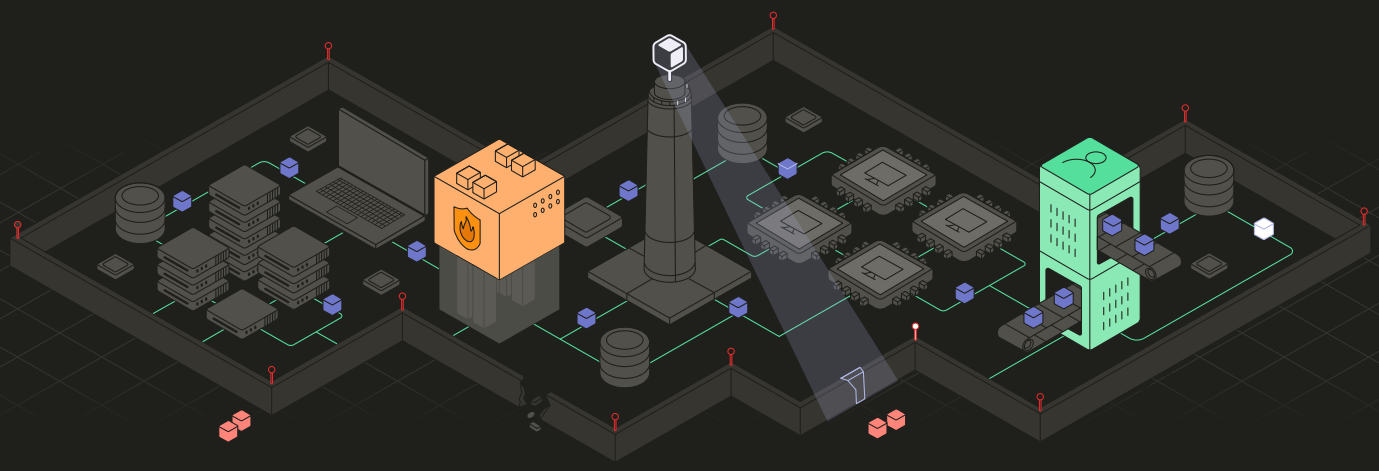
Reach

REACH TELEMETRY DATA REPORT AUGUST 2026

SECURITY INTENT VS. SECURITY REALITY:

Configuration Drift in the Age of AI

What a Year of Production Telemetry Reveals About Configuration Drift and Hidden Exposure



To uncover the scale, impact and underlying causes of cybersecurity tool configuration risk, Reach Security commissioned an independent survey in April 2026 involving 250 cybersecurity professionals across the US, representing a broad range of industries. The research explored respondents' experiences with configuration drift, its contributing factors, and the strategies being adopted to detect and manage the associated risks.

To complement and validate these insights, Reach Security conducted a telemetry-based analysis in July 2026 using an anonymized 12-month dataset spanning more than 50 production environments and millions of security control domains. While the independent survey captured how cybersecurity leaders perceive and address configuration drift, this telemetry report provides a real-time view of conditions in live production environments. Covering a full year of operational activity, the data in this report delivers a ground-truth assessment of configuration drift, exposure risk and security outcomes across customer deployments.

All rights reserved. Reach Security 2026.

Executive Summary

The independent survey revealed the scale of the challenge facing organizations as they struggle to manage increasingly complex cybersecurity environments. This complexity is translating directly into risk: 97% of respondents reported experiencing either a confirmed breach or a near miss caused by a cybersecurity tool misconfiguration within the past 12 months.

Despite this, the findings suggest that cybersecurity investment remains heavily weighted towards detection and response rather than prevention, reinforcing a largely reactive approach to risk management. Configuration assurance, by comparison, remains relatively immature. Organizations review security tool configurations an average of 6.5 times per month, yet when issues are identified, remediation takes an average of 8.3 days. This creates a significant window during which risk can accumulate and exposures can remain unaddressed. The urgency of this gap has grown sharply. AI-powered adversaries can now probe environments, identify misconfigurations, and move laterally at machine speed. Meanwhile, the control weaknesses that enable them to persist for days remain. That mismatch — between how fast attackers move and how slowly defenders find and fix the exposures they exploit — is the defining risk of this era.

Reach's telemetry data adds essential context to this picture. Across a robust sample of production environments, configuration drift was observed daily. The average organization generated 13 drift alerts per day, with 12 linked to genuine, risk-prioritized security exposures. These are not low-priority hygiene issues or false positives; they represent meaningful gaps between how security controls are configured and how they should be configured based on the organization’s actual threat profile.

The data also highlights that network security controls produced the highest concentration of drift alerts, with firewall technologies accounting for the largest share of exposure. This suggests that the pace of operational change has outstripped the ability of traditional assurance processes to keep configurations aligned with intended security policies. Collectively, the telemetry findings point to three important realities: configuration drift is continuous, not sporadic; materiality is highly contextual, meaning not every change introduces meaningful risk; and security exposure is concentrated in a relatively small number of operationally complex controls.

The challenge for security teams is therefore no longer simply identifying configuration changes, but rapidly determining which changes create genuine exposure and require immediate action. In an era where AI is enabling threat actors to operate at unprecedented speed and scale, organizations cannot rely on periodic reviews and reactive processes. Effective cyber resilience now depends on continuous configuration assurance, risk-based prioritization and the ability to respond to drift as quickly as it occurs.

SECTION 1

What the Data Reveals about Drift Risk

Independent Research

97%

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

8.3 days

The average time taken to remediate an identified misconfiguration

6.5

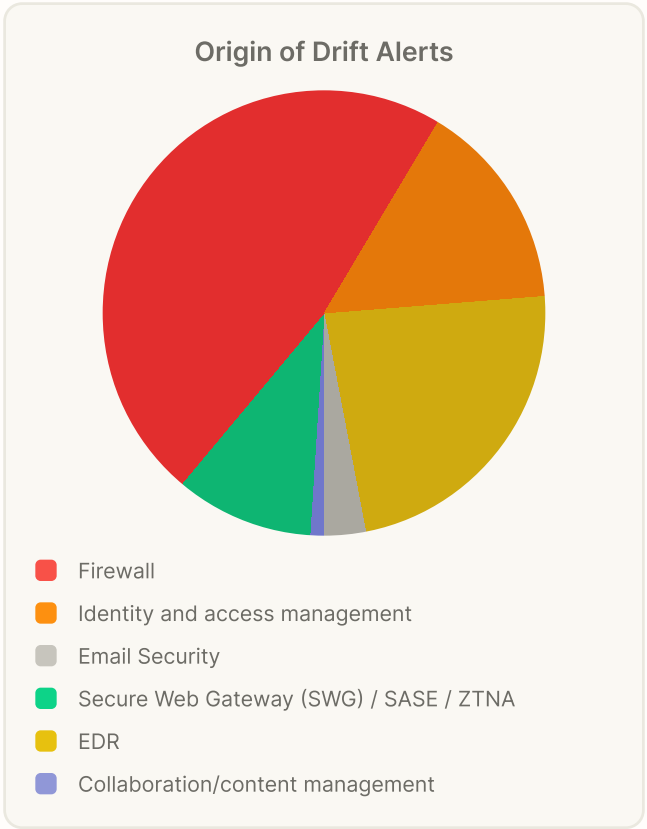
The number of times per month organizations check for cybersecurity misconfiguration

42%

Said they suffered an incident originating in the firewall making it the most frequently cited source of configuration delated drift exposure.

Reach Telemetry

- 12-month anonymized telemetry dataset spans a robust sample of production environments across multiple industries incorporating millions of security control observations.
- 13 drifts per day with 12 tied to real exposure
- This means 92% of identified configuration drift is tied to real security exposure – nine in ten drift alerts require a response.
- On average, 47% of the drift alerts detected by Reach over the 12-month period originated in firewall technology, 23% originated in EDR tools, 15% originated in IAM tools, 10% in Secure Web Gateway (SWG) / SASE / ZTNA tools, and 3% in email security.
- Reach scans daily for configuration changes, identified and remediated issues are automatically recorded helping build a picture of the cadence of drift in the organization.



SECTION 2

The Challenge of Continuous Configuration Drift

Reach's telemetry data, collected across production environments over the course of a year, shows that configuration drift is not an occasional compliance issue, it is a continuous operational challenge. On average, Reach customers generate 13 drift alerts each day, with 12 linked to genuine security exposure. These findings demonstrate that drift is not simply background noise. Security teams are routinely dealing with changes that have the potential to weaken security controls and create exploitable gaps.

Origins of drift alerts detected by Reach over 12 -month period

Firewall	47%
Secure Web Gateway (SWG) / SASE / ZTNA	10%
IAM or IAM governance	15%
EDR	23%
Email security	3%
Collaboration / content management tools such as MS Teams, Sharepoint, Google Workspace	1%

The data also highlights the concentration of risk within network security controls. Firewall tools and management platforms generated by far the highest volumes of both drift alerts and security findings, confirming that firewall environments are particularly susceptible to configuration drift. This is not necessarily due to weaknesses in the technologies themselves. Rather, firewalls sit at the heart of network enforcement and rely on complex rule sets that constantly evolve in response to business requirements, troubleshooting activities and security exceptions.

Every rule change has the potential to move a configuration further from its intended state. Over time, stale rules accumulate, shadowed policies obscure visibility, and overly permissive access controls can emerge, creating hidden attack paths that are difficult to detect through manual reviews alone.

Risk resides in firewalls

Reach's analysis found that, overall, customers typically experienced 13 drift events per day, with 12 tied to genuine security exposure. In other words, 92% of alerts relate to risk. This further validates our industry research findings – if 47% of drift alerts relate to firewall technology, the percentage tied to genuine risk is 43% - close to the 42% of our industry respondents who said they'd suffered a breach or near miss tied to a firewall misconfiguration.

When combined with the independent survey findings, the telemetry data paints a clear picture. Security practitioners recognize configuration drift as a growing challenge, and the operational data confirms that risk is continuously accumulating beneath the surface.

SECTION 3

Identifying Material Configuration Drift

While configuration drift is widespread, not every configuration change creates meaningful risk. The challenge for security teams is not simply detecting drift but determining which deviations materially weaken security controls and require remediation.

Reach's analysis reveals an important distinction between drift alerts and security findings. Drift alerts indicate that a configuration has changed from its intended state, while security findings identify instances where those changes create genuine security exposure. This distinction is critical because some controls may generate frequent drift alerts yet have little impact on security outcomes, whereas seemingly minor configuration changes can significantly weaken an organization's security posture.

Identity and Access Management tools offer a case in point. They generated 15% of all drift alerts recorded by Reach among this customer sample over 12 months but only represented 0.3% of priority security findings.

Security Findings (%)

Firewall	87.98
EDR	9.64
Email Security	1.24
Collaboration/content management tools such as MS Teams, Sharepoint, and Google Workspaces	0.63
Secure Web Gateway (SWG) / SASE / ZTNA	0.09
IAM	0.34

The findings also demonstrate that materiality is highly contextual. The same configuration change may be considered acceptable in one environment while representing an unacceptable risk in another, depending on factors such as threat exposure, business objectives, compliance requirements and overall risk appetite.

This challenge persists despite widespread investment in security management, compliance and automation tools. Organizations have become increasingly effective at identifying that changes have occurred, but many still struggle to determine whether those changes introduce meaningful risk and warrant immediate attention.

Configuration Risk: Not Just a Security Tool Issue

Reach tracks security-related configurations for all products, not just security-specific tools. Collaboration and content management tools generated 1% of alerts and .06% of material security findings. Though small, these figures still matter, as unused controls and configuration drift can be easily exploited.

This is the case for the UNC6692 exploitation, as Microsoft's April 2026 advisory confirms: the campaign abuses legitimate external collaboration features, and victims override clearly communicated warnings to allow it to succeed. There is no issue with the controls and guidance Microsoft built, but when its product enters busy work environments and external access needs to be granted to vendors and clients, temporary exceptions to original policies are introduced. The only way to address this risk is through continuous observability. Without it, you'll have no idea how many gaps are in your environment, ready for an attacker to walk right on through. This is equally true as AI models become a primary interface for security operations. A frontier model asked to assess exposure or recommend remediation is only as accurate as the control-state context it receives. Reach provides that foundation — deterministic, continuously validated, and sourced directly from the controls themselves rather than inferred from logs or telemetry.

Reach addresses this challenge by evaluating drift through a threat-informed lens, enabling organizations to distinguish between routine configuration variance and exposures that materially increase risk. By prioritizing findings based on security impact rather than change volume, security teams can focus remediation efforts where they will deliver the greatest risk reduction.

Another key Reach advantage is its broad environment coverage. While specialist monitoring tools covering individual products are effective at a granular level, they are unavoidably siloed and do not deliver a comprehensive picture of configuration risk across the full technology stack. This makes it harder to prioritize remediation at a global level, meaning security teams are often pulled in multiple directions because relative risk is not assessed in the context of the entire technology landscape.

Ultimately, the goal is not to eliminate every instance of configuration drift – that isn't possible in a live environment. Instead, it is to identify the drift that matters most and ensure that limited security resources are directed towards exposures that genuinely threaten the organization.

SECTION 4

Recognizing Drift Danger Zones

One of the most important findings from the telemetry analysis is that configuration drift does not occur randomly. Instead, drift events consistently cluster around several predictable operational periods and activities, creating identifiable "danger zones" where organizations are significantly more likely to experience configuration changes that lead to security exposure.

Platform changes and external updates

The data shows that some of the largest drift spikes occurred simultaneously across multiple, otherwise unrelated environments. In one instance, four separate organizations experienced an identical surge in drift alerts on the same day and at the same volume. This pattern strongly suggests that external factors, such as software updates, platform changes or other ecosystem-wide events, can act as triggers for widespread configuration drift. These findings highlight the importance of maintaining visibility not only over internal changes, but also over the broader technology landscape on which security controls depend.

Seasonal/holidays

Seasonal patterns were equally revealing. The holiday period, particularly during the second half of December, produced the two largest spikes in drift activity recorded during the year-long analysis. This may reflect a combination of factors, including organizations rushing to implement changes before operational freeze periods, reduced staffing levels, delayed review processes and lower levels of oversight during holiday breaks. Whatever the underlying cause, the data suggests that periods traditionally associated with reduced business activity can actually create elevated risk from a configuration management perspective.

Weekly cadence

The research also identified clear weekly patterns. Drift activity increased noticeably between Wednesday and Friday, indicating that configuration changes become more frequent as organizations move towards the end of the working week. This trend is particularly significant because changes implemented late in the week often have less available time for validation, testing and remediation before staff availability decreases over the weekend. Small configuration deviations introduced during these periods can therefore remain undetected for longer, increasing the likelihood of exposure.

Software releases

Industry-wide maintenance and patching cycles also appeared to influence drift levels. Several of the largest multi-customer drift events occurred shortly after scheduled software release and patch deployment periods. While patching remains a critical security best practice, these findings reinforce the reality that any large-scale change activity introduces opportunities for configuration variance. Organizations should therefore treat major update cycles as periods requiring heightened monitoring and validation rather than assuming successful deployment automatically equates to a secure outcome.

Month-end

Month-end activity presented a slightly different pattern. The analysis found that month-end was not a frequent driver of drift, but when spikes did occur during these periods, they tended to be larger than average. This may reflect concentrated operational activity, project milestones, reporting deadlines or scheduled infrastructure changes that are often aligned to monthly business cycles. As a result, month-end periods can still represent a meaningful source of risk despite occurring less consistently than other drift danger zones.

These findings demonstrate that configuration drift is closely linked to operational behavior rather than occurring in isolation. Holiday periods, major technology updates, concentrated change windows and end-of-week deployment activity all create conditions where drift is more likely to emerge. By understanding these patterns, organizations can proactively strengthen monitoring, validation and change-control processes during higher-risk periods, reducing the likelihood that routine operational activity evolves into genuine security exposure.

#	DANGER ZONE	WHAT DRIVES IT	INTENSITY
1	Vendor Updates	Identical simultaneous spikes across multiple unrelated environments on the same day — external triggers are invisible to internal change controls	●●●●●
2	Holidays	The two largest annual drift spikes both occurred in late December — reduced oversight, rushed pre-freeze changes, and delayed review cycles	●●●●○
3	End of week	Drift rises consistently Wednesday through Friday — changes made late in the week have less time for validation before the weekend	●●●○○
4	Patch cycles	Scheduled maintenance windows generate multi-customer events — successful deployment doesn't mean controls are correctly configured	●●●○○
5	Month-end	Lower frequency but larger magnitude when it occurs — concentrated project deadlines and infrastructure changes aligned to monthly business cycles	●●○○○

SECTION 5

Recommendations & Conclusion

The findings in this report make one reality clear: configuration drift is no longer simply an operational challenge, it is a persistent security risk. As environments grow in scale and complexity, even well-designed controls frequently diverge from their intended state, creating hidden exposure that often goes unnoticed until a security review, audit, or incident uncovers the problem. While organizations continue to invest heavily in security technologies, our analysis shows that the effectiveness of those investments ultimately depends on whether controls remain aligned with security intent over time.

To address this challenge, security leaders must:

- **Adopt continuous security assurance across the entire technology stack.** Rather than assuming that deployed controls are functioning as intended, security teams must continuously verify that configurations remain effective, identify where drift has occurred, and understand whether that drift creates meaningful exposure. Reach Network Security Assurance is purpose-built for this requirement — continuously validating whether network security controls are configured, enforced, and aligned with security intent across firewalls, SASE, and adjacent enforcement points.
- **Understand drift materiality.** Not all drifts carry the same level of risk. A threat-informed strategy that prioritizes material exposures enables teams to focus remediation efforts where they will have the greatest impact on reducing organizational risk.
- **Prioritize remediation based on real-world exposure,** drift danger zones and attacker relevance rather than the volume of configuration changes alone.
- **Reduce the time between drift detection and remediation** by providing clear, actionable, and prioritized guidance to the teams responsible for resolving issues.

At the heart of this analysis lies a fundamental question: are the security controls across your environment actively protecting the organization right now? If gaps exist, where are they, how significant is the risk they create, and how quickly can they be addressed? The persistence of configuration drift demonstrates that traditional approaches based on periodic reviews and point-in-time assessments are no longer sufficient. Existing tools highlight the problem, but they do not eliminate it.

Reach's approach is designed to close this gap. By continuously validating security intent, detecting drift and hidden exposure, prioritizing weaknesses through a threat-informed lens, and guiding remediation to the exact rule, device or profile creating risk, Reach enables organizations to move from reactive security management to continuous assurance. The result is greater confidence that security controls are operating as intended, security investments are delivering measurable protection, and risk is being reduced where it matters most.

Looking ahead, the pressure on configuration assurance will only intensify. As AI agents take on more active roles in security operations — adjusting controls, triaging findings, and driving remediation — the need for accurate, continuously validated control-state context becomes foundational. Organizations that build that foundation now will be better positioned to deploy AI-assisted defense safely and cost-effectively. Those that don't will find that automation amplifies their exposure as readily as it reduces it.

Ultimately, the goal is not to eliminate every instance of configuration drift. It is to identify the drift that creates material security risk, understand its potential impact, and act before it can be exploited. Organizations that adopt this approach will be better positioned to maintain effective security controls, strengthen resilience and keep pace with an increasingly dynamic threat landscape.

CONTINUOUS SECURITY ASSURANCE

Key Questions Every Organization Should Ask

The following questions can help security leaders evaluate the maturity of their continuous security assurance program and identify potential gaps in visibility, prioritization and governance.

Are We Clear on Our Security Intent?

Before organizations can determine whether security controls have drifted, they must first understand what those controls are intended to achieve.

- Have we clearly documented how each critical security control is expected to protect the organization?
- Can we differentiate between a configuration that is approved and one that is simply technically valid?
- Do our security assessments reflect our specific threat landscape, business requirements and risk appetite?

Do We Have Complete Visibility into Change?

Effective assurance depends on knowing what has changed, when it changed and why.

- Can we identify configuration changes across all major security control domains, not just firewalls?
- Do we know who made each change, when it occurred and whether it was authorized?
- Can we distinguish newly introduced drift from long-standing or recurring configuration issues?

Are We Prioritizing Risk Effectively?

Not every configuration change creates meaningful risk, making context and prioritization essential.

- Can we determine which configuration changes have introduced genuine security exposure?
- Are findings prioritized using threat relevance, asset exposure and business impact rather than static severity ratings alone?
- Can we identify when a single change affects multiple policies, systems or layers of security controls?

Can We Remediate and Validate Quickly?

The ability to detect drift is valuable only if teams can respond efficiently and verify that issues have been resolved.

- Does every finding point directly to the rule, policy, object, device or configuration requiring attention?
- Are remediation responsibilities and service-level expectations clearly defined?
- Do we automatically verify that corrective actions have restored the intended security outcome?

Are We Operationally Prepared for High-Risk Periods?

The telemetry analysis highlighted several predictable "drift danger zones" where organizations are more likely to experience configuration issues.

- Do we increase monitoring and assurance activities around major technology rollouts, patch cycles, maintenance windows, holiday periods and change freezes?
- Are security assurance insights incorporated into change management, operational reviews and post-incident assessments?
- Do we track key metrics such as recurrence rates, time to detect, time to remediate and the percentage of changes that receive post-change validation?

Does Our Coverage Extend Across the Entire Security Estate?

Security assurance should provide a comprehensive view of security effectiveness rather than focusing on individual technologies in isolation.

- Are we validating security outcomes consistently across different vendors, platforms and control categories?
- Do our existing tools manage configuration workflows, validate security effectiveness, or both?
- Can leadership confidently answer the question: *"Are our security controls protecting the organization as intended right now?"*

Ultimately, organizations that can answer these questions with confidence are better positioned to identify meaningful drift, reduce unnecessary noise and maintain a security posture that remains aligned with business objectives as their environments continue to evolve