nagomi

The Illusion of Maturity

2025 Enterprise Exposure Snapshot

Why security programs look strong while exposure compounds

Table of Contents

Introduction and Executive Summary

Security dashboards signal maturity. Exposure data shows fragility.

Across enterprises, patch rates improve and coverage charts look strong. Yet breaches still originate from familiar weaknesses: incomplete MFA, inconsistent endpoint protection, lingering misconfigurations, and users who complete training but still fall for modern social engineering.

This disconnect defines one of the central challenges facing modern security teams. The issue does not stem from a single failed control or threat. It stems from how organizations measure and manage exposure in isolation while attackers exploit the gaps it converges.

Risk-based vulnerability management (RBVM) highlights this tension clearly. RBVM has helped enterprises prioritize exploitable vulnerabilities and reduce patching noise. As a result, vulnerability programs have become some of the most mature and operationalized functions in security. Vulnerability counts fall. Tickets close. Patch SLAs improve. RBVM reduces noise and improves prioritization, but cannot eliminate exposure when structural controls that govern exploitability remain weak.

This report shows where reported maturity collapses under live exposure conditions.

Drawing on control assessment and exposure data from thousands of endpoints across dozens of enterprises and over 70 security controls integrated with Nagomi, the analysis addresses questions most dashboards cannot. Measuring control effectiveness by outcomes rather than coverage exposes where security programs break down.

- Where exposure concentrates even when vulnerability metrics look strong.
- Which control failures recur most consistently across organizations.
- Which misconfigurations and degraded protections create the largest blast radius.
- Where multiple control failures overlap to form reliable attack paths.

This report is for security leaders and teams who recognize the gap between reported maturity and operational reality. The findings prioritize work that removes exposure rather than activity that improves dashboards.

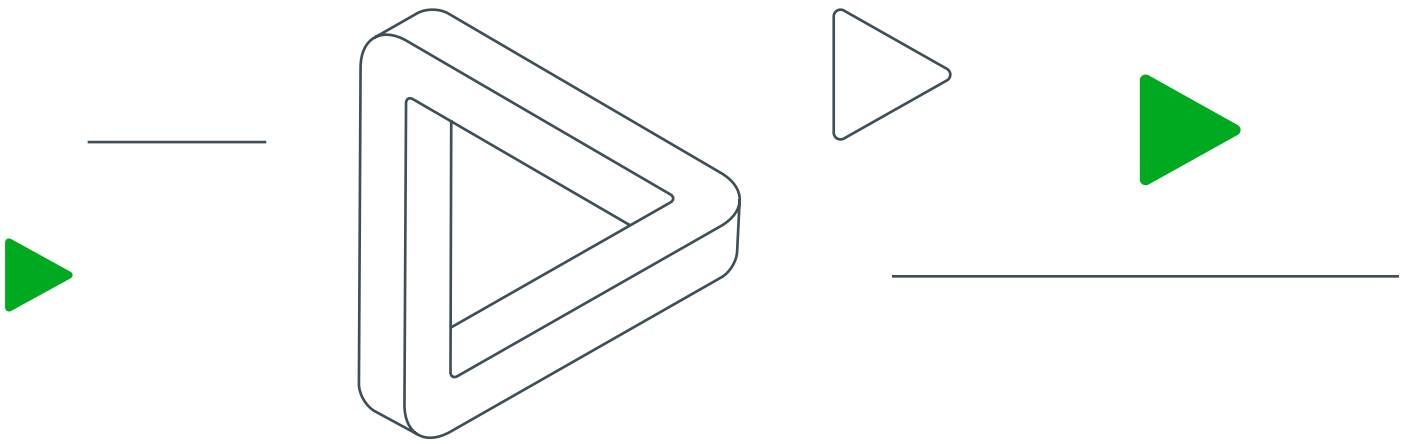

Key findings

Vulnerability programs perform well, but foundational controls lag.
Vulnerability management passes on 91% of assets, while identity, endpoint, and user readiness controls pass on only 30–52%, leaving structural weaknesses that determine exploitability.

Real risk forms where multiple controls fail together.
Fewer than 30% of assets demonstrate effective coverage across identity, endpoint, and awareness simultaneously, allowing attackers to exploit the intersections single-control metrics never reveal.

Controls leaders assume are "done" often are not.
Despite widespread deployment, 75% of organizations show incomplete MFA or weak EDR enforcement, with advanced endpoint protections failing in more than 60% of environments.

Risk concentrates in misconfiguration-driven exposure conditions.
Across organizations, 20–40 findings collapse into ~7 high-impact exposure conditions, where a single control failure can expose thousands of assets and outweigh dozens of isolated CVEs.

Training completion does not equal readiness.
Fewer than 30% of assets pass security awareness controls, and human-driven failures repeatedly align with weak identity and endpoint enforcement to form reliable intrusion paths.

Exposure reflects operating discipline, not tool quality.
Controls that require sustained tuning, integration, and cross-team ownership consistently underperform, creating blind spots regardless of vendor selection.

These findings show that exposure does not stem from isolated failures. It concentrates where multiple weaknesses align, and it persists when organizations measure controls independently rather than managing attack paths end to end.

Observed Failure Pattern #1: Strong Vulnerability Scores, Weak Structural Controls

Enterprises demonstrate mature behavior in patching and scanning, yet continue to struggle with the structural controls that determine whether vulnerabilities and misconfigurations become exploitable. Control assessment data shows a significant performance gap between Vulnerability Management and foundational controls in identity, endpoint protection, and user readiness.

Control performance snapshot: Asset Pass Rate by Security Program

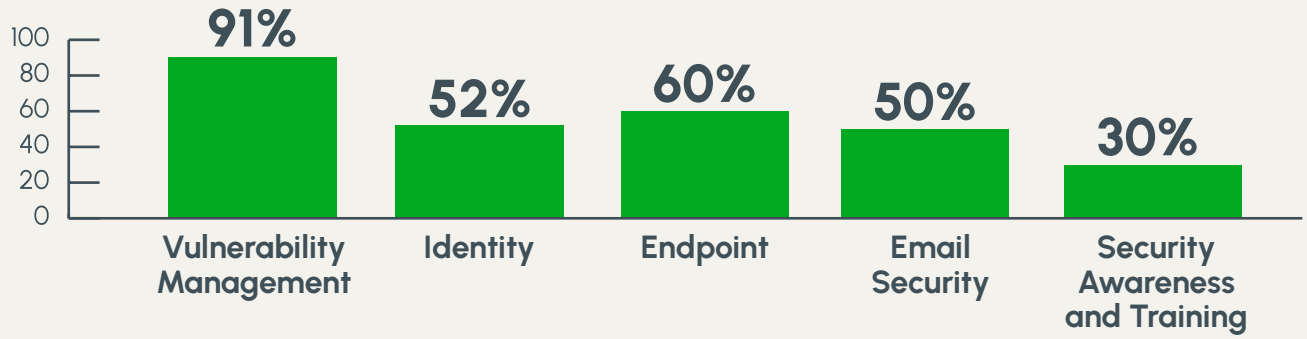

Vulnerability Management significantly outperforms every other control area. Identity and endpoint controls hover around the midpoint. Training and readiness fall well below the others.

What fails most often

Password strength controls fail on 50% of assessed assets. Organizations often document policies, yet enforcement remains inconsistent across systems and directories.

Document and script execution prevention passes at 60%. Many endpoints still allow macro- and script-based exploitation paths that attackers repeatedly use for initial access.

Authentication policy enforcement performs the worst at 30% pass rates, including gaps in MFA coverage, session baselines, and access control enforcement.

Cloud identity configuration baselines fall below 30%, indicating widespread infrastructure weaknesses.

The automation divide

Control performance diverges along operational lines. Organizations perform well where execution is reinforced by automation, such as vulnerability scanning and patching, which benefit from centralized workflows, clear ownership, and measurable outcomes.

Identity, endpoint, and awareness controls operate under a different model. Their effectiveness depends on sustained coordination across multiple teams. As coordination weakens, enforcement becomes uneven. Authentication coverage fragments, endpoint policies drift, and training completion is recorded without validating behavioral readiness.

These breakdowns do not occur independently. They frequently align on the same assets, forming convergent exposure conditions that are not visible through single-control performance metrics.

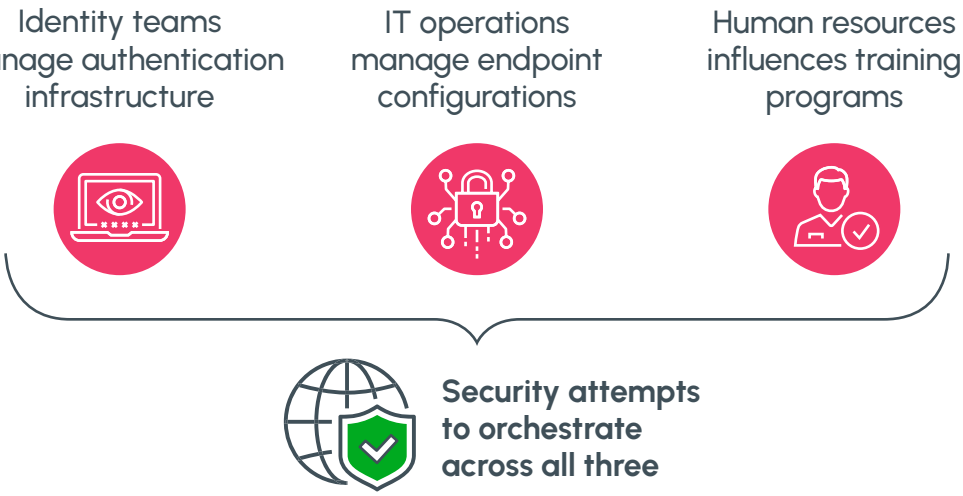

Why this matters

Breaches rarely stem from unpatched CVEs alone. They succeed when authentication is weak, execution controls remain permissive, or prevention features fail in practice. The gap between vulnerability success metrics and real-world resilience often sits in these structural controls.

Observed Failure Pattern #2: The Controls Everyone Claims Are "Done"

MFA and EDR are commonly reported as complete, inconsistent enforcement leaves critical authentication paths and endpoint execution flows exposed, enabling lateral movement in practice.

75% of organizations show missing EDR coverage or weak/ incomplete MFA.

Deployment vs. Operationalization

Organizations deploy controls in broad strokes, then stall in the details.

MFA deploys for primary user access paths but remains inconsistent across service accounts, legacy applications, administrative interfaces, and high-risk workflows.

EDR agents are present, yet advanced protections remain disabled or misconfigured.

Analysis of findings across organizations reveals consistent patterns

- Missing EDR coverage ranks among the most common hygiene findings across the dataset.
- Weak or incomplete MFA appears in more than 30% of organizations.
- More than 60% of organizations fail advanced EDR policy tests, even when agents are widely deployed.
- Endpoint policy failures affect hundreds of devices per environment on average.

Why this matters

The controls that boards and executives assume are complete often fail in practice. MFA and endpoint protections form the foundation of lateral movement prevention and modern zero trust strategies. When enforcement degrades, exploitable vulnerabilities and user-driven entry points become reliable attack paths.

Observed Failure Pattern #3: Tool Sprawl Without Operational Discipline

Security investment is intended to reduce exposure. In practice, expanding toolsets often increases it. As environments scale, each additional control raises enforcement complexity faster than it reduces attack surface.

Controls that depend on sustained tuning and cross-team coordination break down first. This reveals an operational problem, not a product problem.

The tools that underperform share common traits:

- Complex policy models
- Ongoing tuning requirements
- Reliance on agent enforcement
- Shared ownership across multiple teams

The tools that underperform operationally

The lowest performance consistently appears in controls tied to user behavior and endpoint enforcement, where effectiveness depends on continuous coordination rather than one-time deployment.

Tool Category	Operational Challenge
Email Security & Phishing Detection	Requires ongoing tuning, user behavior correlation
Security Awareness & Training	Completion tracking, behavioral validation needed
Endpoint Configuration Management	Policy enforcement, multi-team coordination
Endpoint Privilege Management	Complex policy models, sustained operational attention

This reinforces a central theme of the data. Exposure stems from configuration inconsistency and operational gaps, not from product selection. Buying best of breed does not guarantee coverage, enforcement, or resilience.

Coverage blind spots

- Incomplete scanner coverage indicates visibility gaps rather than low vulnerability prevalence.
- Identity platforms show objects without baseline enforcement.
- Cloud and infrastructure security tools often show partial integration depth.

Why this matters

Exposure persists not because organizations lack tools, but because operating models often do not sustain enforcement across the full environment. Buying best of breed does not guarantee coverage, consistency, or resilience.

Observed Failure Pattern #4: Misconfigurations Drive More Risk Than Vulnerabilities

Risk is not evenly distributed across findings. A small number of misconfiguration-driven conditions account for a large portion of enterprise attack surface. These conditions often span multiple control domains, demonstrating how vulnerabilities, endpoint protections, identity controls, and policy enforcement failures converge into exploitable exposure.

Highest-impact exposure conditions (average assets impacted)

Rank	Exposure Condition	Avg Assets Impacted
1	Exploited Critical RCE + EDR Exploit Protection Disabled	**2,000**
2	Active Exploited Vulnerability + Missing/Ineffective Endpoint Controls	**1,400**
3	Widespread Vulnerabilities + Suppressed Security Baselines	**800**
4	Misconfigured Execution Policies + Vulnerable Script Engines	**500**
5	Privilege Path Exposures (vuln + permissive escalation)	**250**

Findings represent individual control or configuration gaps. Exposure conditions represent correlated combinations of findings that form realistic attack paths.

Risk is concentrated, not uniform

Fixing a single high-impact condition often reduces more exposure than closing dozens of narrower CVE findings. This changes prioritization and executive expectations. Maturity is measured by elimination of high-blast-radius exposure conditions, not by closure volume.

Misconfigurations scale faster than vulnerabilities

Narrow CVE-specific findings often impact a small number of systems. Broad policy failures and baseline drift can impact thousands. In some organizations, a single training configuration failure affected more than 20,000 assets.

Why this matters

Teams that prioritize by asset impact and convergence reduce risk faster than teams that prioritize by vulnerability severity alone. Exposure management should focus on eliminating the few conditions that expand attack surface at scale.

Observed Failure Pattern #5: Training Done on Paper, Failing in Reality

Human-focused controls rarely fail in isolation. Across the dataset, email security, user behavior, and identity enforcement consistently degrade together. This convergence forms the most reliable initial access path in modern attacks.

<30% of assets pass training-related controls.

What the data shows

Awareness tools record the lowest performance across all programs.

Advanced training completion tests fail across thousands, and in some cases tens of thousands, of users in a single organization.

Email security and human behavior tests rank among the most frequently failed across the dataset.

Endpoint tests tied to phishing protection and execution controls also show widespread failure.

These patterns indicate that awareness programs are often deployed without full operational rigor. Training gets assigned and tracked, yet failure counts show large user populations remain exposed.

A repeatable high-risk profile emerges

- Baseline training completed but advanced simulations failed
- Weak password enforcement or incomplete MFA
- Permissive execution controls on endpoints

Each issue appears manageable in isolation. Together, they create a reliable intrusion path that bypasses single-control defenses.

Why this matters

Human error remains a frequent breach catalyst, yet the underlying issue is operational. Completion does not equal readiness. Until organizations treat behavioral readiness with the same rigor as technical controls, this path remains consistently exploitable.

Bottom Line: Risk Concentrates Where Multiple Controls Fail Simultaneously

Organizations measure controls independently: patch rates, identity scores, endpoint deployment, training completion. Each metric may show progress, yet none quantify where risk actually concentrates.

Risk concentrates where controls fail together on the same assets.

> 70%

of assets are exposed because at least one control fails along the same attack path.

< 30%

of assets demonstrate effective control coverage across identity, endpoint, and awareness domains at the same time.

The convergent risk problem

Security teams structure work by domain:

- Vulnerability teams manage scanning and patching
- Identity teams manage authentication and access
- Endpoint teams manage configurations and protection
- Awareness teams manage training programs

This model works for controls with clear ownership. It fails when risk spans systems, identities, endpoints, and people simultaneously.

A repeatable high-risk asset profile emerges:

- Baseline training completed but advanced simulations failed
- MFA missing on critical access paths
- Passwords fail strength requirements
- Endpoint lacks execution controls
- EDR agent present but exploit protections disabled
- Lower-severity vulnerabilities become critical in combination

Each issue looks manageable alone. Traditional metrics show progress. The attack succeeds because the gaps align across the same path.

Why this matters

Improving a single control area delivers limited risk reduction when other controls remain weak on the same assets. Reducing risk requires measuring and managing exposure where controls intersect.

Recommendations

Reducing exposure requires shifting from control-centric execution to coordinated elimination of the conditions attackers actually exploit.

Restructure ownership around exposure conditions, not control domains

- Establish cross-functional exposure squads with dedicated ownership of specific high-impact conditions.
- Define accountability for integrated outcomes rather than domain-specific activities.
- Create escalation paths when cross-domain coordination fails.

Measure integrated exposure, not siloed control performance

- Correlate data to identify assets where multiple control gaps converge.
- Report exposure conditions that combine vulnerability, endpoint, identity, and awareness failures on the same systems.
- Rank exposures by asset impact and convergence rather than vulnerability severity alone.

Treat hygiene with the same rigor as vulnerability closure

- Track missing EDR coverage, incomplete MFA deployment, failed endpoint baselines, and configuration drift.
- Define SLAs for hygiene remediation based on severity and asset criticality.
- Include hygiene trends in executive dashboards alongside exposure metrics.

Connect security awareness to identity and endpoint posture

- Correlate training completion and simulation performance with identity and endpoint telemetry.
- Define high-risk users based on combined signals: training failures, weak authentication, permissive access patterns, endpoint policy violations.
- Implement targeted interventions for users with persistent convergent risk profiles.

Align executive metrics with integrated outcomes

Report three core trends:

- high-impact exposure count
- average closure time
- total affected assets

Frame security posture as attack surface eliminated, not control compliance achieved.

Closing Perspective

Security dashboards will continue to show maturity as long as programs measure individual controls in isolation. Actual attack surface concentrates where multiple controls fail on the same systems. Misconfigurations, coverage gaps, and degraded enforcement drive the majority of real exposure. No amount of additional tooling or incremental optimization within silos will close that gap.

Only when organizations measure and manage exposure at the intersections where attackers operate will reported maturity begin to reflect actual resilience. That shift requires more than better reporting. It requires an operating model built for execution across domains, where high-impact exposure conditions have clear ownership, measurable closure, and continuous reassessment.

This is the next phase of exposure management. As agentic automation becomes more practical across security operations, organizations can reduce coordination drag by using systems that continuously identify high-impact exposure conditions, route work to the right owners across identity, endpoint, and vulnerability domains, validate remediation actions, and confirm closure over time. The advantage is not replacing teams. The advantage is making exposure reduction operational at scale.

The programs that mature fastest will not be the ones with the most tools or the greenest control dashboards. They will be the ones that consistently eliminate high-blast-radius exposure conditions, shorten time to closure, and maintain those gains as environments change.

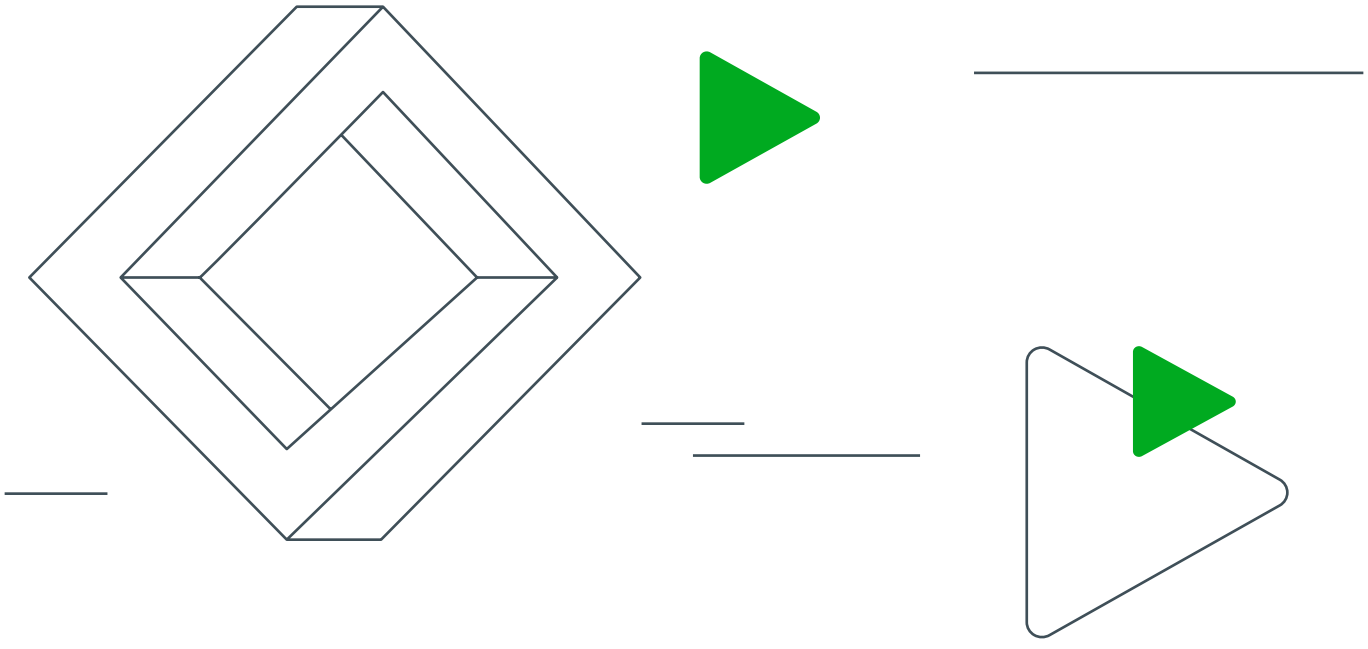

Methodology

This report analyzes aggregated control assessment, hygiene, and exposure data collected across enterprise environments using Nagomi Security during the reporting period. The goal of the methodology is to make clear what was measured, how exposure was derived, and where the analysis draws its boundaries.

Data Sources

The analysis draws on multiple classes of security telemetry and assessment data, including:

- **Control assessment results** across six security program areas: Vulnerability Management, Identity, Endpoint and Email Security, Network Security, Data Protection, and Security Awareness and Training
- **Hygiene findings,** including missing EDR coverage, incomplete MFA deployment, failed endpoint and identity policy baselines, and configuration drift
- **Exposure rules,** which correlate vulnerabilities with configuration state and control coverage to identify realistic attack chains.
- Aggregated organizational metrics, including pass rates by program and capability, number of findings per organization, and average assets affected per finding
- **Convergent risk signals,** identifying assets where multiple control failures appear simultaneously across vulnerability, identity, endpoint, and human domains

Analysis Scope and Inclusion Criteria

The analysis reflects the **most recent completed assessment scoring run per organization** at the time of data collection in order to represent current operational posture rather than historical trends. Organizations without sufficient integrated data sources for a given control domain were excluded from analysis of that domain to avoid skewed results.

Exposure Definition and Correlation Logic

Exposure findings represent **correlated conditions,** not isolated control failures. Each exposure condition combines two or more of the following elements:

- Presence of an exploitable or high-risk vulnerability
- Degraded or missing security controls (for example, disabled exploit protection or missing MFA)
- Configuration states that increase blast radius or likelihood of successful exploitation

This correlation approach is designed to surface **realistic attack chains,** rather than enumerate all possible weaknesses. Exposure rules intentionally prioritize signal quality and impact over completeness.

Asset Impact Measurement

Asset impact represents the **average number of unique assets affected per exposure condition per organization** across the dataset. This metric is used to distinguish narrow, low-impact findings from high-blast-radius conditions that materially expand attack surface.

Interpretation and Limitations

The analysis focuses on **patterns and distributions,** not individual organizations. Results are presented in aggregate to identify systemic behaviors and structural gaps rather than rank customers, industries, or tools.

Because the dataset reflects integrated tooling and available telemetry, the findings may underrepresent exposure in environments with limited visibility. As integrations expand and exposure rules evolve, future reporting may refine or extend these conclusions.

Intent of the Analysis

This report is designed to support exposure reduction, not compliance reporting. Metrics are selected to illustrate how vulnerabilities, misconfigurations, identity posture, endpoint enforcement, and user behavior interact to create exploitable conditions. The methodology favors practical signal over theoretical completeness, aligning measurement with how attackers actually operate.

nagomi

Learn more at: nagomisecurity.com