HORIZON3.ai

The State of Assumed Security

Why measuring activity is *not* the same as measuring resistance.

2026

TL;DR **This report examines how organizations measure security and whether those measurements reflect real resistance to attacker behavior.**

Across 750 security leaders and practitioners, the findings reveal a consistent pattern. Executive confidence in prevention and detection capabilities is high. Remediation workflows are mature. Performance indicators are widely tracked. Yet fewer organizations consistently test whether vulnerabilities can still be exploited, whether detection approaches stop lateral movement in time, or whether fixes actually remove risk. The issue is not effort. It is confirmation.

As attackers increasingly chain identity, infrastructure, and cloud weaknesses, the difference between having a control in place and knowing it stopped an attack becomes consequential.

This research explores where measurement systems emphasize activity over confirmation and what that means for the next phase of cybersecurity maturity.

1. Executive Summary

Measuring security, or measuring activity?

Modern security programs are busy. Assets are scanned. Alerts are generated. Findings are prioritized. Patches are deployed. Dashboards are updated. Tickets are closed. For years, this activity has been treated as evidence of security.

But it raises a harder question: **Are organizations measuring whether they can stop a real attacker, or whether work was completed?**

Across the United States, United Kingdom, and Europe, confidence in security controls is high:

93%

of CISOs say they could prove their organization took reasonable, validated steps to prevent a breach.

97%

are confident their endpoint protection would detect lateral movement or privilege escalation.

96%

believe their Security Operations Center could identify an attacker operating inside the environment.

Operational practice reveals a different standard.

Only 30% of CISOs report that their organizations patch and then test to ensure risk has been remediated. Nearly half patch and rescan with a vulnerability scanner instead. Just 12% say they have validated EDR effectiveness within the last three months. Only 26% use red team exercises or pentesting to assess SOC detection capability. Among practitioners, 33% assume scanner findings are accurate without further testing, and 17% do not validate findings at all.

Response to known exploited vulnerabilities shows similar lag. Only 11% of practitioners report validating or patching within 24 hours of a CISA or ENISA alert. Many require a week or more to confirm if they're exposed.

The pattern is consistent. Security programs are structured around workflow completion: scan, patch, rescan, close. Detection tools are deployed and monitored. Implementing automation increases speed. What is less consistent is confirmation.

These five structural themes define the current state of enterprise security:

1
Confidence in controls exceeds the frequency of adversarial testing

2
Severity-based prioritization remains common, even when exploitability is unverified

3
Detection capabilities are widely trusted despite limited routine stress testing

4
Automation is expanding faster than independent validation

5
Metrics emphasize remediation speed, but not always remediation effectiveness

The issue is not effort. Security teams face competing priorities, approval delays, staffing shortages, and increasing threat velocity.

Security programs are proving that work was completed. They are not always proving that attacks will fail.

Modern attackers do not exploit a single missing or delayed patch. They reuse credentials, pivot between environments, and combine small weaknesses into meaningful impact.
In that environment, having a control in place is not the same as knowing it can stop an attack.
Closing a ticket is not the same as eliminating an attack path.

The findings in this report point to a structural gap between assumed security and demonstrated resistance.

The next phase of cybersecurity maturity will not be defined by how many controls are deployed or how quickly tickets close. It will be defined by whether organizations systematically confirm that their defenses can withstand realistic attacker behavior.

Confidence does not stop an attacker. *Confirmation does.*

2. The Illusion of Low Risk

When “low risk” meets high exposure

How secure is your organization today? The answer depends on who you ask. Among CISOs surveyed:

MORE THAN
30%
describe their organization as low or minimal risk if targeted by a skilled attacker.

ONLY
41%
classify themselves as high risk.

71%
describe their current exposure as high risk.

The divergence is clear.

At the executive level, a meaningful portion of leaders believe their organization would withstand a determined attack. At the operational level, the majority of practitioners see significant exposure. This difference in perception often shapes decision-making.

When leadership perceives risk as moderate or low, funding decisions, staffing urgency, and remediation timelines follow that signal. When practitioners know exploitable conditions still exist, operational friction increases.

The issue is not optimism. It is *alignment.*

Executives often evaluate readiness through policies, tooling, and reporting structures. Practitioners evaluate readiness through active attack paths, detection gaps, and weaknesses still present in the environment.

YET, ONLY
17%
of practitioners report conducting weekly pentesting or some form of adversarial testing.

For most organizations, this kind of evaluation is periodic rather than routine. Confidence that is not regularly challenged is not proven resilience. It is assumption.

When testing is infrequent, confidence and exposure are measured on different scales. That misalignment sustains the illusion of low risk.

3. Patch, Scan, Assume

Vulnerability management still prioritizes severity over exploitability.

For most organizations, vulnerability management remains the backbone of security operations. The workflows are mature and well instrumented. The question is not whether work is being done. It is what that work actually confirms.

When asked to describe their remediation process:

NEARLY HALF

of CISOs say their orgs patch and then rescan with a vulnerability scanner.

30%

report patching, then testing for exploitability using real exploits or adversarial tools.

17%

say they patch and assume closure.

When asked how they validate vulnerability findings and confirm exploitability among practitioners:

33%

assume scanner findings are accurate without further validation.

ONLY 27%

verify exploitability through a follow-on pentest.

17%

do not validate findings at all.

At the same time, 40% of CISOs believe scanner results always accurately reflect real exploitability in their environment. Practitioners report that, on average, roughly 25% of vulnerability scanner output is low-value or false positive. Remediation workflows therefore operate on imperfect signal.

Limited validation before remediation often carries forward into limited confirmation after remediation. This creates a dependency on severity scores and version checks instead of confirming whether an attacker can actually succeed.

A rescan confirms that a patch was applied.

It does *not* confirm that
AN ATTACK PATH WAS CLOSED.

It does *not* confirm that
OTHER WEAKNESSES CAN'T BE CHAINED.

It does *not* confirm that
PRIVILEGES WERE REDUCED.

Remediation activity is often treated as risk reduction.

Closing a ticket does not mean closing an attack path.

This does not reflect indifference. Fewer than 1% of practitioners report having no obstacles to faster remediation, and no CISOs report operating without barriers. Approval delays, competing priorities, staffing constraints, and uncertainty about exposure are persistent realities.

Vulnerability management programs are optimized to track closure. They are less consistently designed to confirm that exploitable conditions no longer exist.

Patching reduces exposure. Testing confirms that exploitable conditions no longer exist. Without confirmation, remediation becomes patch, scan, and assume it's fixed.

4. Detection Confidence vs Detection Testing

Trust in detection tools outpaces measured performance

Modern security programs should always assume a breach is possible. Detection and response capabilities are expected to identify attacker behavior before material damage occurs.

As noted earlier, executive confidence in detection controls is extremely high. The vast majority of CISOs express assurance that endpoint protection and SOC monitoring would identify lateral movement, privilege escalation, or an attacker operating inside the environment.

That level of confidence implies detection is assumed to be working.

Testing frequency tells a different story.

ONLY 12%

of CISOs say they have tested EDR effectiveness within the last 3 months.

JUST 26%

report using red team exercises or pentesting to assess SOC detection capability.

ONLY 29%

of practitioners test their SOC's ability to detect attacks monthly or more frequently.

Alert quality tells a more complicated story.

ROUGHLY 52%

of EDR alerts are true positives reported by practitioners on average.

33%

describe detection noise levels as overwhelming.

ONLY 39%

say noise levels are manageable.

Detection is widely trusted, but it is less frequently tested under pressure.

Detection systems may be deployed and monitored and dashboards may show alerts are being acknowledged. Metrics may show response times improving. But without adversarial testing, organizations cannot determine whether detection interrupts an attack before meaningful impact occurs.

Deploying a detection tool does not guarantee it will stop lateral movement or privilege escalation in time. Alerting after an attacker has already escalated privileges or moved laterally is different from interrupting an attack path in time.

Many organizations trust detection controls because they are operational. Fewer consistently test whether those controls block realistic attacker behavior under normal conditions. As with vulnerability management, the gap lies in confirmation. Detection capability is widely assumed, but it is less frequently demonstrated under realistic conditions.

5. Identity: The Underweighted Risk Multiplier

Credential abuse is recognized but not consistently tested

Modern attacks rarely end with initial access. They expand through credential abuse, privilege escalation, and lateral movement. Compromised credentials change the equation. Once authentication is bypassed, attackers often inherit legitimate access paths, making malicious activity look indistinguishable from normal user behavior.

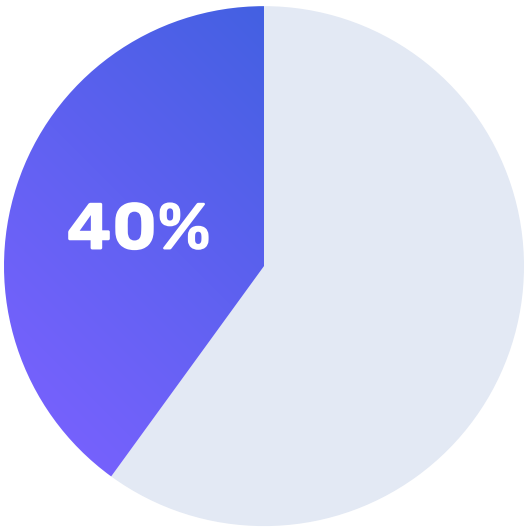

Yet when asked which data types would cause the most damage if stolen, only 40% of CISOs identified credentials as the primary concern. Other categories, including financials and trade secrets, were ranked as equally or more damaging.

Testing frequency shows a similar gap.

ONLY 38%

of practitioners say they frequently test detection of credential-based attacks in Active Directory.

54%

report testing only occasionally.

LESS THAN HALF

of CISOs express full confidence that they could reliably detect attacks targeting directory services.

Alert quality reinforces the concern.

Nearly half of practitioners report that only 26%–50% of Active Directory alerts are confirmed as true positives. Credential compromise is widely recognized as high impact. Routine testing of directory attack paths is less common.

Once credentials are abused, traditional perimeter controls and vulnerability patching offer limited protection. The attacker is already operating with authorized access. Because credential compromise enables lateral movement and escalation, detection testing in this area should be routine. For many organizations, it remains periodic at best. Awareness is high, but embedded testing remains limited.

6. Automation Without Assurance

AI adoption is accelerating faster than validation

Security teams are under pressure to move faster.

Vulnerability backlogs grow. Alert volumes increase. Skilled talent is limited. AI-driven systems are positioned as a way to scale prioritization, remediation, and reporting.

Adoption is already widespread.

60%

of CISOs report that AI is fully integrated into their vulnerability management or remediation workflows.

51%

of practitioners say they are piloting AI-driven patching, ticketing, or remediation systems.

35%

report widespread use.

AI is no longer experimental. It is being embedded in production workflows. But independent validation has not kept pace.

ONLY 17%

of CISOs say they independently test AI-generated recommendations using their own tools.

ONLY 32%

of practitioners say they are fully confident in automated prioritization and remediation, while 54% are somewhat confident.

52%

believe it would save time, while 43% believe it could introduce new risk, when asked about automated retesting after fixes.

AI can accelerate prioritization and ticket closure. It does not confirm that exploit conditions have been eliminated. An automated decision may close a vulnerability faster, but without independent testing, teams cannot determine whether exposure was actually reduced.

The question is not whether AI improves efficiency. It is whether automated decisions are independently verified as effective.

As AI becomes embedded in remediation workflows, confirmation becomes the control mechanism that determines whether efficiency translates into reduced risk. **Acceleration increases speed. Verification determines outcome.**

7. Metrics Without Measurement

Remediation speed is tracked. Resistance is not always tested.

Security programs rely on performance indicators to demonstrate maturity. Mean Time to Remediate. Mean Time to Mitigate. Reoccurrence rates. Patch cycle adherence. SLA compliance. Dashboards track how quickly vulnerabilities are closed. Reports show trend lines. Metrics show whether teams are meeting deadlines.

But speed does not confirm that exposure has been removed.

Only 10% of practitioners report using a dedicated exposure management or pentesting platform to track exploit conditions in an automated way. Most organizations rely on spreadsheets, scanner exports, or ticketing systems to monitor closure rates.

CISOs widely report tracking MTTR and related remediation metrics. Yet when asked about their biggest overall cybersecurity challenge going into 2026, 22% cite validation of fixes and 21% cite demonstrating measurable risk reduction. Both rank ahead of budget constraints and talent shortages. Executives measure remediation speed while acknowledging that confirming risk reduction remains difficult.

When practitioners were asked which about their biggest overall cybersecurity challenge going into 2026:

22%

cite validation of fixes.

21%

cite demonstrating measurable risk reduction.

MTTR reflects response speed. It does not confirm that privileges were reduced, credentials secured, or chained weaknesses removed. A vulnerability may be patched within the SLA while related exposure remains. Improving metrics does not necessarily mean reducing exposure. Dashboards may show faster closure rates while underlying attack paths remain reachable. This does not reflect indifference. Tracking workflow performance is easier than confirming whether an attacker can still succeed.

Performance indicators are necessary. They must also demonstrate that fixes eliminate exploit conditions, not simply that tickets were closed on time.

Without that shift, teams optimize for closure speed while assuming effectiveness.

8. The KEV Reality Check

Speed of confirmation lags behind known exploited threats

Known exploited vulnerabilities are not theoretical. They are weaknesses actively being used in real attacks.

When exploitation is confirmed in the wild, response timing becomes critical. Patching is necessary but confirming that exposure has been eliminated is equally important.

Rapid confirmation remains uncommon.

As previously mentioned, **only 11%** report that their organization patches or confirms exposure within 24 hours of receiving a CISA or ENISA known exploited vulnerability alert.

42%
report taking up to 7 days.

30%
require up to two weeks.

16%
report more than 2 weeks.

In many organizations, a vulnerability known to be actively exploited remains unverified for days or weeks. Approval workflows, competing priorities, maintenance windows, staffing constraints, and the initial effort required to determine exposure all affect timing.

When exploitation is active and confirmation is delayed, exposure remains. Attackers do not wait for maintenance windows.

KEV response is not only about deploying patches. It is about verifying that the vulnerability is no longer exploitable in the environment. The challenge is not simply how fast teams patch. It is how fast they confirm that risk has been reduced.

9. What Security Leaders Want to Automate

Validation and retesting lead the list

Security teams are not asking for more alerts. They are asking for clearer evidence that fixes work.

When practitioners were asked which workflow they most want to streamline:

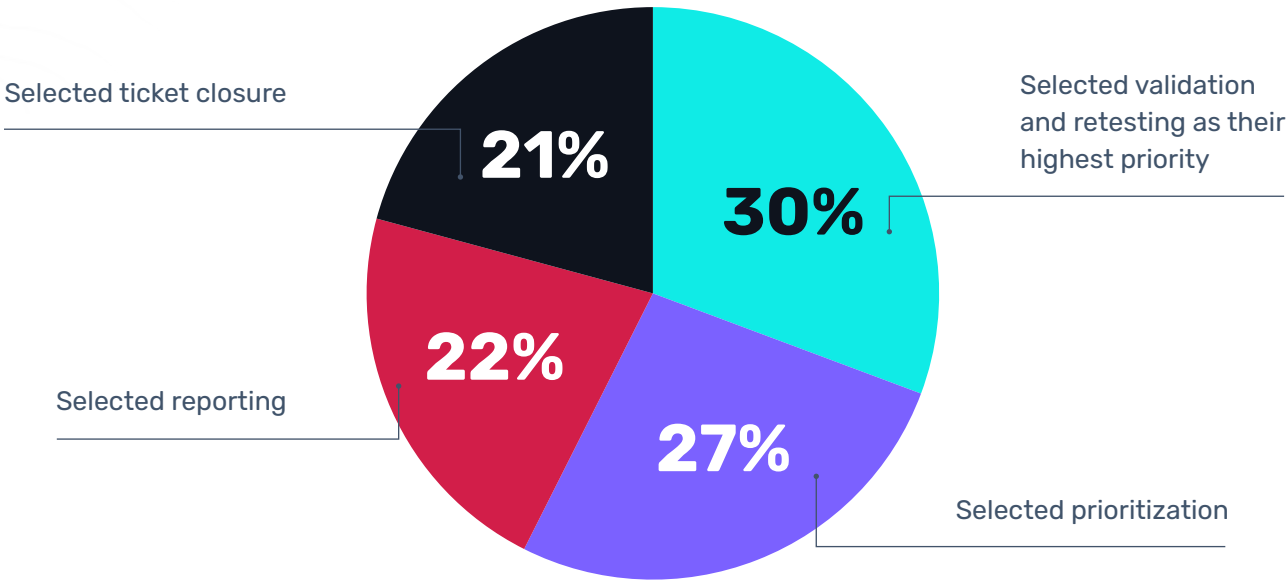

Validation and retesting ranked highest.

Teams are not primarily seeking additional findings. They want to know whether a vulnerability can be exploited, whether a fix eliminated that exploitability, and whether the exposure remains closed over time.

Executive responses point in the same direction. CISOs identified automated patch verification, prioritization, and reporting as areas where efficiency gains would be most valuable. All three are tied to demonstrating that remediation actions reduce exposure.

Notably, 70% of CISOs agree that integrating vulnerability scanner data with adversarial testing results would improve prioritization decisions. The demand signal is clear. Organizations are looking for ways to connect findings with confirmation.

Many workflows still assume closure based on scans, version checks, or ticket completion. Retesting remains inconsistent and often manual. Under operational pressure, confirmation is frequently deferred. The pattern is consistent. Demand for automation now centers on validation that confirms exposure has actually been reduced.

10. The Structural Confidence Gap

Five patterns that reinforce assumed security

When viewed together, the findings describe a structural gap between how security programs measure progress and how often they confirm real resistance to attacker behavior.

Across roles, regions, and maturity levels, five reinforcing patterns emerge.

1. **Confidence outpaces testing**
   CISOs report high assurance in prevention and detection capabilities. Routine adversarial testing, exploit confirmation, and embedded validation occur far less frequently.

2. **Severity guides remediation more than exploitability**
   Severity scores and version checks drive prioritization and closure. Direct confirmation that an attacker can no longer succeed is less consistent.

3. **Detection is trusted more than it is measured**
   Detection platforms are widely deployed. True-positive rates, noise levels, and limited stress testing suggest that interruption of real attack paths is not routinely validated.

4. **Automation increases speed faster than verification**
   AI-driven systems now influence prioritization and remediation. Independent confirmation of their decisions remains limited.

5. **Metrics emphasize closure over resistance**
   MTTR, SLA adherence, and ticket velocity are widely tracked. Fewer programs systematically measure whether exploit paths have been eliminated or whether exposure reappears.

Individually, each pattern reflects a reasonable operational choice. Together, they reinforce a system that rewards activity more visibly than resistance. This is what defines assumed security.

Assumed security is not negligence. It is what happens when measurement favors activity over resistance. But mature security processes do not automatically produce demonstrated resilience.

When attackers reuse credentials, pivot across environments, and chain weaknesses, confirmation becomes the differentiator. The next phase of cybersecurity maturity will not depend on additional tools or faster ticket closure. It will depend on how consistently organizations test what they trust.

11. What This Means for Security Leaders

From assumed security to *demonstrated resilience*

The findings in this report do not point to a lack of effort. Security programs are active, instrumented, and increasingly automated. The gap lies in confirmation. Security maturity depends on how clearly organizations can demonstrate that their actions reduce real exposure.

Moving from assumed security to demonstrated resilience requires deliberate shifts in behavior.

Test controls as routinely as you deploy them

Deployment does not equal assurance.

If a control is trusted, it should be tested at the same frequency as that trust.

Endpoint protection, identity monitoring, segmentation policies, and remediation workflows should be tested at a cadence aligned with the confidence placed in them. If leadership expresses high assurance in a control, testing frequency should reflect that assurance.

A once-a-year red team exercise is not the same as routinely testing whether controls stop real attacks.

Prioritize exploit conditions over severity alone

Severity scores support triage.

They do not confirm whether an attacker can succeed.

Security leaders should ask: how do we know we are at risk, and how do we know this condition can no longer be exploited in our environment?

Closing a ticket is not the same as closing an attack path.

Measure detection timing, not just alert volume

Alert counts and response times show activity.

They do not show interruption.

Effectiveness is determined by whether detection occurs before lateral movement, privilege escalation, or credential abuse succeed.

Detection timing should be measured and tested under realistic scenarios.

Pair acceleration with verification

AI-driven prioritization and automated remediation increase speed.

Speed must be paired with independent confirmation.

Automated decisions should trigger retesting. Efficiency without validation leaves exposure uncertain.

Align executive reporting with attack reality

Dashboards that report MTTR, SLA adherence, and ticket closure rates demonstrate responsiveness.

They do not demonstrate whether exposure has been removed.

Security programs are mature. The opportunity lies in how consistently they verify outcomes. Organizations that test what they trust will outperform those that rely on process completion as evidence.

Confidence does not reduce exposure. *Testing* does.

12. Conclusion

Security programs today are more capable, more instrumented, and more active than at any point in the past. The question facing leaders is not whether work is being done, but whether that work consistently reduces real exposure. Organizations that align measurement with confirmation, and activity with demonstrated resistance, will operate with a clearer understanding of their risk.

In an environment where attackers adapt quickly and chain small weaknesses into meaningful impact, the ability to show that defenses hold under realistic conditions becomes a defining characteristic of mature security leadership.

13. Methodology

Survey scope and demographics

This research is based on survey data collected from cybersecurity leaders and practitioners across the United States and Europe. Sample size: the survey included 750 cybersecurity professionals, evenly segmented across two audiences:

375

Chief Information Security Officers and senior security leaders

375

Security practitioners, including engineers, analysts, and operations personnel

This structure enables direct comparison between executive perception and operational practice.

Regions

Respondents were drawn from the **United States, United Kingdom, France, Germany, Spain, and Italy.** The sample reflects major enterprise markets across North America and Europe.

Roles

Participants were grouped into two categories:

- **CISOs and senior security leaders**
  Responsible for strategic direction, executive reporting, and risk governance
- **Security practitioners**
  Responsible for vulnerability management, detection and response, identity security, and remediation workflows

This design supports comparison between leadership perspective and operational experience.

About the survey

Survey data for this report was collected by Censuswide, an independent research firm, from 750 cybersecurity leaders and practitioners across the United States and Europe.

Field dates

Data was collected during November 2025. All responses were gathered within the same field window.

Margin considerations

Findings reflect self-reported perceptions and practices at the time of response. Results are subject to standard sampling variability. Percentages may not total 100% due to rounding. This research measures reported behavior and stated confidence. It does not independently validate technical outcomes. All findings in this report are derived from participant responses at the time of collection.