


THE 7 DEADLY SINS OF CYBERSECURITY

2026 Cyber Protect Report

Every vendor publishes a Threat Report. We publish a Protect Report — because our job isn't to catalog what went wrong. It's to help you prevent it.



A Message from Michael Crean

SVP and GM of SonicWall Managed Security Services

Small and mid-market businesses (SMBs) are the backbone of the United States economy. They represent 99% of all US businesses and nearly half of private sector employment while contributing roughly 44% of GDP. What they may not know is that they are facing the same cyber risks as large enterprises; however, they lack the same levels of expertise, budget or resources. For SMBs, cybersecurity is no longer a technical concern. It is a business necessity.

The consequences of security gaps are now measured in downtime, missed paychecks, lost customers, public embarrassment, insurance claims denied for lack of adequate defenses, and reputational damage that can threaten the very existence of a business. Unlike large enterprises that can absorb and recover from major incidents, most SMBs cannot survive a serious breach.

Today, the most common cause of security incidents is not advanced malware or exotic attacks. It is security misconfiguration. Security misconfiguration is a rapidly escalating risk, rising from #6 in the Open Worldwide Application Security Project (OWASP) Top 10 in 2017 up to #2 in 2025—highlighting how foundational weaknesses are becoming one of the most critical drivers of cyber exposure.

Across the environments we monitor and the incidents we investigate, the same patterns surface repeatedly: exposed services, unmonitored access, inconsistent policies and gaps that grew quietly over time. These are rarely the result of negligence. They are the inevitable result of complex environments, tight timelines and stretched teams trying to keep up with a threat landscape that never slows down.

The threat landscape is also shifting in ways that demand attention. Nation-state actors increased their targeting of SMBs and mid-market organizations throughout 2025, recognizing that smaller organizations often serve as entry points into larger supply chains and critical infrastructure. These are no longer threats reserved for governments and large enterprises.

Compounding the risk further, AI is accelerating threat actors' ability to automatically scan for weaknesses at a scale and speed that manual attackers could never achieve, rapidly identifying exposed services, overly

permissive access and administrative gaps across thousands of targets simultaneously.

In preparing this year's report, we found ourselves returning to a familiar truth: the organizations that suffer the most are not failing because of sophisticated attacks. They are failing because of predictable, preventable gaps. We came to think of these as the seven deadly sins of cybersecurity; seven patterns that we see repeatedly across breach investigations, security assessments and incident reviews. They are not obscure vulnerabilities. They are operational failures hiding in plain sight.

This report is not focused on threats alone. It is focused on protection outcomes that matter to business leaders. It is designed to help SMBs and the partners who protect them understand what really keeps them operational, resilient and trusted in their markets. Rather than focusing solely on the threats organizations face, we made a deliberate decision to reframe our annual research around protection outcomes. That is why we are proud to introduce the **SonicWall 2026 Cyber Protect Report**.

At SonicWall, we deeply believe that partners deliver the best security outcomes. SMBs should not go it alone. MSPs and MSSPs play a critical role in delivering protection at scale, and this report is designed to equip them with the language and data they need for strategic conversations with decision makers.

Protecting systems, networks and data ultimately protects people, businesses and entire communities. It enables organizations to pay their employees on time, serve their customers reliably and grow without fear of disruption. The window to act is open. SonicWall is honored to help SMBs and their partners close the gaps before attackers find them, because the businesses that keep our economy moving and our communities strong deserve nothing less.











Michael Crean
SVP and GM of Managed Security Services
SonicWall

Executive Summary

The data tells a consistent story: attackers are getting more precise while most organizations are still closing yesterday's gaps.

HITS: The total number of times a threat or rule was triggered. This is the raw count of detections, including repeated activity from the same source.

TICKS: The number of distinct time intervals (or unique events) in which a threat was observed. This measures how often or how consistently a threat appeared over time, rather than just the volume.



KEY TAKEAWAY

Attackers are working smarter, not just harder, so SMBs can't ignore even "low-diversity" attacks.

By the Numbers

13B

High/Medium (HM) Intrusion Prevention System (IPS) hits increased to **13B** (+20.8%).

Exploitation attempts targeting high and medium severity vulnerabilities increased 20.8% in 2025. Total intrusion attempts stayed flat. The number of serious, actionable attacks increased sharply, meaning attackers are getting more precise, and defenders have less room for error.

THE KEY INSIGHT

Attackers aren't attacking more — they're attacking better.

- | Same overall volume
- | Much higher **quality and intent**
- | More attacks that matter, fewer junk probes

THINK OF IT LIKE THIS...

Last year: lots of random door-knob jiggling

This year: fewer jiggles, more people trying real keys

Why this matters operationally (especially for SMBs)

ONE

Alert fatigue gets worse, not better. Even if total alerts are flat...

- teams still face **more real decisions**
- more chances to miss something critical

TWO

"Just respond faster" breaks down. If **1 in 5** alerts are serious...

- you can't just speed up humans
- prevention and safe defaults matter more

THREE

Noise filtering is now a security requirement. Security systems must...

- separate junk from danger
- block high-confidence threats automatically
- reduce human involvement where possible

One-sentence executive takeaway

Total intrusion attempts stayed flat, the number of serious, actionable attacks increased sharply — meaning attackers are getting more precise, and defenders have less room for error

.47%

Out of thousands of detection rules, .47% account for 80% of what customers actually see. Only a very small fraction is responsible for the vast majority of alerts. Fewer than 1% of rules generate most of the alert volume, and a single rule alone accounts for billions of detections.

A small number of signatures dominate what customers experience.

This means that out of thousands of detection rules, only a few are responsible for the majority of the alerts or hits that customers actually see.

Why this matters

Operational impact: Security teams often focus on the alerts they see most. If a tiny fraction of signatures is responsible for almost all hits, it makes prioritization easier.

Performance: These high-volume signatures might strain the IPS or require tuning, because billions of hits could slow things down or generate alert fatigue.

Risk Management: The top signatures often target very common attack patterns (like malformed web requests), so they're critical to monitor.

Big picture

Think of it like a store where thousands of products are sold, but only a few products make up 80% of all sales — and one product alone is the single best-seller with billions of units sold. For practical purposes, most of the “action” is coming from a tiny subset.

825M

Even four years after it was discovered, the Log4j vulnerability was still targeted over 825 million times in 2025 - proving that “old” vulnerabilities continue to drive massive real-world attack traffic and remain operationally relevant.

THIS MEANS...

- | Even **4 years after the vulnerability was discovered**, attackers are still actively targeting it.
- | **825 million** hits is enormous—indicating there’s still a lot of automated scanning or exploitation attempts.
- | So even “old” vulnerabilities aren’t just historical—they continue to impact security operations.

It’s like a famous cheat code that was published decades ago. Players are still using it because the game never patched it out, and it still works every single time.

It’s a reminder that patching and monitoring aren’t one-time tasks—old vulnerabilities can continue to drive massive network traffic and alerts.

36,000

per second

Bad bot traffic alone has surged to 37% of all global internet traffic. Hackers don't pick targets individually, automated bots now generate over 36,000 vulnerability scans per second, making up more than half of all internet traffic as they relentlessly probe every public website for simple flaws like directory traversal and malformed requests.

THIS MEANS...

- | Hackers are constantly scanning the internet, looking for easy targets.
- | The attacks they use most often are **web-based and simple**, like:
 - Trying to access files they shouldn't (directory traversal)
 - Sending weird or broken website requests (malformed URLs) to crash servers or find weaknesses
- | These attacks are **automated** and happen all the time, not just once in a while.

The real threat to SMBs

Your website or web apps are the biggest target, especially if they're online and publicly accessible.

Even small or old vulnerabilities can be exploited if not patched.

THESE ATTACKS CAN LEAD TO

- Data leaks (sensitive files accessed)
- Service disruption (website crashes)
- Ransomware or malware infections if attackers find a way in

Basically, attackers don't need to specifically target your business—they just scan everything online and hit whatever is vulnerable.

11%

IoT hits rose +11.0% YoY and peaked in December, aligning with continued botnet and device exploitation activity throughout 2025.

THIS MEANS...

- | **IoT attacks are growing:** 609.9 million IPS hits in 2025, which is **11% more than the previous year.**
- | **Peak in December:** Attack activity isn't constant—it spikes at certain times.
- | **Botnets and exploitation:** Most of these attacks are automated, with compromised IoT devices being scanned, hacked, or added to botnets (networks of infected devices used to launch attacks).

Why this matters

IoT devices are easy targets: They often have weak passwords, outdated firmware, or default settings.

SMBs and households are vulnerable: Even a single exploited IoT device can give attackers a foothold in your network.

Botnet growth: Compromised IoT devices are often recruited into botnets, which are used for:

- Distributed Denial of Service (DDoS) attacks
- Spreading malware
- Proxying other attacks

Key takeaway for SMBs

SECURE ALL IOT DEVICES

Change default passwords, update firmware, and isolate them from critical networks.

EXPECT CONSTANT PROBING

Attackers are continuously scanning for vulnerable devices.

EVEN SMALL DEVICES CAN CREATE BIG RISKS

One insecure device can compromise your whole network.

The 7 Deadly Sins

Ignoring the Fundamentals

The security industry has spent most of its time worried about the scariest, most advanced threats, like attacks powered by artificial intelligence (AI), brand-new vulnerabilities that no one has seen before, and hacking operations run by nation-state threat actors. While that focus is not entirely wrong, it has obscured a more uncomfortable truth: the majority of successful breaches in 2025 didn't require any of that. They exploited what was already missing.

Weak authentication. Unpatched systems. Accounts with more access than they should ever have. These are not exotic

vulnerabilities, rather they are operational failures. The kind that accumulate quietly in busy environments where security competes with a hundred other priorities.

According to SonicWall data, identity, cloud and credential compromise account for 85% of actionable security alerts. That number tells the whole story: the attacker's preferred front door isn't a zero-day. It's a stolen password walking through an unguarded entrance.

Most breaches don't start with a sophisticated attack. They start with the basics left undone.

85%

of actionable security alerts involve identity, cloud, or credential compromise

SONICWALL 2025

61%

of exploits occur within 48 hours of proof-of-concept disclosure

SONICWALL 2025

66%

of SMBs globally have not implemented MFA at all

CYBER READINESS INSTITUTE

THE PATCH GAP

Attacker Speed vs.
Defender Response

~2 days

exploit released before disclosure

102 days

is the average time to patch a high-severity vulnerability

77%

of organizations need **7+ days to patch**

32%

of ransomware incidents in 2025 started with an exploited vulnerability

Weak or missing MFA

The single highest-return control in security — and the most commonly skipped by SMBs. Once credentials are stolen, MFA is the last line of defense.

Poor patch discipline

Attackers weaponize known vulnerabilities within hours. Most organizations take weeks or months to patch creating a wide, preventable window of exposure.

Excessive admin privileges

Default credentials and over-privileged accounts give attackers the ability to rewrite rules, disable logging, and pivot freely — often within minutes of initial access.

The Patch Window Problem

SonicWall found that **61% of exploits happen within 48 hours** of a vulnerability being made public. That acceleration reflects a fundamental shift: AI-assisted tooling, automated scanning and underground exploit markets have compressed the time from vulnerability disclosure to weaponized attack to a matter of hours.

The defender's timeline has not kept pace. **77% of organizations need more than a week to deploy patches** enterprise-wide, and **14% require more than four weeks.**¹ In financial services, the average time to patch a high-severity vulnerability is **102 days.**²

That gap, which is hours on the attacker's side (compared to weeks or months on the defender's) is not a technology problem. It is a process problem. Change management approvals, testing requirements, fear of disrupting production systems and limited staff all create bottlenecks that leave known vulnerabilities sitting open long after a fix exists.

According to the Sophos State of Ransomware 2025 report, **32% of ransomware incidents in 2025 started with an exploited vulnerability**, making it the single most common technical cause, ahead of compromised credentials (23%) and phishing (18%).

Default and Excessive Administrative Access

The third foundational failure is one of scope. When admin accounts have more access than they need, or when default passwords are never changed, a single compromised account can open everything. Attackers don't just get into one system. They get into all of them.

In investigations into SonicWall device compromises, attackers were able to pivot to domain controllers within hours of initial intrusion, not because of advanced techniques, but because administrative access was insufficiently hardened and credentials had never been rotated during infrastructure migrations.

The principle of least privilege is not a new idea. But in practice, it requires active governance: regular access reviews, role-based controls, prompt removal of stale accounts, and a deliberate decision that no account should carry more access than its function requires. In organizations where IT staff manage everything from endpoints to cloud infrastructure, that discipline is hard to maintain under pressure.

The consequence is predictable. When a single compromised account can reach every system, the blast radius of any breach becomes total.

The Bigger Reality

The fundamentals are not glamorous. They don't generate conference talks or vendor marketing. But they remain the primary attack surface, not because they are technically complex, but because they are operationally difficult to sustain.

Addressing them doesn't require new tools. It requires enforcement. Consistent MFA deployment with no exceptions. Patch workflows that prioritize speed on critical vulnerabilities. Access policies that are actively maintained rather than assumed. And platforms that make all of this visible, reportable, and auditable — so that gaps surface before attackers find them.

SonicWall's threat data makes clear that attackers today are leveraging automation and AI-driven attacks, making it clear that SMBs and organizations of all sizes can't fight this battle alone. The fundamentals are where that fight starts, and too often, where it is lost before it begins.



More than 3 in 5 cloud compromises traced to poor access management

Including **excessive permissions** and **weak access controls** — the leading enablers of cloud breach escalation.

¹ [Adaptiva](#): The 2025 State of Patch Management

² [Gitnux](#): Patch. Management Statistics

SMB EXPOSURE

A Perfect Storm

SMBs are where this problem hits hardest. For most SMBs, one or two people manage everything, and those people almost certainly have admin access to everything because it's easier that way.

That convenience is also a serious risk.

When a single admin account gets compromised, the attacker inherits whatever that account can touch. In an SMB where privileges are broadly shared and rarely reviewed, that usually means everything: email, file servers, cloud storage, backups, customer data.

Default credentials make it worse. Many SMBs deploy routers, firewalls, and software platforms and never change the factory passwords. Attackers know this, and automated tools scan for them constantly.

The fix isn't complicated. Restrict admin privileges to only what each person actually needs, change default credentials on every device at deployment, and enforce MFA on every elevated account. None of this requires expensive tools. It requires follow-through, which is exactly what gets squeezed out when IT is a one-person operation juggling a dozen other priorities.

SOC POV

In post-incident reviews involving compromised SMB environments, one pattern appears consistently: a single admin account was the entry point, and from there, the attacker moved without resistance.

Security teams investigating these cases reported the same finding:

"There was no lateral movement to speak of. They were already everywhere."

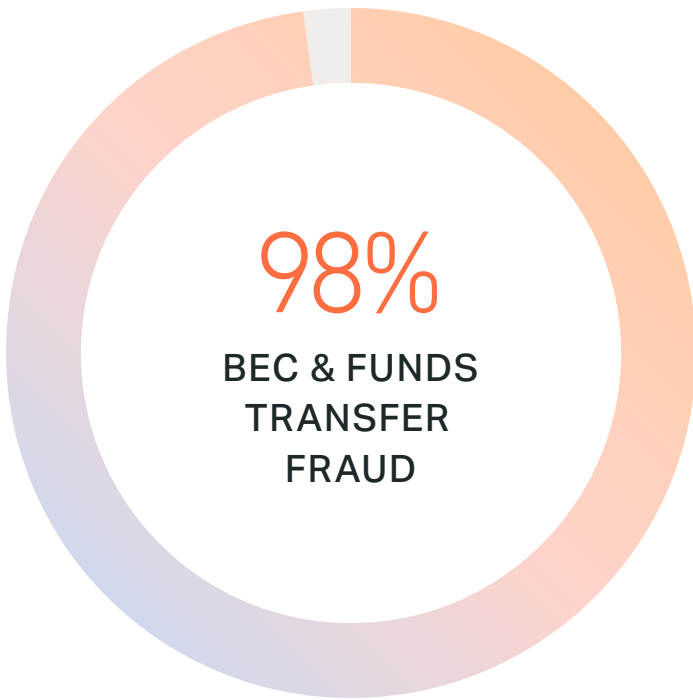
Default credentials on network devices, shared admin passwords, and accounts that hadn't been reviewed in years gave attackers immediate, broad access. By the time suspicious activity was flagged, the question wasn't how far the attacker had gotten. It was whether anything had been left untouched.

Cyber Insurance POV

Cysurance, a leading cyber insurance provider and SonicWall partner, reports that 98% of its claims stem from Business Email Compromise (BEC) and funds transfer fraud. Not ransomware, not data breaches, not advanced persistent threats. Just social engineering. An attacker sends a convincing email, someone updates payment details, and money moves to an account it should never have reached. In many cases, no system was ever compromised. No credentials were stolen. No malware was deployed. Someone was simply deceived.

The consequences are severe and personal. Cysurance is now seeing employees lose their jobs as a result of authorizing fraudulent payments. Employees who were acting in good faith, following instructions that appeared legitimate.

The fix is straightforward and costs nothing: any change to payment information must be verified by voice, every single time, without exception. Not by email. Not by chat. A phone call to a known number. That one step would prevent the majority of these claims from ever being filed. In a landscape full of complex threats, this is one of the most important and most actionable things any organization can do today.



What cyber insurance claims actually look like

BEC and funds transfer fraud	98%
all other claim types	2%

Most of these attacks involved no malware, no stolen credentials, and no technical exploit. Just social engineering.

How BEC attacks succeed



KEY INSIGHT

Any change to payment information must be verified by voice, every time, without exception. That one step would prevent the majority of these claims.

Source: Cysurance — SonicWall cyber insurance partner claims data

REMEDIATION

Restoring Meaningful Control

Addressing weak authentication, poor patching, and excessive privileges does not require new infrastructure. It requires consistent execution.

Effective remediation begins with an honest inventory:

1. Audit every account with administrative privileges and remove access that cannot be justified by their current role.
2. Change default credentials on every device, application, and platform in the environment.
3. Enforce MFA on all accounts, starting with admin and remote access, with no exceptions.
4. Establish a patch prioritization process that treats internet-facing and critical systems as urgent, not routine.
5. Set a maximum acceptable patch window for critical vulnerabilities and hold to it.
6. Review and test MFA coverage quarterly to catch gaps from new applications or user onboarding.
7. Remove or disable accounts that are no longer active.

Organizations that enforced MFA broadly and implemented structured patch cycles reduced credential-based breach risk by more than half within the first year.

Most environments don't fail because someone made a bad decision. They fail because no one made any decision at all.

False Confidence



Attackers don't care who you are. Only what they can access.

There is a particular kind of risk that doesn't show up on a threat dashboard. It doesn't trigger an alert. It doesn't generate a ticket. It lives in the gap between what an organization believes about its security posture and what is actually true.

That gap has a name: false confidence.

In 2025, it remained one of the most consequential and least discussed problems in cybersecurity. Organizations invested

in tools, checked compliance boxes, ran annual trainings, and walked away believing they were protected. Many of them were not.

"We're Too Small to Be a Target"

This is the belief that keeps security budgets low and controls weak in thousands of organizations. It is also demonstrably false.

According to the Verizon 2025 Data Breach Investigations Report, ransomware was present in 88% of SMB breaches,

Attackers don't care who you are, only what they can access.




Ransomware Presence in Breaches: SMB vs. Enterprise

Ransomware in breaches



Confident in readiness



Confidence gap reflects the difference between the percentage of C-Suite respondents and mid-level managers who rated their organization as fully prepared for a cyberattack. A larger gap indicates greater misalignment between leadership perception and operational reality.

Sources: Verizon 2025 DBIR - Dell Cyber Resilience Insights 2025 - WanAware / IBM Cost of a Data Breach 2025 Bitdefender 2025 Cybersecurity Assessment • SonicWall 2025 Cyber Threat Report

Overestimating What You Have

The second dimension of false confidence is subtler. It applies to organizations that do invest in security, but overestimate how well their controls are actually working.

A survey of 600 IT leaders found that **80% claim they can detect and contain a cyber incident in under eight hours.**³ IBM data shows attackers dwell inside environments undetected for an average of **181 days.** Those two numbers cannot both be true.

69% of IT professionals say their **leadership overestimates their readiness for a cyber event.** And while 99% of organizations report having cyber resilience strategies in place, only 46% successfully contained and recovered from an attack or drill with minimal impact.⁴

The Readiness Illusion

WHAT IT LEADERS BELIEVE

8hrs

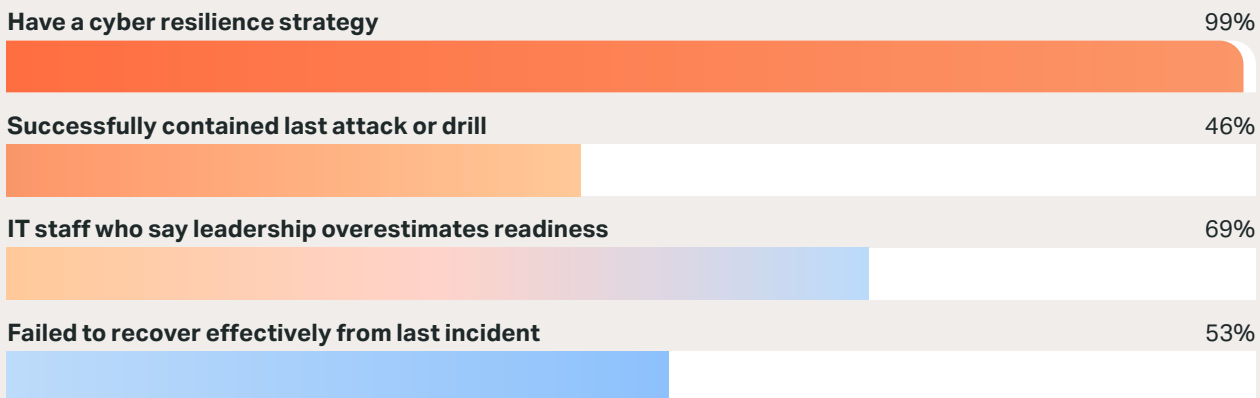
80% say they can detect and contain a breach in under 8 hours

WHAT THE DATA SHOWS

181 days

Average attacker dwell time inside environments before detection

The Gap Behind the Confidence



Sources: WanAware 2025 Cyber Response & Resilience Study • IBM Cost of a Data Breach 2025 • Dell Cyber Resilience Insights 2025

³ WanAware 2025 Cyber Response & Resilience Study

⁴ [Cybernews](#): Your Company most probably can't recover from a cyberattack and doesn't know it

DEADLY SIN #2: FALSE CONFIDENCE

Confidence isn't manufactured. It's genuine. Leaders see tools deployed, reports filed, and training completed, and they reasonably conclude the organization is protected. Nearly half of C-level respondents describe themselves as "very confident" in their organization's readiness, compared to fewer than one in five mid-level managers who run day-to-day operations.⁵ The people closest to the risk see it most clearly.

Assuming Resilience Without Testing It

The third form of false confidence is the most avoidable. Organizations that have never stress-tested their defenses have no real basis for confidence in them.

Crisis simulation data reveals a consistent and dangerous pattern: teams are paradoxically most confident in the exact areas where they perform the worst. They are not just failing. They are confidently failing.

A security plan that has never been tested under realistic conditions is a hypothesis, not a capability. Tabletop exercises that never simulate real pressure, incident response playbooks that have never been run against an actual scenario, and backup systems that have never been restored all create the appearance of preparedness without the substance of it.

SonicWall's threat data reinforces that threat actors are leveraging automation and AI-driven attacks at scale, making it clear that SMBs and organizations of all sizes can't fight this battle alone. That fight requires knowing what you have, not what you assume you have.

SMB EXPOSURE

The Confidence Tax

For SMBs, false confidence carries a specific and compounding cost. When leadership believes the organization is protected, security investment stalls. When investment stalls, gaps widen. When gaps widen, attackers find them.

The pattern is consistent across breach investigations: the organization had some security controls in place. Those controls had never been tested against a real scenario. When the attack came, the gaps that leadership didn't know existed became the gaps that defined the outcome.

SMBs rarely have a dedicated security team to push back against overconfidence at the executive level. The IT generalist who manages everything is often too busy keeping systems running to conduct a rigorous assessment of what's actually working. The result is a posture built on assumptions rather than evidence.

Assumptions don't stop ransomware.

SOC POV

In post-incident reviews at SMB environments, one conversation happens repeatedly. The leadership team is surprised. Not surprised an attack occurred. Surprised by how far it got.

Security teams investigating these cases consistently reported the same finding:

"They had tools. They just didn't know what the tools were actually covering. Closing the confidence gap doesn't require more tools. It requires an honest assessment of the ones you have."

Endpoint protection was deployed on most devices, not all. Backups existed but hadn't been tested in months. Logging was enabled, but no one was reviewing it. Each gap on its own seemed minor. Together, they formed the path the attacker took.

⁵ [Bitdefender](#): The Cybersecurity Perception Gap

REMEDIATION

Replacing Assumption With Evidence

The path forward isn't investment. It's verification. **Start by finding out what is actually working.**

1. Conduct a security posture assessment against your actual environment, not your documentation.
2. Run a tabletop exercise that tests decision-making under realistic pressure, not just process recall.
3. Test backup restoration on a defined schedule and document the results.
4. Verify that logging and alerting is functioning and that someone is actively reviewing it.
5. Close the gap between what leadership believes and what operations teams know by building regular, honest reporting into the security program.

Organizations that moved from assumed readiness to validated readiness reduced mean breach impact significantly within the first year. The controls didn't change. The knowledge of whether they worked did.

Confidence is not a security control. Verification is.

Overexposed Access

Excessive access is still the fastest path to compromise.

Getting in is only half the problem. What happens next is determined almost entirely by how much access an attacker finds waiting for them on the other side.

In too many environments, the answer is: everything. Flat networks with no meaningful segmentation. VPNs that grant broad access to internal systems the moment credentials are verified. Service accounts running with permissions accumulated over years. Users with access to systems they haven't touched in months. The perimeter is treated as the only line of defense, and once it falls, there is nothing left to slow anyone down.

This is the architecture that turns a single compromised credential into a catastrophic breach. And in 2025, it remained one of the most common configurations in SMB environments.

The "Any/Any" Mindset Lives On

Overly permissive access rules are not limited to firewall configurations. The same logic shows up across network design, identity systems, and cloud environments. The path of least resistance during setup or troubleshooting is to allow broad access, and that broad access rarely gets cleaned up.

SaaS environments in 2025 largely run on implicit trust. Once a user or application is authenticated and given access, it is trusted indefinitely. Tokens issued to third-party apps rarely expire, integrations often get more permissions than they truly need, and automations execute with minimal human oversight.⁶

The irony is sharp. Organizations that talk about Zero Trust principles often operate on the opposite model in practice: trust once, then never verify again. The initial authentication is thorough. Everything after it is assumed. Attackers often target small businesses because, with the right approach, they can achieve disproportionately high returns for relatively little effort. Modern ransomware operations are industrialized. Automated scanning tools don't filter by company size. They filter by vulnerability. If your systems are exposed, you are a target regardless of your revenue, your industry, or your headcount.

The "we're not interesting enough" mindset is not just wrong. It is operationally dangerous. It justifies underinvestment and delays the controls that would have made a difference.

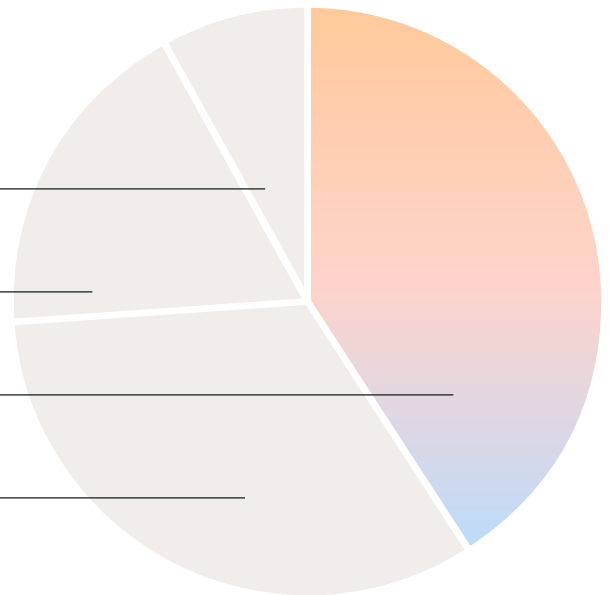
Percentage of Breached Environments with Overly Permissive Firewall Rules in 2025

8% Properly segmented environments

18% Excessive inbound exposure

41% Any/any present

33% Broad subnet-to-subnet rules



⁶ [The Hacker News](#): The Problem with 'Trust but Verify' is That We Don't Verify

Authentication Is Not the Finish Line

Traditional security thinking treats authentication as the problem to solve. Once identity is verified, access is granted. The session is trusted. The network is open.

That model was always imperfect. In 2025, it was actively exploited.

According to SonicWall data, **48% of breaches involved compromised VPN credentials as the initial access method.** In many cases, attackers obtain these credentials by compromising unmanaged or user-controlled devices, with more than 23 million exploited to extract login information, often using session cookies to bypass multi-factor authentication entirely.⁷

Authentication confirms who someone claims to be at a single point in time. It says nothing about what they should be able to reach, how far they should be able to move, or whether the session that follows is legitimate. When networks are flat and access is broad, a verified identity becomes a master key.

The Segmentation Gap

According to Illumio's 2025 Global Cloud Detection and Response Report, **92%** of organizations experienced security incidents involving lateral movement, with each incident resulting in a global average of over seven hours of downtime.

Seven hours is the average. In flat environments with no meaningful segmentation, attackers don't need seven hours. Current data shows average lateral movement occurring in 48 minutes from initial compromise, with the fastest observed attacks achieving full network propagation in just **18 minutes.**⁸

Segmentation doesn't prevent the initial breach. It limits what the breach becomes. When an attacker compromises one endpoint in a properly segmented environment, they reach that segment. In a flat network, they reach everything. The difference between a contained incident and a full ransomware deployment often comes down to whether boundaries existed to slow the spread.

Average Time to Lateral Movement

Security impact of firewall rule configurations




Chart 3

KEY INSIGHT

Restrictive firewall rules slow attackers by 7x compared to any/any rules, buying security teams critical time to detect and respond to threats. The 72-minute difference between restrictive and permissive configurations can mean the difference between containment and compromise.

⁷ SC Media: Zero trust, zero progress? Why some say the identity perimeter is still full of holes

⁸ Vectra: Lateral movement in cybersecurity

SMB EXPOSURE

Open Floor Plans

For SMBs, network segmentation is frequently treated as an enterprise concern. The environment is small, the team is small, and building security zones feels like overhead that larger organizations can worry about. That reasoning is exactly what attackers count on.

Small networks are not safer because they are small. They are more exposed because complexity has been traded for convenience. A single flat network connecting endpoints, servers, cloud applications, IoT devices, and administrative systems gives an attacker who gets past the front door a clear view of everything worth taking.

56% of organizations reported breaches exploiting VPN vulnerabilities last year, a notable rise from the prior year.⁹ VPNs built on implicit trust are a particularly acute risk for SMBs, where a single VPN connection often provides access to the full internal environment without restriction.

The fix requires rethinking what access actually means. Not who is allowed in, but what they are allowed to reach, for how long, and under what conditions.

SOC POV

In investigations where attackers moved quickly from initial access to full compromise, one characteristic appeared consistently: there was nothing slowing them down.

Security teams reviewing these cases reported the same observation:

"The credential worked, and from that point the network was theirs. There were no internal boundaries, no re-authentication requirements, nothing that forced them to work harder to get to the next system."

The attacker didn't need sophisticated tools. They needed time, and a flat network gave them all the time they needed.



Cyber Insurance POV

As AI agents become embedded in everyday workflows, the risk surface isn't just expanding, it's shifting in ways traditional controls weren't built to handle.

One of the most underestimated exposures right now: browser-saved passwords. What was once a minor convenience risk has become a primary attack vector. Infostealer malware routinely harvests credentials directly from browsers, and with AI agents increasingly granted broad system permissions, a single compromised credential can now cascade into something far more damaging than it could have a year ago.

⁹ CIO: Why 81% of organizations plan to adopt zero trust by 2026

REMEDIATION

Building Walls Inside the House

Reducing overexposed access doesn't require a full Zero Trust implementation on day one. It requires introducing friction where there currently is none.

1. Audit what every account, service, and integration actually has access to versus what it needs.
2. Segment the network into logical zones that limit movement between systems with different risk profiles.
3. Replace broad VPN access with application-level access controls that grant only what the user's role requires.
4. Remove or disable unused accounts, expired tokens, and dormant service accounts on a defined schedule.
5. Treat administrative access as a time-limited privilege, not a persistent state.

Excessive access doesn't look like a vulnerability until an attacker uses it. By then, the cost of not acting is already set.

Reactive Security Posture

If you're only reacting, attackers are dictating the timeline.

Reactive security feels like security. Alerts fire, teams respond, incidents get logged, and reports get filed. The machinery of security operations is running. The problem is that by the time most of that machinery kicks into gear, the attacker has often already done what they came to do.

Reactive security is built around response. Proactive security is built around prevention and early detection. The gap between those two approaches is measured in time, and in 2025, time was the most exploited resource in cybersecurity.

Waiting for the Alert That May Never Come

Six months. That is the average window an attacker operates undetected inside an environment while the organization believes nothing is wrong. In a purely reactive security model, there is nothing to respond to until an alert fires. And alerts only fire when a tool detects something it was configured to detect. Sophisticated attackers know how detection tools work, and they are specifically designed to operate below those thresholds.

If you're only reacting, attackers are dictating the timeline.

181 days average days to identify a breach before containment begins.

IBM COST OF A DATA BREACH 2025


44% of all security alerts go uninvestigated due to overload and talent scarcity

CYBERSIERRA / OSTERMAN RESEARCH 2025

75% of SOC analysts say they have no time for proactive threathunting

OMDIA SURVEY 2025

DAILY ALERT LOAD | WHERE IT GOES

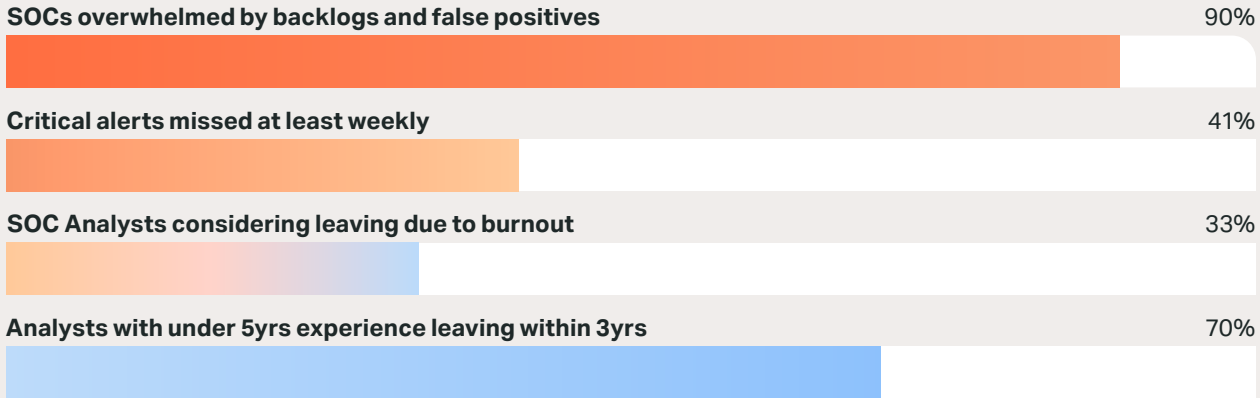


Enterprises over 20,000 employees see 3,000+ alerts daily. Nearly half go unreviewed. Most SMBs have no dedicated analyst at all.

Source: AI SOC Market Landscape 2025 • Osterman Research

Source: IBM Cost of a Data Breach 2025 • AI SOC Market Landscape 2025 • Osterman Research • OMDIA 2025 • SANS 2025 SOC Survey • Expel /IDC Research

The Alert Fatigue Spiral



Source: IBM Cost of a Data Breach 2025 • AI SOC Market Landscape 2025 • Osterman Research • OMDIA 2025 • SANS 2025 SOC Survey • Expel /IDC Research

The Alert Fatigue Problem

Reactive security also fails in a second, more systemic way. The tools that are supposed to generate the alerts teams respond to are producing more noise than any team can reasonably process.

Organizations face an average of 960 security alerts daily, with enterprises over 20,000 employees seeing more than 3,000. Nearly 90% of SOC's are overwhelmed by backlogs and false positives, and 80% of analysts report feeling consistently behind in their work.¹⁰

With 44% of all alerts going uninvestigated due to a combination of talent scarcity and alert overload, organizations face significantly increased breach risk.¹¹

Alert fatigue is not laziness. It is a structural problem. When the ratio of noise to signal becomes unmanageable, the human response is predictable: triage becomes faster, thresholds for investigation rise, and real threats start looking like the false positives that surround them.

The alert that matters is buried. The attacker is still moving.

No Incident Readiness

Reactive security also assumes that when an incident does occur, the organization knows how to respond. That assumption is almost always wrong.

Without documented playbooks, practiced response procedures, and clear escalation paths, the response to a real incident becomes improvised. Improvised responses are slower, more expensive, and more likely to miss critical containment steps.

Incident readiness is not just about having a plan on paper. It requires regular testing, defined roles under pressure, and coordination between security, IT, legal, and leadership that has been rehearsed before it is needed. Most organizations have none of that in place.

¹⁰ Dropzone AI: Alert Fatigue in Cybersecurity

¹¹ Cyber Sierra: What is Alert Fatigue and How to Combat in Your SOC

SMB EXPOSURE

No One Watching the Clock

For SMBs, the reactive security problem is compounded by a simple reality: there is often no one watching at all.

75% of SOC analysts report a lack of time for strategic work like threat hunting or professional growth. Security operations become purely reactive, constantly responding to yesterday's threats while tomorrow's attacks go undetected.¹²

That dynamic describes large organizations with dedicated security teams. SMBs typically have no dedicated SOC analyst at all. Security monitoring is handled by the same IT generalist who manages endpoints, cloud accounts, and user support. Alerts that go unreviewed for hours in an enterprise go unreviewed for days in a small business, if they are reviewed at all.

Attackers know this. 24/7 monitoring is not a luxury for SMBs. It is the basic requirement for detecting threats that do not announce themselves. Without it, reactive security becomes no security at all.

The Alert Threshold Problem

Many organizations have drawn a hard line at high/critical alerts, leaving medium-severity events unactioned. This is one of the most dangerous blind spots in a network environment today. We are actively seeing firewall rule changes logged as medium alerts (the kind that get triaged, queued, and forgotten) directly preceding full account takeovers and funds transfer fraud. The attacker didn't bypass the controls. The alert fired. Nobody acted.

SOC POV

In investigations involving SMBs that had suffered extended breaches, one finding repeated consistently: the alerts were there. They just weren't seen.

Security teams examining post-incident logs found the same pattern:

"The signals were in the data. There was lateral movement, unusual outbound traffic, failed authentication attempts at odd hours. All of it logged. None of it reviewed. By the time anyone looked, the attacker had been operating for weeks."

The tools were working. The monitoring was not.

Cyber Insurance POV

The most preventable losses we see aren't caused by sophisticated attacks. They're caused by gaps in enforcement, alert prioritization, and readiness and attackers know exactly how to exploit them.

¹² **Abnormal:** Alert Fatigue: The Hidden Cost Draining Your SOC

REMEDIATION

Stop Waiting to Be Found: How to Close the Security Detection Gap

Closing the reactive gap requires both better tooling and a change in how security operations are structured.

1. Implement 24/7 monitoring through a managed security service provider if internal capacity doesn't exist.
2. Establish a formal threat hunting cadence, even if monthly, that assumes compromise and actively looks for evidence.
3. Build and test incident response playbooks for the scenarios most likely to affect your environment: ransomware, credential compromise, and data exfiltration.
4. Reduce alert noise by tuning detection rules and consolidating tools so analysts spend time on signal, not noise.
5. Set mean time to detect and mean time to respond as tracked metrics, not aspirational goals.

Speed of detection is not just an operational metric. It is a financial one. The attackers are already setting the timeline. The only way to change that is to stop waiting for them to announce themselves.

Cost-Driven Security Decisions

Cheap security creates expensive problems.

Security budgets feel like costs. Breaches feel like crises. The problem is that organizations consistently compare the known, visible cost of security investment against the unknown, hypothetical cost of a breach. The known cost almost always wins the argument.

That calculation is wrong. And in 2025, the data made it undeniable.

Prioritizing Cost Over Outcomes

The most common form of cost-driven security failure is not cutting security entirely. It is making decisions that feel

financially responsible but create outsized risk. Deferring a penetration test. Skipping security awareness training because "people are busy." Choosing the cheapest endpoint protection without considering whether it integrates with anything else. Passing on an incident response retainer because nothing has happened yet.

Organizations with comprehensive incident response plans save \$1.23 million per breach compared to those without.13 The cost of building that plan is a fraction of that figure. The cost of not having it shows up when it is too late to write one.

Tools don't fix problems, people do. Underinvesting in people is not a budget decision. It is a risk decision with a measurable price tag.

Cheap security creates expensive problems.

WHAT YOU SAVE CUTTING CORNERS	WHAT IT COSTS WHEN IT GOES WRONG
Skipping incident response planning saves a few thousand dollars in consultant fees	\$1.23M saved on average by orgs with incident response plans (IBM 2025)
Deferring security training saves time and budget this quarter	95% of breaches involve human error — training is the control (IBM / Verizon 2025)
Buying the cheapest tool avoids a larger platform investment	\$4.44M global average cost of a data breach (IBM 2025)
Underinvesting in staff keeps headcount costs low	\$1.57M extra cost added to breaches by cybersecurity skills shortage (IBM 2025)

The Tool Sprawl Problem

45
average cybersecurity tools operated by enterprises (Gartner 2025)

46%
of security teams spend more time managing tools than defending against attacks

74%
of repeat ransomware victims say they're juggling too many security tools (Barracuda 2025)

Source: IBM Cost of a Data Breach 2025 • AI SOC Market Landscape 2025 • Osterman Research • OMDIA 2025 • SANS 2025 SOC Survey • Expel /IDC Research



Tool-First Instead of Strategy-First

The second dimension of cost-driven failure is counterintuitive. It is not always about spending too little. It is often about spending on the wrong things.

When a breach happens or a threat makes headlines, the instinctive response is to buy something. A company gets scared about security, executives itch to "do something," and the purchase of a new tool feels like something is being done. It is the cybersecurity version of buying a treadmill after a bad checkup.

The result is tool sprawl. According to Gartner's 2025 research, the average enterprise now operates 45 different cybersecurity tools. Nearly half of security professionals surveyed say they spend more time maintaining tools than defending against actual attacks.¹⁴

More tools do not mean more security. 74% of organizations hit by multiple ransomware attacks say they are juggling too many security tools, while 61% report their tools don't integrate properly.¹⁵ Fragmented tools create fragmented visibility, which creates the gaps attackers exploit.

65% of organizations say they have too many security tools, and over half say their tools can't be integrated. Yet 65% also say consolidation would improve their overall risk posture.¹⁶

They know the problem. They haven't fixed it.

¹⁴ [Security Boulevard](#): Too Much Time Being Spent on Managing Cybersecurity Tools

¹⁵ [Insights from Analytics](#): Why Security Tool Sprawl Is Making Organizations Less Secure

¹⁶ [The Stack](#): The Risks of Cybersecurity Sprawl

Underinvesting in People and Process

The third failure is the most persistent. Tools get budget approval because they appear on invoices. People and process improvements are harder to quantify, easier to defer, and rarely generate a visible deliverable the week they are funded.

Despite record investments in cybersecurity technology, the weakest link in 2025 remains the same: people. Human error drives the majority of breaches, a sobering reminder that security is ultimately a human challenge, with tools acting as enablers rather than solutions.

Security awareness training, documented processes, governance frameworks, and clear ownership of security responsibilities are consistently underfunded relative to technology. The result is environments where the tools are deployed but the people using them don't know how, the processes for responding to alerts don't exist, and no one is accountable for maintaining either.

SMB EXPOSURE

The False Economy

For SMBs, cost-driven security decisions feel unavoidable. Budgets are real constraints. The question is not whether to spend carefully, but whether the savings from underinvestment are genuine or deferred costs that will arrive with interest.

SonicWall estimates that a single breach at an SMB could exceed \$4.91M when downtime and recovery are included. For many small businesses, that amount could be a matter of survival. The cost of the breach dwarfs the cost of the controls that would have prevented it.

The calculation that leads to underinvestment tends to be: nothing has happened yet, so nothing needs to change. That reasoning holds right up until it doesn't. And by the time it doesn't, the savings are gone and the damage is done.

The smarter framing is return on resilience. What does it cost to implement MFA, maintain patched systems, run quarterly security reviews, and have an incident response plan? Compared to what does it cost when those things are absent and an attacker finds the gap?

The math is not close.

SOC POV

In post-incident reviews at cost-constrained environments, one theme appeared consistently across investigations. The organization had made a series of individually defensible decisions that collectively removed the controls that mattered most.

Security teams examining these cases reported the same observation:

"Every decision made sense in isolation. The training was skipped because it was a busy quarter. The tool wasn't purchased because last year was quiet. The consultant wasn't engaged because nothing had happened. Together, those decisions removed every meaningful layer of protection."

The breach did not happen because of one bad choice. It happened because a series of reasonable-seeming cost decisions accumulated into an environment with no depth.



Cyber Insurance POV

One of the most common questions we receive after an incident: will coverage apply if the impacted system wasn't fully up to date or properly managed?

It's a fair question. IT environments are complex, constantly changing, and no organization achieves perfect hygiene across every endpoint, tool, and backup system. We understand that. Change is continuous, and exceptions are sometimes unavoidable.

But there's a meaningful difference between an exception and a culture of exceptions.

REMEDIATION

Spending Smarter, Not Just More

Addressing cost-driven security failure is not about spending more. It is about spending on what actually reduces risk rather than what looks like security.

1. Build a simple risk register that maps your most likely threat scenarios to the controls that address them and fund those controls first.
2. Consolidate security tools around integrated platforms rather than accumulating point solutions.
3. Treat incident response planning as a non-negotiable line item, not a deferred project.
4. Invest in security awareness training on a regular schedule, not as a one-time annual event.
5. Audit tool utilization before purchasing new tools. Half of what most organizations own is underused.
6. Engage an MSP/MSSP if internal expertise gaps exist. The cost is predictable; the cost of a breach is not.

Smart investment consistently reduces breach cost. The question is whether that investment happens before the breach or after it.

Cheap security saves money today, but could potentially cost much more tomorrow.

Reliance on Legacy Access Models

If your access model hasn't evolved, neither has your defense.

There was a time when the perimeter made sense. Employees worked from offices, applications ran in data centers, and the network boundary was real and defensible. Build a strong enough wall and you controlled what came in.

That world is gone. The workforce is distributed, applications live in the cloud, and data moves across environments that

no single firewall can see. Attackers figured this out long ago. The fastest way into an organization is not through the wall anymore. It is through the door, using credentials that belong to someone who is supposed to be there.

Legacy access models were designed for a threat landscape that no longer exists. Organizations that haven't updated them are defending a perimeter that has already dissolved.

If your access model hasn't evolved, neither has your defense

NETWORK-FIRST THINKING

- Trust everyone inside the perimeter
- Authenticate once, access broadly
- VPN grants full network access
- Security boundary = firewall edge
- IP address determines trust



IDENTITY FIRST THINKING

- Never trust, always verify — regardless of location
- Continuous verification throughout the session
- Application-level access only — least privilege
- Security boundary = every access request
- Identity, device, and context determine trust

The VPN Problem

VPNs became the standard for remote access in a different era, and for decades they worked reasonably well. Today, they are one of the most exploited entry points in enterprise security.

SonicWall data shows that 48% of breaches were attributed to compromised VPN credentials as the initial access vector, with weak access controls such as absent MFA contributing to those compromises. The reason is architectural. VPNs authenticate once and then grant broad network access. An attacker who acquires valid credentials doesn't just reach one application. They reach the network, and from there, everything the network can reach.

VPN CVEs grew by 82.5% over the analyzed period, with roughly 60% of those vulnerabilities rated high or critical.¹⁷ VPN concentrators must be publicly visible to accept connections, which makes them permanent targets for automated scanning. Attackers don't need sophisticated techniques to find and probe them. They just need to look.

90% of organizations have one or more issues with their current VPN. And 83% of users report being willing to bypass security measures just to stay productive,¹⁸ which means the controls the VPN is supposed to enforce are being routed around regularly by the people the organization is trying to protect.

¹⁷ CIO: Why 81% of organizations plan to adopt zero trust by 2026

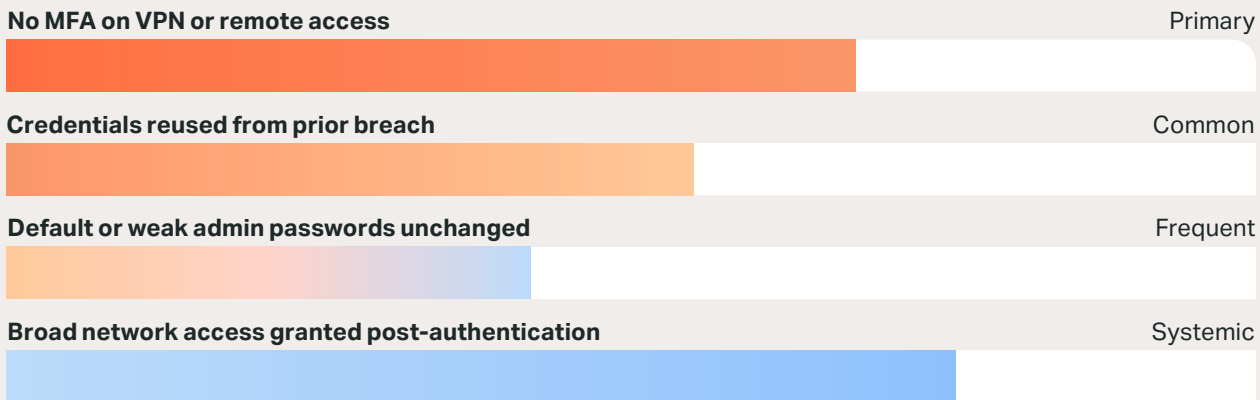
¹⁸ Tailscale: The State of Zero Trust

Initial Access Vector: Credential Compromise

48% of breaches traced to compromised VPN credentials as the initial access vector.

Weak access controls (including absent MFA) were a contributing factor in the majority of those compromises.

Contributing Weaknesses Behind the 48%



Network-First vs. Identity-First

The deeper problem is not the VPN itself. It is the thinking behind it. Network-first security assumes that controlling where someone connects from is the primary security question. Identity-first security starts with a different question: who is this, what do they actually need, and should this specific request be granted right now?

Attackers are not breaking through firewalls. They are logging in. And once they are logged in with valid credentials in a network-first environment, they look identical to legitimate users. Threat Actors are not hunting for zero-days. They are buying or stealing the credentials that get them past the only check the network-first model performs.

Resistance to Architectural Evolution

Knowing the model needs to change and actually changing it are very different things.

The resistance is understandable. Changing access architecture is complex, potentially disruptive, and requires coordination across IT, security, and the business. The path of least resistance is to maintain the existing model and patch around the edges.

But the cost of that resistance is compounding. Every year the legacy model stays in place is another year that attackers have a predictable, well-understood attack surface to exploit. The tools to find exposed VPN infrastructure, identify valid credentials, and move laterally through flat networks are widely available and increasingly automated.

The architecture that was supposed to prevent those incidents is the same architecture that made them possible.

SMB EXPOSURE

Stuck on Old Infrastructure

For SMBs, legacy access models are often the result of a single decision made years ago that was never revisited. A VPN was deployed, it worked, and no one had the time or budget to change it. The question of whether it still represented the right security model was never asked.

The risk is acute. SMB networks are typically flat, VPN access often grants access to the full environment, and there is rarely anyone monitoring whether the credentials being used are legitimate or stolen. An attacker with a valid set of VPN credentials in an SMB environment faces almost no friction between initial access and full compromise.

The shift to identity-first security does not have to happen all at once. It starts with asking the right questions: what does each user actually need to access, how is that access being verified, and what happens when those credentials are stolen? Many SMBs can begin addressing these questions through existing cloud identity platforms they already own, without major infrastructure investment.


SOC POV

In investigations involving legacy access environments, one observation appeared consistently. The breach was not technically sophisticated. The attacker looked legitimate from the moment they connected.

Security teams examining these cases reported the same finding:

"They had valid credentials and VPN access. From that point, everything they did looked like normal user activity. By the time the behavior became obviously anomalous, they had been inside for weeks."

The network-first model had no mechanism to ask whether the session that followed authentication should have continued. It had already decided to trust. It never asked again.



REMEDIATION

Evolving the Access Model

Moving from network-first to identity-first security is a journey, not a single project. But the direction is clear, and meaningful progress is possible without replacing everything at once.

1. Audit what your VPN access currently grants and whether that scope is justified by business need.
2. Begin replacing broad VPN access with application-level access controls for the highest-risk use cases first.
3. Implement continuous session monitoring so that authentication is not the only checkpoint.
4. Enforce least privilege access so that credentials, if stolen, reach only what they need to reach.
5. Adopt phishing-resistant MFA, such as FIDO2 or hardware keys, for all remote and administrative access.

Enterprises that have transitioned from VPN to Zero Trust cite improved security and compliance as the primary advantage, reporting a reduction in ransomware, credential theft, and lateral movement risk was the main benefit.

The perimeter has moved. It is no longer the edge of the network. It is every access request, every session, every identity. Organizations that have recognized that shift are significantly better protected. Those still defending the old perimeter are guarding something that attackers stopped caring about years ago.

Chasing Hype Over Execution

Tools don't create outcomes. Execution does.

Every year brings a new wave of security technology that promises to change everything. AI-powered detection, autonomous response, next-generation platforms with capabilities that previous generations couldn't touch. The marketing is compelling, the demos are impressive, and the urgency feels real.

And in many cases, the technology is genuinely powerful. The problem is not the tools. It is the belief that buying them is the same as being protected by them.

The AI Reality on Both Sides of the Fight

Artificial intelligence has fundamentally changed the speed and scale of cyberattacks. That is not hyperbole. It is documented.

According to the CrowdStrike Global Threat Report 2026, there was an 89% increase in attacks by AI-enabled adversaries in 2025 compared with the previous year. Attackers deployed AI to aid with social engineering, malware development, and disinformation campaigns.

This acceleration is real, and it matters. Campaigns that once took weeks to prepare now take hours. Phishing emails that once required manual crafting are now generated, personalized, and localized at scale. Reconnaissance that previously consumed significant attacker resources is now largely automated.

AI is also a powerful defensive tool. Organizations using AI-driven security platforms report detecting threats up to 60% faster than those using traditional methods.¹⁹ The case for AI in security is strong on both sides of the equation. But it is not unconditional.

When AI Becomes an Excuse

SonicWall's Michael Crean, General Manager of Managed Security Services, has spent over two decades running an MSSP and investigating breaches across SMB environments. His perspective is direct.

"Attacks are getting faster, and in some instances, they're getting a little more sophisticated. But the vast majority of the attacks that we're seeing and investigating are basic fundamentals that are still being missed. It's like we've gotten so quote-unquote smart with AI that we're allowing it — or pretending to allow it — to overcompensate for the things that are still probably the most important to do."²⁰

**Michael Crean, GM of Managed Security Services,
SonicWall**

That observation cuts to the heart of the seventh deadly sin. Organizations are chasing AI capabilities while leaving MFA unconfigured, patches unapplied, and access policies undisciplined. The new tool gets deployed. The fundamentals stay broken. And the attacker, who doesn't care what's on your dashboard, walks through the same gaps as before.

90% of companies currently lack the maturity to effectively counter today's advanced AI-enabled threats.²¹ That maturity gap is not primarily a technology problem. It is an execution problem.

¹⁹TechAdvisors: AI Cyber Attack Statistics 2025

²⁰Channel Insider: SonicWall's Michael Crean on State of Managed Security

²¹DeepStrike: AI Cybersecurity Threats 2026



AI Threat Maturity Gap 2025

lack maturity to counter AI threats	90%
have adequate maturity	10%

That maturity gap is not primarily a technology problem, but an execution problem.

Source: DeepStrike — AI Cybersecurity Threats 2025

Tool-First Thinking vs. Outcome-First Thinking

The cycle is familiar: a breach makes headlines, a new category of tool emerges to address it, organizations buy it, add it to an already complex stack, and move on. Meanwhile the tool is misconfigured, underutilized, or generating alerts that no one reviews.

The organizations that get the most from AI-powered security are the ones that have already done the foundational work:

- Their logs are centralized and reviewed.
- Their access policies are enforced.
- Their patching is disciplined.
- AI makes those controls faster and smarter.
- In environments where the controls are absent, AI has nothing to work with.

Tools amplify what is already there. They do not replace what is missing.

Constantly Shifting Instead of Maturing

Constant context-switching is its own form of risk. When the security program never stays still long enough to be executed, it never matures. Last year's platform gets abandoned for this year's. Integration work goes unfinished. The team never gets proficient with anything, because there is always something newer demanding their attention.

Security maturity is not measured by how current your tooling is. It is measured by how consistently your controls are enforced, how quickly you detect anomalies, and how reliably you can respond when something goes wrong. Those outcomes come from depth and discipline, not from chasing the next promising technology.



SMB EXPOSURE

Distracted by the Shiny

For SMBs, the hype cycle is particularly dangerous because resources are limited. Every dollar spent on a tool that isn't properly deployed or integrated is a dollar not spent on the training, the process, or the managed service that would have actually reduced risk.

The most common security failure pattern in SMB breach investigations is not a missing tool. It is a present tool that wasn't configured correctly, wasn't monitored, or wasn't connected to anything else. The endpoint protection that didn't cover all devices. The SIEM that was deployed but never tuned. The firewall that was managed by the vendor's default settings.

Buying is easy. Operating is hard. The organizations that protect themselves most effectively are not necessarily the ones with the most advanced tools. They are the ones that have decided to do a smaller number of things completely and correctly.

SOC POV

In post-incident reviews across SMB environments, the same pattern surfaced repeatedly. The organization had recently invested in a new security platform. The breach happened anyway.

Security teams examining these cases found the same observation:

"They had good technology. But the deployment was incomplete, the alerts were landing in an inbox nobody watched, and the basics had been pushed aside while everyone focused on standing up the new platform. MFA on remote access, a current patch on an internet-facing system. The simple things that would have stopped this were never implemented."

REMEDIATION

Execution Over Acquisition

The path forward is not about buying less technology. It is about ensuring that what is bought gets fully deployed, properly configured, and actively used.

1. Audit every security tool in your environment and assess whether it is deployed completely and functioning as intended.
2. Before purchasing new technology, define the specific security outcome you are trying to achieve and verify that existing tools cannot deliver it.
3. Treat AI-powered security capabilities as force multipliers for disciplined operations, not substitutes for them.
4. Integrate AI-assisted detection and response into a workflow that humans are actively managing, not passively monitoring.
5. Build a security roadmap that prioritizes maturity and stability over constant adoption of new categories.

AI-enabled SOCs that can facilitate faster alert triage, combined with next-generation approaches to detection and response, will transform how MSPs and MSSPs protect their customers. But that transformation requires the foundation to be solid first.

AI is not the answer to bad hygiene. It is a powerful capability that belongs on top of good hygiene. In 2025, the organizations that understood that distinction were significantly better protected than those that didn't.

The cat-and-mouse game is accelerating on both sides. The organizations that will win are not the ones with the best tools. They are the ones that execute the fundamentals, consistently and completely, while using technology to go faster.

IN SUMMARY OF THE SEVEN DEADLY SINS OF CYBERSECURITY

The Real Cybersecurity Challenges Facing SMBs Today


Understanding the real-world cybersecurity challenges SMBs face is essential for IT teams, service providers, and leadership alike. While headlines focus on AI-powered attacks and nation-state threats, most successful breaches begin with something far more preventable: the wrong priorities, the wrong assumptions, and the wrong approach to fundamentals.

Today's threat landscape does not just exploit zero-day vulnerabilities. It exploits gaps in discipline, governance, and execution. For SMBs operating with lean teams and growing complexity, the seven patterns listed on pages 41 and 42 consistently define the difference between resilience and exposure.

ONE

Ignoring the Fundamentals

Weak authentication, unpatched systems, and excessive admin privileges remain the primary attack surface. Attackers don't need sophisticated tools when the basics aren't covered. Most breaches don't start advanced. They start with a gap that should have been closed.



TWO

False Confidence

Believing you are too small to be a target, overestimating how well your controls are working, and assuming resilience without ever testing it are three of the most expensive mistakes an organization can make. Ransomware was present in 88% of SMB breaches in 2025. The data does not support the assumption.

THREE

Overexposed Access

Overly permissive rules, implicit trust after authentication, and flat networks with no meaningful segmentation give attackers a clear path once they are inside. Excessive access is still the fastest path to compromise, and once credentials are stolen, broad access turns a single breach into a total loss.



FOUR

Reactive Security Posture

Waiting for alerts to fire, drowning in noise that buries real threats, and having no incident readiness plan means attackers are always setting the timeline. The average breach goes undetected for 181 days. In environments without 24/7 monitoring, that window is even wider.

FIVE

Cost-Driven Security Decisions

Prioritizing the lowest cost option over the right outcome, accumulating disconnected tools, and underinvesting in people and process creates the illusion of security without the substance. Cheap security is not cheap. It is just cheap upfront.



SIX

Reliance on Legacy Access Models

Network-first thinking, VPNs that grant broad access on a single credential check, and resistance to architectural evolution leave organizations defending a perimeter that attackers stopped caring about years ago. Identity

is the new perimeter. Organizations still treating the firewall edge as the primary boundary are already behind.

SEVEN

Chasing Hype Over Execution

Buying the latest tools without fully deploying them, expecting technology to fix process gaps, and changing strategy before anything matures creates risk. As SonicWall's Michael Crean notes, most breaches still come down to missed fundamentals. Tools alone do not create outcomes. Execution does..

SMBs often already have the tools they need. The real challenge is operational discipline: enforcing the basics, honestly assessing what is working, updating access controls to match today's business, and validating security rather than assuming it.



Strategic Actions SMBs Must Take to Eliminate Critical Security Gaps

The seven deadly sins of cybersecurity are not caused by insufficient technology. They are caused by insufficient discipline, visibility, and follow-through. Most of them are fixable without a major budget overhaul. What they require is intentional action, applied consistently.

Close the fundamentals gap

- Enforce MFA on every account with no exceptions, starting with remote access and admin
- Establish a patch process that treats critical and internet-facing systems as urgent
- Audit admin privileges and remove access that can't be justified by current role
- Change default credentials on every device at deployment

Replace assumption with evidence

- Assess your actual environment, not your documentation
- Test backup restoration on a defined schedule and document the results
- Verify that logging and alerting is functioning and someone is actively reviewing it
- Run tabletop exercises that simulate realistic pressure, not just process walkthroughs

Reduce your blast radius

- Segment your network so a single compromised credential can't reach everything
- Replace broad VPN access with application-level controls tied to role
- Remove unused accounts, expired tokens, and dormant service accounts regularly

Move from reactive to proactive

- Implement 24/7 monitoring through a managed security service if internal capacity doesn't exist
- Build and test incident response playbooks before an incident occurs
- Establish a threat hunting cadence that actively looks for compromise rather than waiting for alerts

Spend on outcomes, not appearances

- Fund controls that address your most likely threat scenarios first
- Consolidate tools around integrated platforms rather than accumulating point solutions
- Audit tool utilization before purchasing anything new

Evolve your access model

- Begin shifting from broad network access to identity-based, application-level controls
- Enforce least privilege so stolen credentials reach only what they need to reach
- Inventory and govern non-human identities, service accounts, and API tokens

Execute before you acquire

- Audit every security tool you own and assess whether it is fully deployed and working
- Define the security outcome you need before evaluating any new vendor
- Treat AI-powered capabilities as force multipliers for good fundamentals, not substitutes for them

Security maturity is not defined by how many tools are deployed. It is defined by how well existing controls are configured, monitored, and maintained. The difference between protection and exposure is often not technology. It is execution.

SonicWall helps organizations execute across all seven areas. From managed security services that deliver 24/7 monitoring and threat hunting, to network security and segmentation that reduce blast radius, to identity-aware access controls that enforce least privilege. SonicWall's portfolio is built around the fundamentals that matter most to SMBs.

About SonicWall

For more than 30 years, [SonicWall](#) has championed a partner-first model that combines purpose-built technology, cloud-delivered security services, and real-time threat intelligence to help businesses prevent breaches, reduce risk, and stay operational in the face of evolving modern threats. The company is committed to delivering the best security outcomes for its customers, while others deliver only features and functions. Through its unified cybersecurity portfolio and global community of more than 17,000 partners, SonicWall enables managed service providers to actively manage, continuously optimize, and measurably protect networks, cloud environments, endpoints, and applications. SonicWall is redefining cybersecurity around outcomes that matter to business leaders, including breach prevention, compliance achievement, cost efficiency, and reduced human error, because protection is not about what a product can do, but about what it actually delivers.



SonicWall, Inc.

1033 McCarthy Boulevard | Milpitas, CA 95035
Refer to our website for additional information.
www.sonicwall.com

SONICWALL®

© 2026 SonicWall Inc. ALL RIGHTS RESERVED.

SonicWall is a trademark or registered trademark of SonicWall Inc. and/or its affiliates in the U.S.A. and/or other countries. All other trademarks and registered trademarks are property of their respective owners. The information in this document is provided in connection with SonicWall Inc. and/or its affiliates' products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of SonicWall products. Except as set forth in the terms and conditions as specified in the license agreement for this product, SonicWall and/or its affiliates assume no liability whatsoever and disclaims any express, implied or statutory warranty relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or non-infringement. In no event shall SonicWall and/or its affiliates be liable for any direct, indirect, consequential, punitive, special or incidental damages (including, without limitation, damages for loss of profits, business interruption or loss of information) arising out of the use or inability to use this document, even if SonicWall and/or its affiliates have been advised of the possibility of such damages. SonicWall and/or its affiliates make no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. SonicWall Inc. and/or its affiliates do not make any commitment to update the information contained in this document.