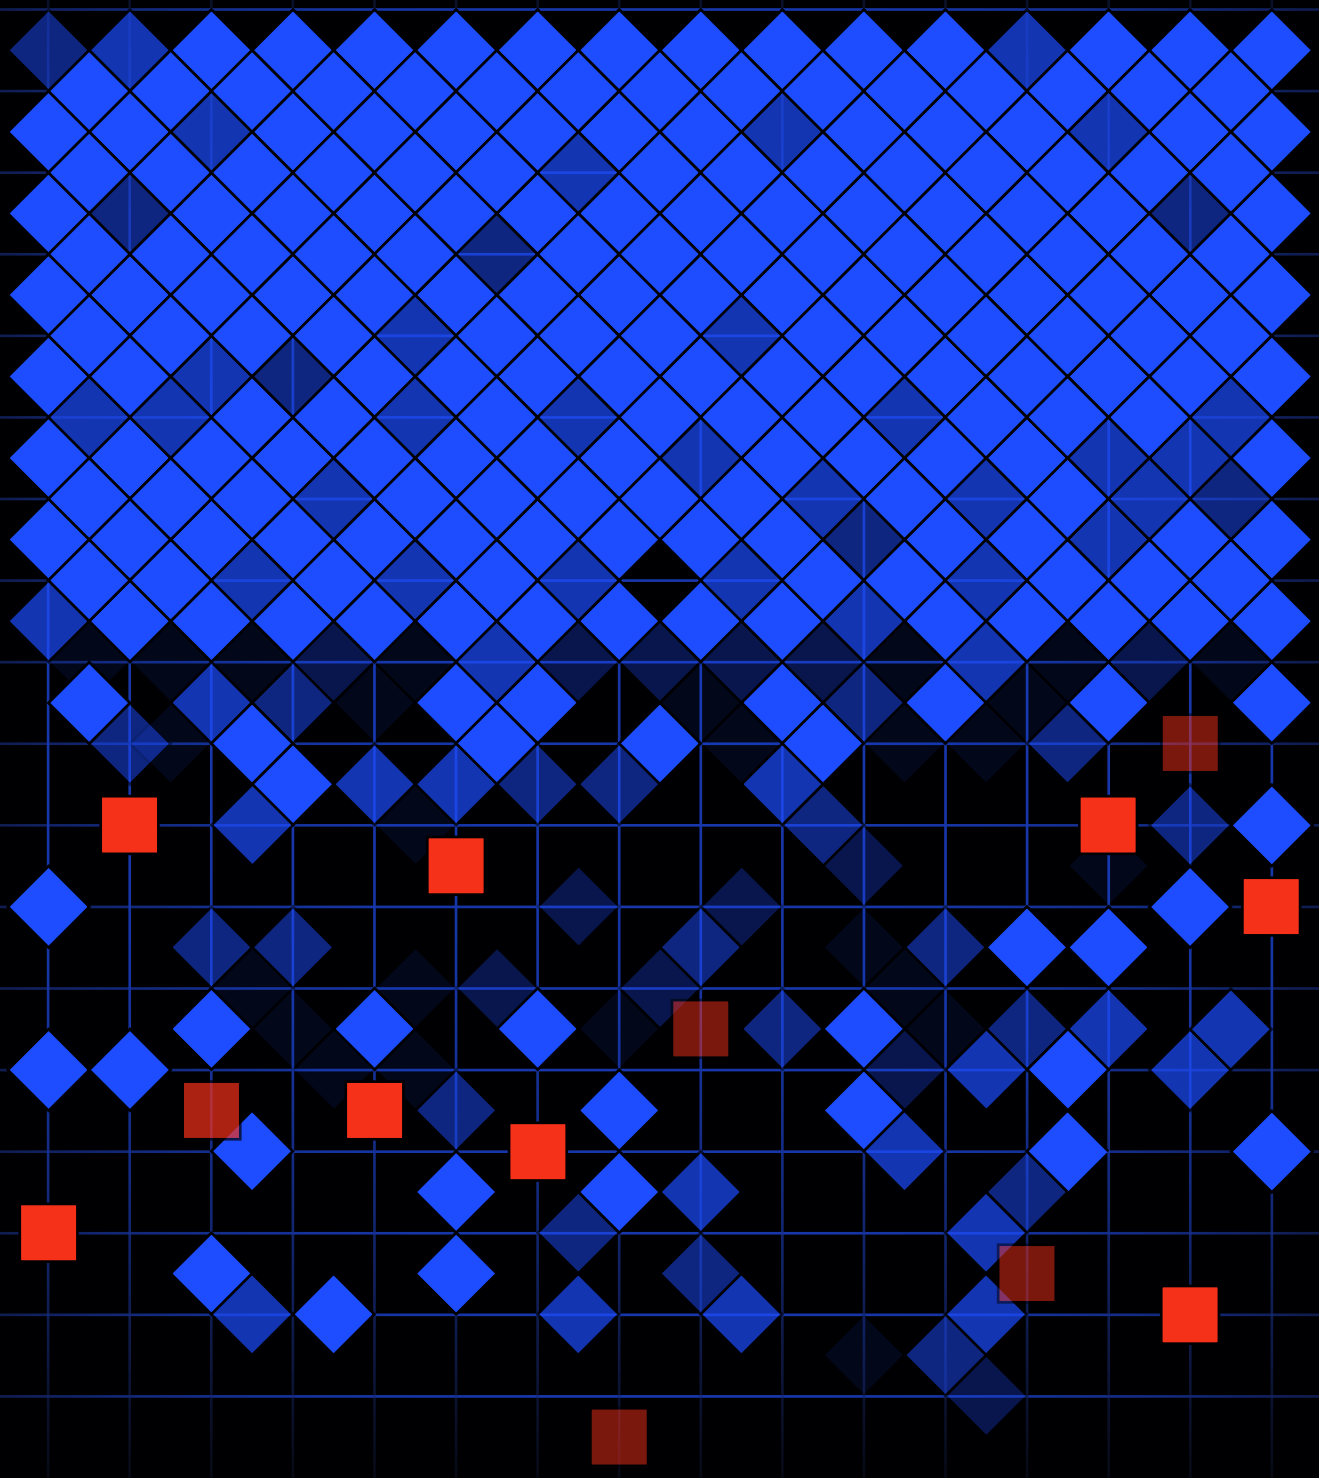

## About Reach Security

Learn more about Reach Security, an AI-native operating system for your security controls, purpose-built to reduce exposure by improving how your existing tools are used. It continuously analyzes how your security controls are configured and operating to identify where protection is breaking down. Using domain-specific AI, Reach understands how controls should work, detects drift and gaps, and pinpoints the highest-impact actions to reduce exposure.

www.reach.security

Reach