Futurum®

N-ABLE

Cybersecurity in the Age of AI: Moving from Fragile to Resilient

FEBRUARY 2026

Executive Summary

The traditional paradigm of cybersecurity is reaching a breaking point. For decades, organizations have relied on a reactive playbook designed for human-speed threats and definable perimeters. However, the rapid integration of artificial intelligence (AI) into the digital ecosystem has fundamentally altered the landscape. AI is now a dual-use technology. While it drives legitimate business efficiency, it also accelerates adversarial tradecraft, granting threat actors the ability to scale attacks with unprecedented speed and sophistication.

To survive this shift, leaders must recognize that in a digital-first economy, cyber resilience is business resilience. It is no longer enough to defend the network. The goal must be to ensure the business can withstand, operate through, and recover from an attack with minimal disruption to revenue and reputation. This requires a transition from fragmented, alert-driven security to a unified program that manages risk across the entire lifecycle of a threat.

This paper outlines a modern framework for business resilience built on three core pillars:

- **Minimizing Exposure:** Proactively managing the attack surface to deny adversaries entry before an alert is ever triggered.
- **Reducing Impact:** Utilizing automated systems to contain threats the moment they bypass the perimeter, preventing minor compromises from becoming major outages.
- **Maintaining Continuity:** Ensuring rapid, verified recovery of operations so the business survives the event.

Central to this strategy is the intelligent application of AI. Rather than treating AI as a generic solution, success depends on using the right technological capabilities – generative AI for interpretation and deterministic AI for action – to balance the speed of automation with the safety of human oversight. This approach provides the blueprint for moving from a posture of fragility to one of active, enduring resilience.

Section 1: THE WHY

THE WHY: The New AI-Driven Threat Landscape

No conversation about the current or future state of technology can avoid the impact of AI; business leaders across every sector are integrating AI to drive operational efficiency. Yet this technology is dual-use. The exact mechanisms driving legitimate productivity are increasingly accelerating adversarial tradecraft.

For small to mid-sized organizations, this creates a distinct vulnerability regarding the balance of expertise versus available resources. Security teams in this segment are typically composed of capable generalists responsible for infrastructure, compliance, and defense simultaneously. It is not realistic to expect that they can match the depth of specialized knowledge maintained by large-scale enterprise SOCs. AI weaponizes this disparity. It grants adversaries easier access to the niche and technical specificity required to exploit edge-case vulnerabilities.

The challenge is compounded by the organization's own adoption of AI. As business units deploy these tools to compete, the IT environment becomes denser and more complex. Security teams must now account for:

- **A Web of External Services:** New APIs and third-party connections that expand the perimeter.
- **Proprietary Data Risk:** Sensitive internal data being fed into public or semi-public models.
- **"Shadow AI":** The inevitable rise of unmanaged tools where employees bypass governance to get the work done.

The attack surface is not just expanding; it is becoming more opaque.

Simultaneously, the threat horizon is shifting toward industrialization. While current attack volumes may not yet reflect the full potential of these tools, the trajectory is unambiguous. AI enables the automation of reconnaissance, social engineering, and, in the future, intelligent interaction, at a scale previously cost-prohibitive. We are moving toward a future where adversaries use machine intelligence to saturate the perimeter, exploiting the widening gap between an increasingly complex internal environment and the finite capacity of human defenders.

Emerging Trends: A substantial percentage of small to mid-size organizations report observing a marked increase in sophisticated, AI-driven social engineering attacks over the last 12 months.

Source: Futurum Research, December 2025, 2H 2025 Cybersecurity Global Enterprise Decision Maker Survey Report

THE WHY: Why the Old Playbook Is Unprepared

The strategic implication of this shift is that the traditional security posture is becoming functionally obsolete. Historically characterized by fragmented tools and reactive monitoring, this approach relies heavily on manual intervention. This "old playbook" was designed for an era of definable perimeters and human-speed attacks. It creates a structural mismatch against the velocity of machine-driven threats.

Reliance on manual detection and response introduces "fragility" into the organization. In this context, fragility is defined by the dangerous latency between the moment a breach occurs and the moment full recovery is accomplished.

Adversaries utilizing automation can move laterally within minutes. Any strategy that depends on human speed to sift through alerts and initiate a response leaves the business exposed. This exposure is particularly acute for small to-mid-sized organizations, where the environment creates specific operational friction:

- **The Skills Gap:** Teams are often staffed by capable generalists rather than specialized threat hunters. In a traditional model, they are forced to manually validate alerts they may not have the depth of expertise to interpret immediately.
- **The Visibility Gap:** The modern IT environment spans on-premises, cloud, and SaaS applications. This complexity forces teams to question if they possess the correct telemetry to detect an incident in the first place.
- **The Signal-to-Noise Ratio:** As complexity grows, the signals that do come through become unmanageable, leading to critical warning signs getting lost in the mix.

"Alert fatigue" is already a critical operational failure point, even before fully AI-driven attack waves arrive. Security teams are currently drowning in noise from false positives and low-fidelity warnings.

Continuing with a reactive posture is not merely a risk, It is a decision to remain fragile in an increasingly hostile environment. The required shift is fundamental: from reactive security that waits for alerts to preemptive security that eliminates vulnerabilities before they can be exploited.

The Risk

Figure 2. Incident Volume: Statistics show that a substantial portion of SMBs continue to experience three or more significant security incidents annually.

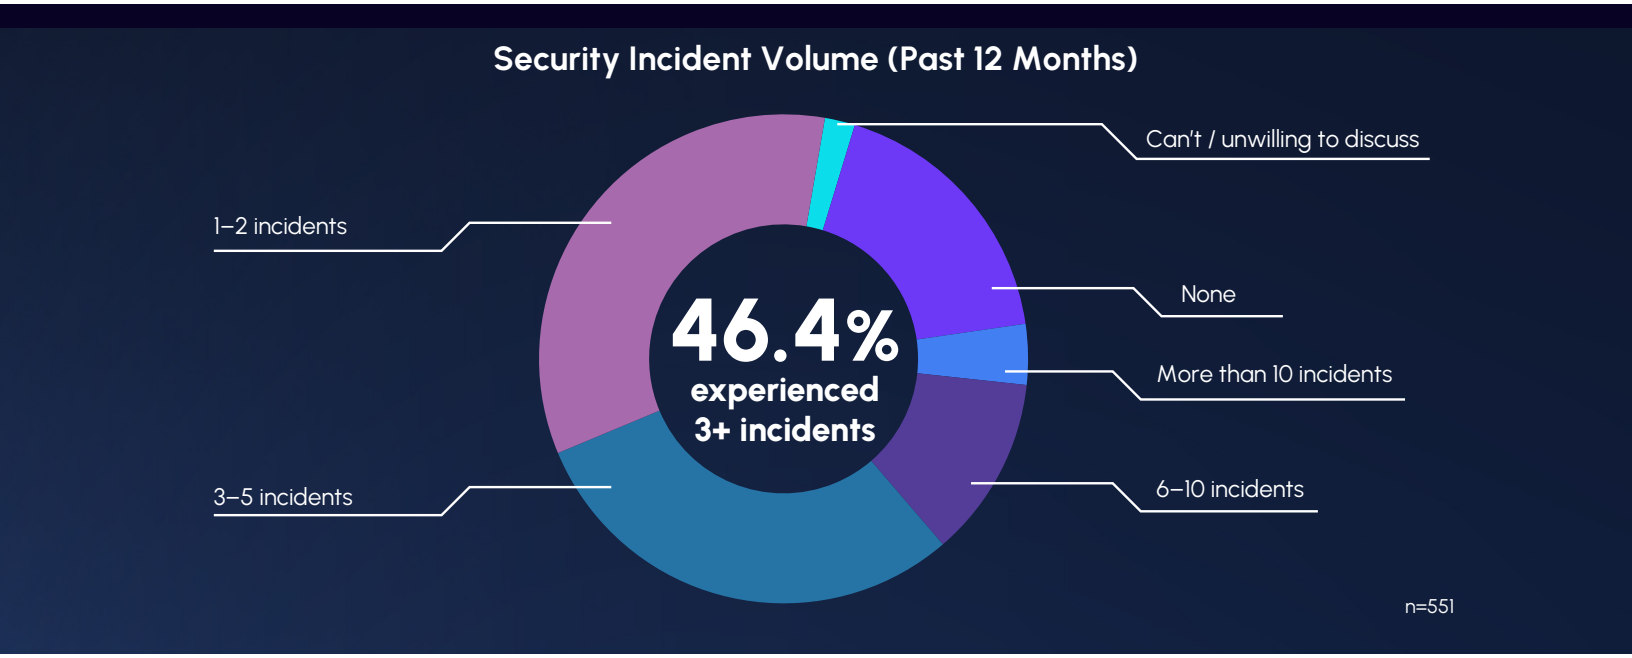

Source: Futurum Research, December 2025, 2H 2025 Cybersecurity Global Enterprise Decision Maker Survey Report

The Impact

Figure 3. Operational Impact: Data loss and "system/operational downtime" are key issues identified by organizations, which directly impact revenue and productivity.

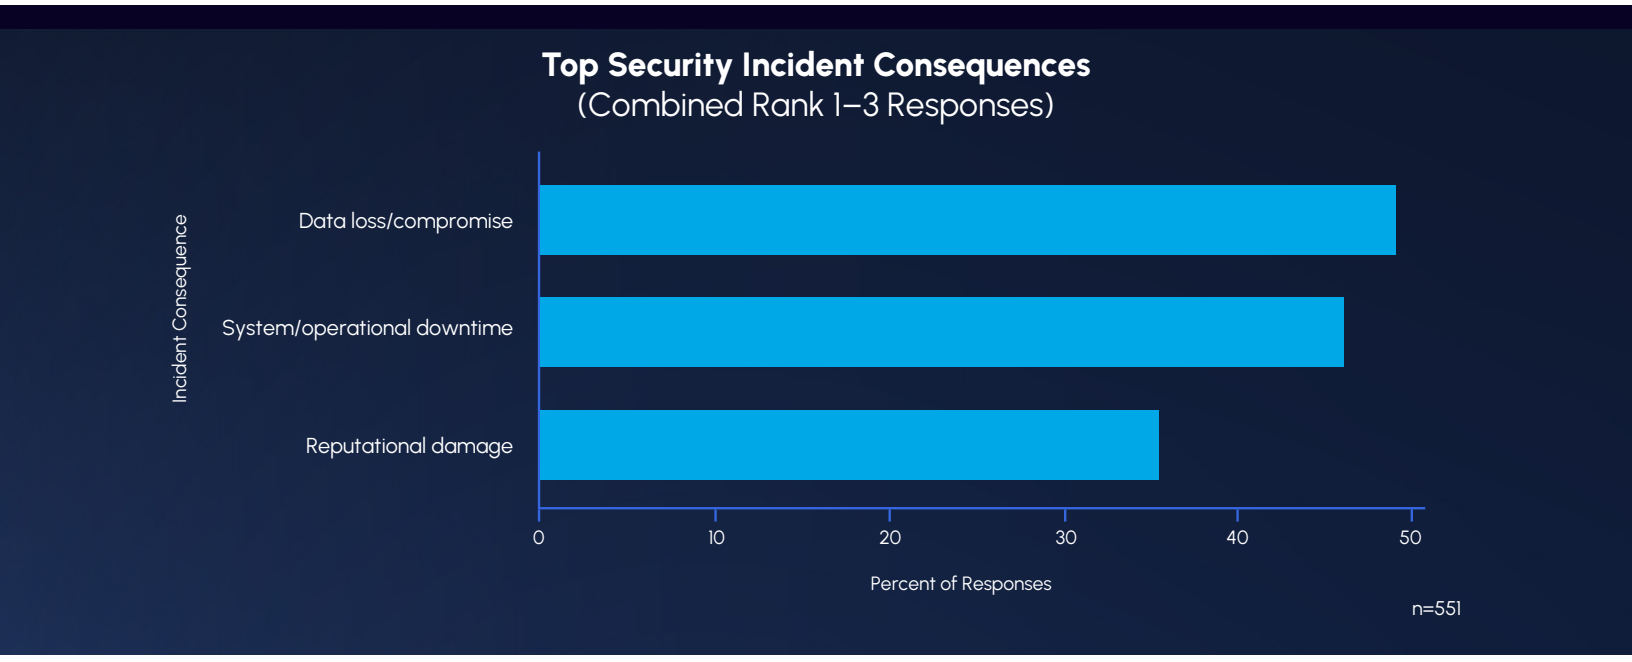

Source: Futurum Research, December 2025, 2H 2025 Cybersecurity Global Enterprise Decision Maker Survey Report

The Need for Augmentation

Figure 4. Operational Strain: "Alert fatigue" and "skills shortage" remain the top challenges for SOC teams.

Q: What do you consider to be the key challenges that your organization faces with respect to risk management and your security operations center (SOC)?

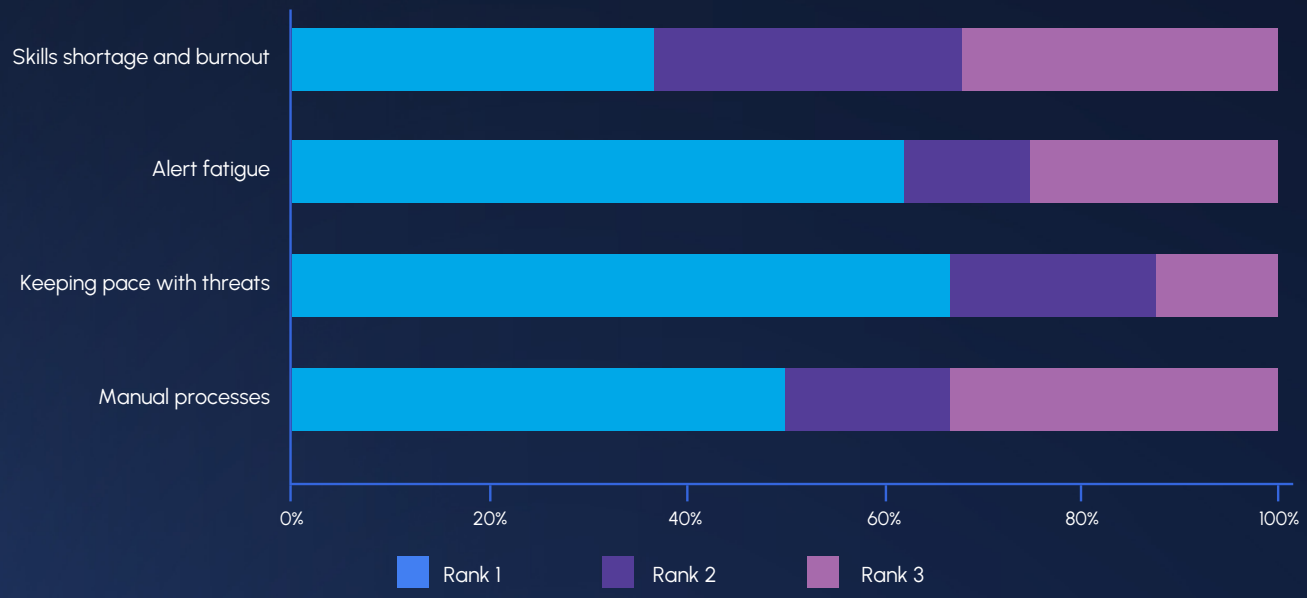

Q: What are your key hurdles to resolving security vulnerabilities and incidents?

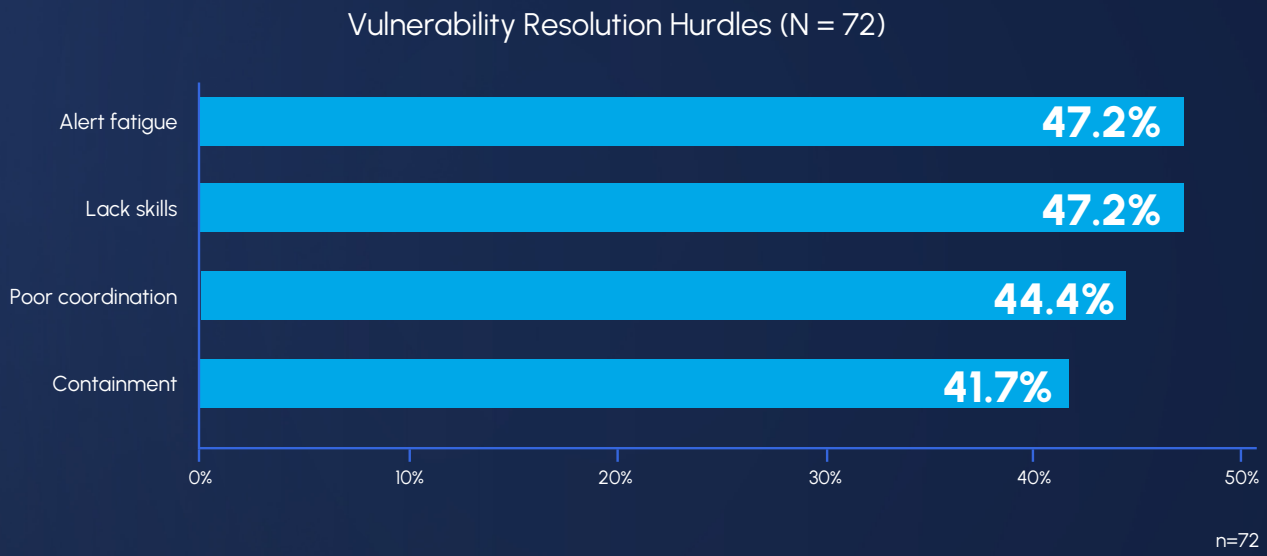

Source: Futurum Research, December 2025, 2H 2025 Cybersecurity Global Enterprise Decision Maker Survey Report

Figure 5. Visibility Deficit: Data highlights show that a low percentage of "Medium" and "Low" priority alerts are actually investigated due to resource constraints.

Source: Futurum Research, December 2025, 2H 2025 Cybersecurity Global Enterprise Decision Maker Survey Report

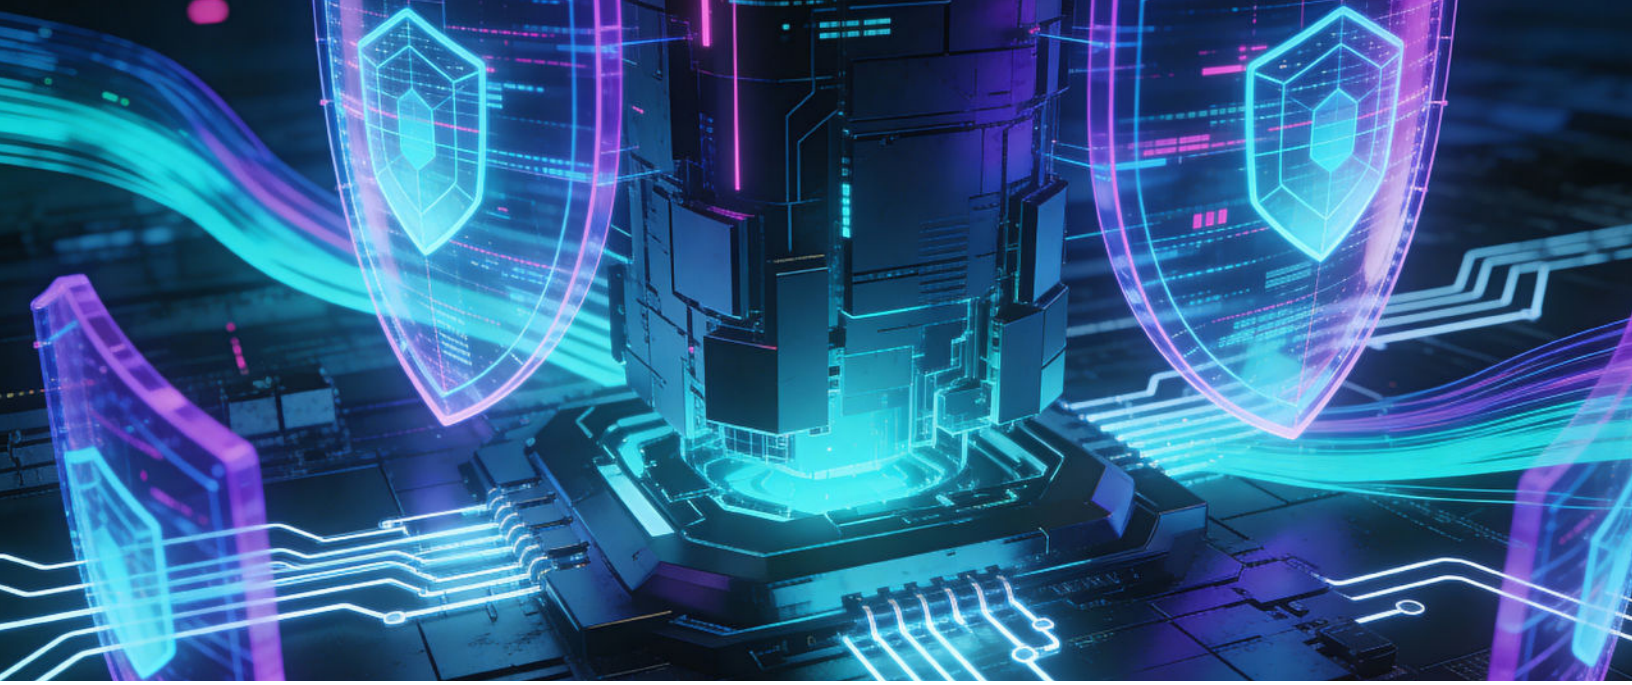

Section 2: THE HOW

THE HOW: A New Framework for Resilience

Escaping the cycle of fragility requires a strategic pivot. Reliance on isolated security tools often leads to gaps in visibility and control, resulting in a disjointed defense that adversaries can easily exploit. A promise of resilience is incomplete if it addresses only a single phase of an attack. While preventing intrusion at the perimeter remains necessary, it is insufficient unless paired with a robust capability to detect and respond when that perimeter is breached.

Crucially, there is no single implementation of resilience that fits every business. Each organization operates with unique resource constraints and specific threat realities. A framework must be adaptable. It should enable the security team to develop a program that aligns with their actual operational environment, rather than forcing them to adopt a rigid structure designed for a different type of enterprise. The goal is to move away from disconnected point products toward a unified program that scales with the organization.

This adaptable strategy is built upon three operational pillars:

1. **Endpoint Resilience - Minimize Exposure (Before):** This pillar focuses on proactive hardening and security measures. It involves strictly managing the attack surface to prevent entry. By ensuring devices are patched, configured correctly, and compliant, the organization reduces the sheer surface area available for an attacker to target.
2. **Security Resilience - Reduce Impact (During):** This pillar focuses on active defense. When an attacker bypasses the perimeter, the system must utilize automated defenses to identify and neutralize threats in real-time. The priority is preventing lateral movement to contain the damage, including any data theft/exfiltration before further deployment of ransomware.
3. **Data Resilience - Maintain Continuity (After):** This pillar focuses on recovery. It ensures that, in the event of a breach, the business can restore operations immediately using verified data. This renders the attack an operational inconvenience rather than a catastrophe. Executing this framework effectively requires speed. Manual processes cannot keep pace with these three pillars quickly enough to match modern threat velocity. To close the gap, defenders need a force multiplier that allows them to operate with the same efficiency as their adversaries.

THE HOW: The Guiding Principle – Use the Right AI for the Job

In the public imagination, "AI" has primarily become synonymous with generative AI. While the surge of interest in Large Language Models (LLMs) is justified, effective cyber resilience relies on a broader definition that includes multiple disciplines. To utilize these tools safely, organizations must distinguish between two fundamental types of intelligence.

- **Generative AI:** This technology is probabilistic. It excels at synthesis and creation, such as interpreting natural language, summarizing vast datasets, or drafting potential remediation scripts. It serves as the engine of creativity and interpretation.
- **Deterministic AI:** This technology includes symbolic logic and rules-based engines built for verification. It operates on strict, binary logic paths where the outcome is predictable and reproducible. It serves as the engine of safety and control.

Layered Intelligence

The most sophisticated security offerings do not choose between these technologies. They layer them to balance speed with safety.

- **The Workflow:** In an incident response scenario, generative AI might be tasked with generating a dynamic playbook or writing a script to remediate a specific vulnerability.
- **The Guardrail:** Executing that script blindly introduces risk. This is where deterministic AI intervenes. It analyzes the proposed code against strict safety policies to ensure it will not crash a critical server or disrupt legitimate traffic.

This interplay allows for a nuanced "human-in-the-loop" strategy. Generative AI handles the heavy lifting of interpretation and drafting to reduce the cognitive load on analysts, while deterministic models provide the automated validation required to ensure safety. This combination allows organizations to adopt AI-driven automation without ceding control over the outcome.

Section 3: THE WHAT

THE WHAT: Before the Attack (Pillar 1: Minimize Risk)

Effective resilience begins before the first alert is triggered. The "Before" phase of a robust security program is defined by the rigorous management of the attack surface, a task that has largely outgrown human scale. In a modern environment, the attack surface is constantly evolving. It expands as devices connect and degrades as configurations drift. For mid-market IT teams, the challenge is not just identifying these gaps but prioritizing them.

AI serves as a critical augmentation tool here, but its value depends on using the right type of intelligence for the specific task.

- First, Generative AI can transform how IT staff interact with their environment, broadly acting as a force multiplier for the human analyst. Instead of manually clicking through dashboards, a technician can use natural language to ask questions such as "Show me all devices that have drifted from the compliance baseline" or "List endpoints missing the critical patch released Tuesday." Beyond simple queries, Generative AI can draft both executive-level communications to help share information as well as remediation scripts to fix these specific gaps. These capabilities help generalist staff to operate with much more efficiency.
- Second, Deterministic AI handles the heavy lifting of predictive risk scoring. While Generative AI handles language, deterministic, often graph-enabled, models excel at correlating vast datasets to identify multidimensional risks that a human might miss. A missing patch on its own might be classified as a low priority. However, these models can correlate that missing patch with a disabled security agent and a spike in CPU usage at 2 am to flag a device as "Critical Risk." This predictive scoring allows teams to address the most dangerous vulnerabilities first, rather than drowning in a flat list of problems.

Finally, these technologies create a continuous feedback loop. Prevention is no longer a static setup. Insights gained during an active incident, such as a specific file hash or behavioral trigger identified in the "During" phase, are fed back into the prevention model. The system effectively "immunizes" the unimpacted fleet against the new threat ensuring that the "Before" state is constantly hardening based on real-world intelligence.

The Hygiene Challenge

Figure 6. Top issues for managing endpoints

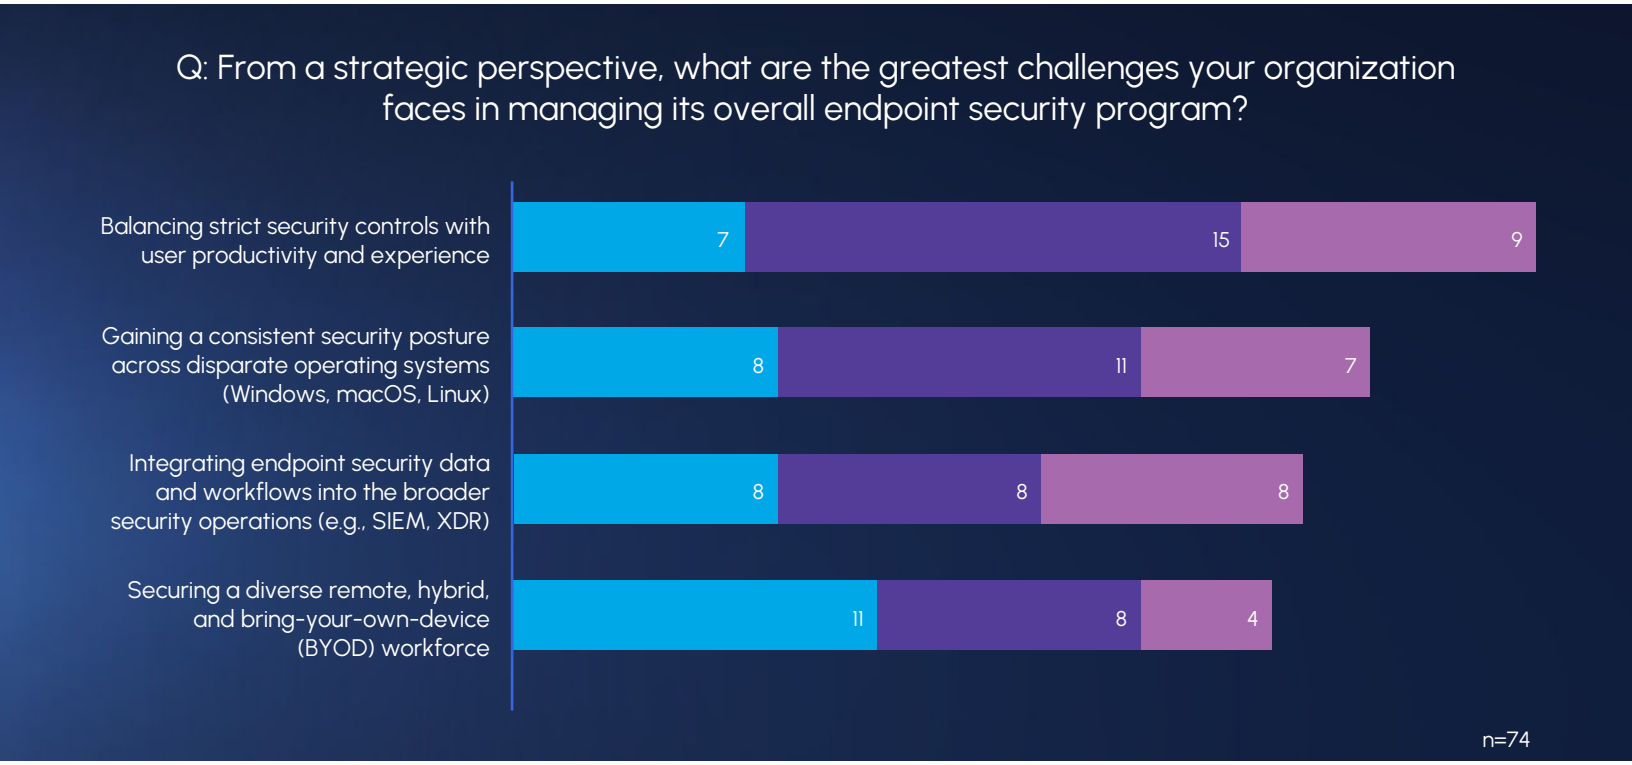

Source: Futurum Research, December 2025, 2H 2025 Cybersecurity Global Enterprise Decision Maker Survey Report

The Visibility Gap

Figure 7. Current Monitoring Coverage Across the IT Environment

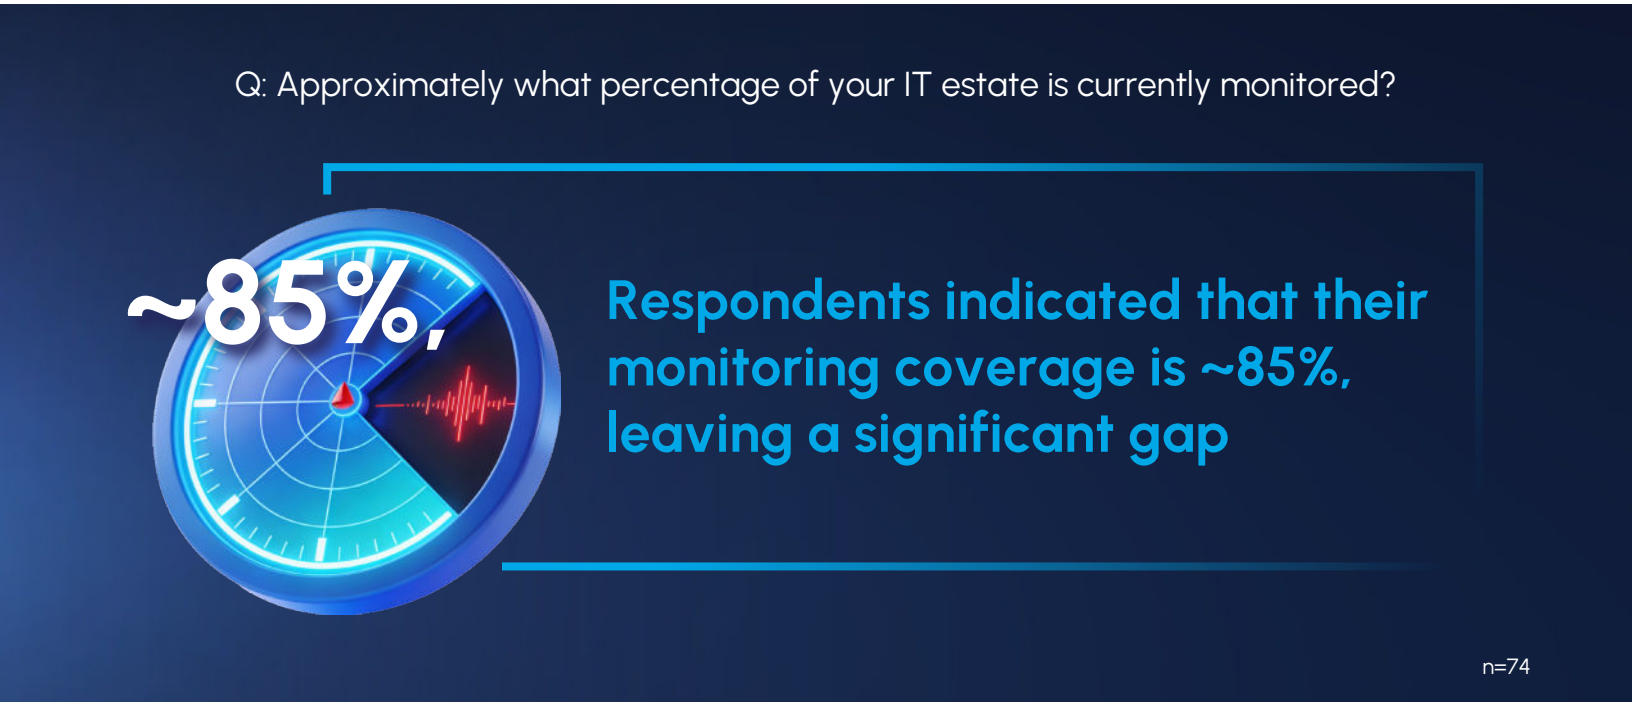

Source: Futurum Research, December 2025, 2H 2025 Cybersecurity Global Enterprise Decision Maker Survey Report

THE WHAT: During the Attack (Pillar 2: Reduce Impact)

When an adversary breaches the perimeter, the strategic objective shifts immediately from avoidance to containment. In this phase, time is the governing variable. The window between initial compromise and lateral movement is steadily decreasing, and it often closes faster than human analysts can react.

Defenders face a compound challenge that creates a dangerous advantage for the attacker.

- The "Fog of War": Security teams often rely on disjointed consoles that lack shared context. This forces analysts to correlate data manually while the attack spreads, introducing "swivel-chair" latency that leaves the business exposed.
- Evolved Tradecraft: Attackers increasingly utilize "living off the land" techniques. By abusing legitimate system tools such as PowerShell, they blend in with regular administrative activity to bypass traditional detection.

The Role of Behavioral AI and Trusted Automation

To counter this, the defense must move beyond static signatures to high-fidelity behavioral monitoring. The requirement is an AI model capable of distinguishing between a legitimate administrator making a change and an adversary executing a similar command with malicious intent.

- **Contextual Intelligence:** Detection without context creates noise. Detection with context creates actionable intelligence. This nuance is vital for identifying threats that hide in plain sight.
- **Trusted Automation:** While human oversight remains necessary for complex investigations, the initial containment of a verified threat cannot wait for a ticket to be opened.

The system requires the autonomy to execute immediate, precise responses based on high-confidence triggers. Whether severing a specific network connection or isolating a host, these actions must occur at machine speed. By neutralizing the threat instantly, the organization prevents a single compromised endpoint from becoming a systemic business event. The goal is not just to stop the attack but to render it manageable.

THE WHAT: After the Attack (Pillar 3: Maintain Continuity)

True resilience requires the uncomfortable admission that a sophisticated adversary may eventually succeed. In that scenario, the definition of success shifts from prevention to the velocity of recovery.

For small-to mid-sized organizations, the recovery phase is often paralyzed by uncertainty. IT teams face "restore anxiety," or the fear that their backups may be corrupt, incomplete, or infected with the same ransomware that took down production. This uncertainty often forces a slow, tentative recovery, extending downtime and losses.

AI does not solve the breach itself, but it acts as a critical force multiplier in mitigating this uncertainty. It shifts the focus from hoping for data integrity to verifying it.

Before the incident, AI-enabled automation already supported data integrity through anomaly detection. By continuously analyzing backup streams for irregularities, such as unexpected spikes in file entropy or mass deletions, intelligent workflows that use more advanced reasoning serve as an early warning system. They flag potential corruption or attempts to encrypt the backup repository. This allows administrators to isolate compromised snapshots before they are needed for restoration.

Second, AI-powered workflows can significantly improve recovery validation activities. Historically, disaster recovery testing was a manual, infrequent "fire drill." AI enables this to potentially run continuously in the background. Systems can,

for example, autonomously mount backups in a secure sandbox to verify that critical applications boot and run correctly. This moves recovery planning from a theoretical exercise to a tested capability.

When a breach occurs, these tools provide the IT team with data-driven clarity rather than blind faith. By combining AI-assisted validation with resilient storage, organizations can turn recovery from a chaotic scramble into a structured, evidence-based operation.

WHAT'S NEXT: The Vision of Active Resilience

The framework described previously represents the current best practice. However, the trajectory of AI suggests a more profound architectural shift is on the horizon. We are moving toward an era where AI doesn't just assist human defenders from the cloud, but enables a network of self-defending devices at the edge.

Currently, most security models rely on a hub-and-spoke architecture: endpoints send telemetry to a central "brain" for analysis. While effective, this introduces latency. The future vision of "Active Resilience" involves decentralized intelligence, where lighter, more efficient AI models reside directly on the endpoint. This allows for complex decision-making without needing constant connectivity to a central server.

In this developing model, devices would utilize agent-to-agent (A2A) protocols to collaborate. Imagine a scenario relevant to many SMBs: a workstation in the finance department executes a script that begins exhibiting early signs of ransomware behavior. In a traditional model, the device waits for cloud validation to occur. In an A2A model, that device immediately communicates the threat signature to its peers on the local network.

The environment effectively creates a digital immune system. Neighboring devices, seeing the warning, proactively update their defenses to block that specific process or isolate traffic from the infected machine. For an SMB with limited staff, this capability would be a significant force multiplier. It distributes the defensive workload across the fleet, allowing the network to react at machine speed to contain lateral movement before an analyst even opens a ticket.

The New Why: Taking Ownership of the Resilience Program

For too long, small-to mid-sized organizations have operated under a narrative of disadvantage. They are told they lack the budget, the staff, and the specialized expertise to defend against sophisticated adversaries. This narrative breeds passivity. It fosters a mindset where security is viewed as a commodity to be purchased – a collection of disconnected tools – rather than a capability to be managed effectively.

Moving from "fragile" to "resilient" requires a fundamental shift in ownership. IT leaders must reframe their goal. The objective is not merely to acquire more software to patch yesterday's problems, but to build a modern resilience program that actively manages risk across the entire lifecycle of a threat. To achieve this, organizations should prioritize three decisive steps:

- **Own the Strategy:** Resilience cannot be achieved through a reactive, alert-driven model that leads to burnout. A successful program is built on three pillars: **minimize exposure** to shrink the attack surface, **reduce impact** to contain threats at machine speed, and **maintain continuity** to ensure recovery.
- **Demand the Right "How":** Security teams must move beyond accepting "AI" as a generic buzzword and interrogate how it is applied. The priority is to insist on the right tool for the job: Generative AI to interpret data and upskill teams, and Deterministic AI to execute critical, autonomous actions with safety and precision.
- **Modernize the Foundation:** Complexity is the enemy of security, but consolidation should not mean compromise. Rather than accepting "walled gardens" that force the replacement of functioning infrastructure, organizations should demand an **open resilience platform**. A modern security core must be capable of integrating with existing trusted investments, such as firewalls, identity providers, and email gateways. By ingesting telemetry from every corner of the environment into a single, open ecosystem, defenders eliminate the blind spots where adversaries hide.

Resilience is not a product feature. It is a strategic outcome. It is entirely within reach for the mid-market, provided the focus shifts from manual reaction to automated, intelligent resilience.

Important Information About This Report

AUTHORS

Fernando Montenegro
Vice President & Practice Lead,
Cybersecurity & Resilience | The Futurum Group

PUBLISHER

Futurum Research

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

N-ABLE

ABOUT N-ABLE

N-able protects businesses from evolving cyberthreats. Our AI-powered cybersecurity platform delivers business resilience to more than 500,000 organizations worldwide, leveraging advanced end-to-end capabilities, simplified workflows, market-leading integrations, and flexible deployment options to improve efficiency and drive critical security outcomes. Our partner-first approach pairs our technology with experts, training, and peer-led events that empower customers to be secure, resilient, and successful. n-able.com

Futurum®

ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

Futurum®
CONTACT INFORMATION: The Futurum Group LLC | futurumgroup.com | (833) 722-5337