PICUS SECURITY

Platform Messaging Document

Status: APPROVED **Last Updated:** March 2026 **Owner:** Kevin Cole **Version:** 2.0

ABOUT THIS DOCUMENT

This is **the foundational framework** for how we articulate the Picus story. It is an internal reference for Marketing teams — not a sales asset, a pitch guide, or a demand gen playbook. It establishes shared language, not (necessarily) public copy.

See here for the Claude skill and our guardrails for all content creation: Guardrails/Prompt

0. Narrative in a Nutshell

Every security vendor has AI now. From vulnerability assessments to pentesting to attack simulations, "AI-powered" has become table stakes. The 2023 question — *are you using AI?* — is settled. The 2026 question is harder: *does your AI actually know your environment well enough to operate autonomously within it?*

Adversaries have already answered that for themselves. They're not running generic scripts against abstract targets. They're probing your specific infrastructure, chaining techniques in your specific environment, exploiting your specific gaps — at machine speed, around the clock, without a human in the loop. It's why the time to exploit a new CVE is nearly 200x faster than it was a few years ago. Why breakout times have collapsed by 94% and some are happening in under 30 seconds.

The answer isn't more tools, more visibility, or more headcount. It's a new operating model: autonomous exposure validation that actually knows your environment and what matters to it.

Picus is the context window for your security stack. Context is what separates the new autonomous era from AI-assisted pentesting or traditional BAS. Anyone can run attack scripts fast or do an automated config check. Only Picus runs them in context.

By ingesting real-time signals across threat intelligence, asset topology, control effectiveness, and business workflows to determine what actually matters, Picus can validate what's genuinely exploitable and push findings into the right teams through the right channels – or autonomously remediate within tunable guardrails you control.

Picus Swarm, our agentic purple team, executes this cycle continuously: synthesizing signals, researching threats, crafting adversary-informed simulations, validating your defenses, and mobilizing remediation in an unbroken, autonomous loop. Crucially, Picus Swarm is autonomy with a chain of custody: every action traceable and auditable, every agent beholden to your rules, and no hallucinated attack paths or yet another opaque exposure score.

The result: validated, context-rich security that moves at machine speed without sacrificing the judgment, traceability, and control that enterprise security demands.

Summary:

Why Change?	**Why Now?**	**Why Picus?**
● Uncertainty about readiness against emerging threats ● Too many context-free and theoretical findings & alerts – can't tell which exposures truly matter to my business ● Traditional human-in-the-loop options are too slow to keep up with agentic adversaries ● Under-resourced and high operational overhead	● AI reducing time-to-exploit to <1 day & breakout to < 30 mins ● Need agentic AI to fight AI given flood of new CVEs ● Teams being asked to keep up with AI adversaries but cannot match machine speed and are already burnt out on too many tools ● Need to defend the spend or prove effectiveness of your security program, including controls, processes, and people	● Only platform converging BAS, APT, EXV for context-rich, signal-driven validation ● Picus Swarm removes operational overhead – autonomously discovers, validates, fixes & optimizes, then re-validates 24x7 ● Tunable AI autonomy so our agents operate safely by your rules not ours ● Trusted leader: 95% recommendation rate, 4.9 on G2, 4.8 on GPI, #1 Leader on Frost Radar

1. What Holds Customers Back

(aka pains and challenges)

- **Validation runs on human time, adversaries run on machine time:** periodic pentests and scheduled red team exercises create unavoidable windows of unvalidated exposure that sophisticated adversaries are built to exploit. Even AI-assisted defensive efforts still depend on manual handoffs between teams, disparate tools & processes, and siloed priorities. Security teams end up in an exhausting hamster wheel that leaves them structurally outpaced by the exponential growth of the attack surface and the accelerating pace of agentic adversaries.
- **Security tools are assumed to be working, not validated:** EDR policies degrade, detection rules go stale, AI guardrails go untested, and firewall configs drift. Many organizations still focus on the existence of security controls (a compliance checkbox) instead of the effectiveness of controls — a dangerous gap as AI-powered attacks scale. Without continuously validated evidence, teams lack the confidence to defend their budget choices, pass regulatory or audit checks, or defend their security posture in the face of the next headline-making threat.
- **Findings flood in without context that makes them actionable**: teams face 135 new CVEs every day with little clarity on what's truly exploitable nor how much that exploitability matters within their business context. Without adversarial context tied to their specific environment, teams can't distinguish what's exploitable and urgent from what's theoretical noise. Worse, adding autonomous pentest or remediation agents to the mix who don't have full context means the noise cranks up at the speed of AI.

2. What Customers Are Doing Today

(aka the painful status Quo)

- **Prioritize vulnerabilities based on technical context alone** without accounting for asset context or business context. It's a whack-a-mole approach that leaves teams scrambling to keep up with threats, patching whatever CVSS flags as "critical," and chasing remediation for vulnerabilities that may not even truly matter to their organization.
- **Run manual and automated penetration tests** that deliver critical depth across specific attack paths but are outdated within days or weeks, can't scale across hybrid estates, and cost $50K-200K per engagement. Teams end up with expensive point-in-time snapshots that tick a compliance checkbox while leaving fatal gaps against AI-powered attackers who probe continuously.
- **Run traditional BAS** tools that use static attack libraries and scripted playbooks. Even with AI automation, they can't match the pace of agentic adversaries who exploit CVEs within a day of disclosure.
- **Operate a patchwork of point solutions**, possibly even FOSS, that can't keep up with the speed of AI, don't have adequate business-aware context, nor have the performance & scale the modern enterprise demands. They often reinforce organizational politics and data silos with little connection to the broader enterprise or security ecosystem. Worse, 77% of teams say adopting too many tools has actively hindered their ability to prevent and detect threats.

- **Overwhelm themselves with brute force security:** throw more bodies at the problem, deploy more tools, trawl through more logs. This works until budget constraints, hiring freezes, or burnout hit – and 76% of security teams report they're already there.

3. What Customers Need Instead

(aka vendor-agnostic critical capabilities)

- **Validation that responds to real-time signals**, not human schedules: AI agents that mobilize autonomously when threats emerge, assets change, or configurations drift — not when someone books a calendar slot.
- **Context-aware validation** that reflects their specific reality rather than generic CVSS and EPSS scores or simulations against an abstract environment.
- **A closed loop from finding to fixing**, without manual handoff, using continuous re-validation and vendor-specific mitigations deeply integrated with their existing stack.
- **Trustworthy AI that security teams can govern and audit:** Full traceability on every test, every finding, every autonomous action with human-in-the-loop options where they matter. No black boxes, no unexplained outputs, no hallucinations.
- **Real-world evidence that answers the three questions everyone is asking**: How would attackers get in and what could they reach? Would our defenses hold? If not, what should be fixed first and how do we close our gaps?

4. What Picus Can Do for Them

(aka value propositions or painkillers & vitamins)

- **Autonomously validate controls, exposures, and attack paths**: Picus Swarm deploys AI agents that operate in parallel 24x7 across network, endpoint, cloud, web app, email, and AI stacks without waiting for human initiation. Picus Swarm operates within customizable guardrails to ensure transparency and safety.
- **Operationalize threat intelligence at machine speed:** As CVEs emerge and TTPs evolve, Picus agents immediately ingest, simulate, and validate whether current controls would detect & prevent those specific threats in your specific environment.
- **Push validated, context-rich findings into the workflows where remediation happens:** Picus does more than simply pull from ITSM and SIEM — we push enriched, validated, prioritized findings back into ServiceNow, Jira, SOAR playbooks, and EDR platforms with the context teams need to act.
- **Close the find-to-fix loop with re-validation built in:** Whether using autonomous or human-validation remediation, Picus provides vendor-specific fixes and fast re-validation after fixes are deployed so teams know the gap is actually closed, not just flagged.
- **Defensible, continuously updated evidence of security effectiveness:** AI-validated posture data mapped to business asset criticality, regulatory frameworks, and board-level risk metrics that replaces assumption-based reporting with real-time, adversarially tested evidence.

5. How Picus Does It

(aka the secret sauce)

- **Picus Platform** converges offensive and defensive tradecraft into one simple, scalable engine for autonomous exposure validation. It unifies previously separate tools – such as breach & attack simulation and automated pentesting – into a continuous, context-rich workbench for the agentic purple team. The platform has five core pillars:
    1. **Signal-driven:** mobilizes autonomously in response to relevant real-world changes, not human input or schedules
    2. **Context-rich:** every action shaped by your specific environment, assets, and business risks, not theoretical threats or best-guess impact
    3. **Real-time:** responds at the speed of threats to validate what matters most in any given moment
    4. **Closed-loop:** from signal to validation to fix and then back to re-validation in a continuous, self-learning loop without manual handoffs
    5. **Tunable autonomy:** from fully autonomous to fully supervised, configured to your risk tolerance and operational requirements with complete transparency, traceability, and chain of custody

- **Picus Swarm** scales your team using our autonomous AI agents to respond to real-time signals – new threats and intel, CVEs, config drift, asset changes – and continuously craft adversary-informed attack simulations, validate your defenses, and automatically identify which exposures need immediate remediation and which can be mitigated with proven compensating controls while you work toward permanent fixes. Crucially, Picus Swarm provides full traceability: details on every autonomous test executed, every finding validated, every remediation mapped to specific controls. No hallucinated attack paths and no unexplainable ratings or scores – and all of it tunable based on your guardrails and your requirements.
- **Picus Score** redefines exposure criticality by combining CVSS, EPSS, and KEV with real-world exploitability proof and your specific asset and business context. Unlike passive vulnerability scanners, Picus synthesizes millions of data points to give a pragmatic matter of what's actually critical in your environment. This gives you defensible evidence of which vulnerabilities genuinely threaten your business so you can confidently prioritize what matters and safely deprioritize theoretical noise.
- **Our security data fabric** leverages 75+ integrations – including CrowdStrike, Microsoft, Palo Alto, and Splunk – to ingest threat intel, vulnerability assessments, and security stack data (EDR, SIEM, firewalls, WAF, cloud security) to build a complete picture of your enterprise's defenses across on-premises, public cloud, and hybrid environments. We ingest context from across your stack and push validated findings back to it to create a closed-loop flywheel rather than yet another silo. The data fabric is paired with our constantly updated threat library (24-hour SLA for critical threats) and vendor-specific mitigation guidance. The result is a dynamic, open ecosystem that doesn't force a false choice between validation and vulnerability assessment & management.
- **The extrapolation engine** uses AI to harness millions of data points to generate a context-rich graph of what matters most to your business. It leverages both your data and the aggregated

findings of all Picus attack simulations to prove exploitability and breach impact without overloading every infrastructure asset, disrupting production, or introducing additional administrative overhead. It leverages the fabric to synthesize the discrete and disparate data into actionable insights that move you beyond noisy alerts, logs, and queues.

6. What Customers Can Achieve as a Result

(aka positive business outcomes)

- **Tune and optimize defenses** with validated evidence of effectiveness and context-rich prioritization for which improvements will make the most difference against both traditional attacks and emerging AI-powered threats. Know within minutes, not days or weeks, how protected you truly are against any new threat.
- **Accelerate mitigation and remediation** by proving exactly which vulnerabilities matter in your exact environments, then autonomously fixing – or providing vendor-specific guidance for humans in the loop – and fast re-validation after gaps were closed.
- **Remove operational burden** and free up teams to focus on more strategic priorities by leveraging Picus Swarm to replace manual validation work.
- **Prove the value of your security investments** and demonstrate readiness with evidence of what's working and what's not so you can confidently validate new tools before purchasing, consolidate or retire underperforming solutions, and defend budget decisions to executives and boards.

7. Proof Picus Can Deliver

(aka Don't Take Our Word For It)

- A **4.9 rating on G2** and recognized as a G2 Leader
- A **4.8 rating on Gartner Peer Insights** (GPI) including the 2025 Gartner Peer Insights Customers' Choice award
- Picus Security is the **#1 Innovation Index leader in the 2026 Frost Radar for Automated Security Validation** and 3rd highest on their Growth Index.
- **SiliconANGLE TechForward 2025 award winner** for Security Validation
- **2025 Award Recipient at the Cybersecurity Excellence Awards** for CTEM and Breach and Attack Simulation
- Customer 1: 98% reduction in critical ticket backlog after Picus
- Customer 2: 89% reduction in MTTR

- Customer 3: 92% fewer SLA violations for high/critical vulnerabilities
- Sutter Health slashed attack simulation times from weeks to <1 hour
- “Picus Customers Prevent 2x More Attacks within Three Months of Deployment”

Boilerplate

100-word version:

Picus Security shows you what's working in your defenses and what isn't, why it matters, and how to fix it. By unifying breach & attack simulation and automated pentesting into one scalable platform, Picus enables autonomous exposure validation to find and fix what truly matters to your business. Picus Swarm scales your team with AI agents that use real-time signals to autonomously validate and remediate within the guardrails you define. With 75+ integrations and the industry's highest customer satisfaction ratings, Picus unlocks the focused, context-rich speed that defenders need to outpace AI-driven adversaries.

50-word version:

Picus Security shows you what's working in your defenses and what isn't, why it matters, and how to fix it. Picus converges attack simulation and automated pentesting into one unified platform for autonomous exposure validation. Security teams use Picus to get the validated, context-rich speed needed to outpace AI adversaries.

What's New for FY26

1. Shifting categories from adversarial exposure validation (AEV; Gartner) or automated security validation (ASV; Frost & Sullivan) to our category of "autonomous validation"

2. Moving from a "portfolio" of "products" to a "platform" with "services." E.g. we might refer to our "security control validation service on the Picus Platform"

3. Clarifying our branded terms/capabilities:

From...	To...
Picus Security Validation Platform; Picus Complete Validation Platform	Picus Platform
Picus Numi AI	Picus Swarm
Picus Exposure Score; Picus Criticality Score	Picus Exposure Score
All other branded terms	Non-branded (e.g. "The Picus data fabric..." vs "Picus Data Fabric")

Key Narratives for FY26

- Validation as the hardest step within CTEM – requires a new approach
- Vulnerability assessments alone are not enough

- Automated pentesting alone is not enough
- With AI, context is everything
- The future is agentic, not just automated
- Real-time > continuous
- Converged BAS + APT + EXV = context-rich platform that's greater than the sum of its parts; just because a competitor says they have a "platform" doesn't make it true

Say This, Not That

Say This...	Not This...	Why
Picus Platform	Picus Security Validation Platform or similar	Need to standardize platform name
Modules	Products	Our platform has modules, such as Security Control Validation. We do not talk about "products" in a "suite" or "portfolio" → they're modules on a platform.
Automated pentesting	APT; pentest automation (PTA)	For our audience, "APT" = "advanced persistent threat". And PTA is not common enough to be used.
Picus Swarm	Picus Numi AI	New brand name
Real-time	Continuous	Use "continuous" selectively → keep in mind we are pivoting from continuous (implies daily or weekly) to real-time
Autonomous	Automated	We are moving from AI-driven automation to AI autonomy via Picus Swarm

Messaging Landmines

Red Flag/Avoid:

- Don't mention any competitors by name unless the asset is specifically competitive-focused and has been approved ahead of time
- Don't pit APT vs BAS directly and force a false choice – treat as complementary, with pros and cons, and push narrative about convergence into a single scalable platform

Yellow Flag/Caution:

- Careful with language around reducing labor hours or headcount – selectively use with higher level audiences (Director or higher) but generally avoid so that we do not give the impression

that Picus could cost you your job. Favor language around freeing up time and allowing cyber teams to refocus on strategic priorities; can also favor messaging around tool consolidation and retiring point solutions.
- Careful with language around using the Picus Score to reprioritize low-scored CVSS vulnerabilities as higher or more critical. This can happen, but we tread carefully in sales & marketing so that we don't give the impression that Picus will create more work for you and add to an already full backlog
- Whenever possible, push Picus and Picus Platform overall. Be selective when talking about specific platform services, such as Security Control Validation or Attack Path Validation. When in doubt, Marketing's efforts should mostly focus on getting buyers excited about Picus and then let Sales qualify with them about which services on the platform are right for them.
- Use the commonly understood terms for our use cases instead of the names of our actual platform services. E.g. say "breach and attack simulation" instead of Security Control Validation; say "automated pentesting" instead of Attack Path Validation, etc.

The Vibes Behind the Messaging

1. We sell to the anxiety of "not knowing," not the anxiety of "being behind." Most vendors exploit fear of falling behind — newer threats, faster attackers, the latest TTP. Picus speaks to a different and more honest fear: the seasoned security leader who has watched "well-protected" companies get breached and knows, in their gut, that tool ownership ≠ actual defense. That buyer isn't anxious about being behind. They're anxious about not knowing. They're not buying Picus because they want to modernize. They're buying it because they want confidence in what they already own.

2. We are the Devil's Advocate, not (just) a vendor. Vendors tell you their product works. Picus asks the uncomfortable question your own team doesn't have time to ask: *but does it actually?* This reframes us from solution-seller to trusted challenger — the one entity in your ecosystem whose entire job is to probe, pressure-test, and prove. That's a fundamentally different relationship than any tool vendor offers.

3. "Faster" on its own is the wrong goal. "Proven" is the right one. The entire industry sells speed: faster detection, faster response, faster patching. We don't – or at least not in the conventional sense. We sell certainty. The implicit challenge in all Picus messaging: what's the point of moving fast if you don't know whether you're moving in the right direction? Speed without validation is just accelerated exposure.

5. Visibility isn't enough. Seeing isn't securing. The visibility vendors made "you can't secure what you can't see" the industry mantra. It's true but incomplete. You can see everything and still be breached — because seeing doesn't tell you whether your controls will actually stop an attack. Picus operates one step further down the chain: from visibility to proof of effectiveness.

6. We respect the intelligence of our buyer. The Picus buyer has been around long enough to be skeptical of every vendor claim. They've been burned by "game-changing" and "paradigm-shifting" tools that under-delivered. Our voice should never talk down to them, never oversimplify, and never promise what can't be proven. The tone is peer-to-peer, not vendor-to-prospect.

Messaging Tone and Style

(Sorry, this is mostly for the LLMs)

Keep the practitioner centered in your mind, even if targeting leaders. Ground every asset in the operational reality of someone who has to actually live with the problem — the analyst drowning in alerts, the architect who knows exactly why the pentest findings sat unactioned for six months. Leaders buy because practitioners push; practitioners advocate when they feel understood, not sold to.

Use direct, plainspoken language — even with experts. Our buyers know the jargon. That's not a reason to use it; it's a reason to earn their respect by not hiding behind it. "Your defenses get tested once a quarter" lands harder than "periodic validation cadences create exposure windows." Concrete beats clinical.

Avoid exaggerated generalities about Picus and the platform. We don't solve all security problems for all organizations at all times — and we don't need to pretend to. This matters most around security outcomes: Picus helps you validate, prioritize, and fix faster. We are not the single answer for achieving resilience. Overclaiming undermines credibility with exactly the buyers we're trying to reach.

The tone is peer-to-peer, not vendor-to-prospect. We are not evangelizing. We're talking to people who have been in security for a decade, have been burned by overpromised tools, and can smell a pitch from a mile away. Write as if the reader is a smart colleague who will immediately notice if something doesn't hold up. Write as if we're both in on it together, we're both knowing insiders who share the same pains, same frustrations, hate the same things about vendors and their overblown promises.

Lead with the problem, not the product. Most Picus copy should spend more words on the pain than on the solution. If a reader gets through the first two sentences and can't tell whether it's a Picus ad or an industry observation, that's often a sign you're doing it right.

Specificity is credibility. "135 new CVEs every day" beats "a growing volume of vulnerabilities." "Breakout in under 30 seconds" beats "adversaries are moving faster than ever." Whenever the doc gives you a number, use it. When it doesn't, don't manufacture vagueness — leave it out or find the real stat.

Avoid the three lazy moves in cyber copy: (1) fear-mongering without a payoff — don't invoke breach anxiety unless you're about to say something specific and useful; (2) "AI-powered" as a standalone claim — it means nothing without context; (3) passive voice around outcomes — "security posture is improved" is not a sentence Picus writes.

Match register to channel and audience. Practitioner-facing copy (emails, social, landing pages) skews shorter, more direct, occasionally wry. Leader-facing copy (exec briefs, sponsored content, board-level materials) can carry more strategic framing and outcome language — but never at the expense of specificity. The voice stays consistent; the altitude shifts.

"Autonomous" is a claim, not a descriptor — use it carefully. We've made it a category term and it carries weight. Don't scatter it as an adjective to make something sound more impressive. Reserve it for contexts where the autonomy is real, specific, and accompanied by the chain-of-custody framing that makes it credible.

Reserve academic register for genuinely academic contexts. Whitepapers, research reports, and technical guides can carry more structural depth — but depth is not the same as formality. Even long-form content should read like a smart practitioner wrote it, not like it was peer-reviewed. If a sentence could appear in a journal abstract without modification, rewrite it. Ask: would someone reading this at their desk, between tickets, feel like this was written for them? If not, it's too academic.

Orient around the user, not the vendor landscape. A common failure mode in cyber content is writing that's more interested in mapping the competitive and analyst ecosystem than in helping the person with the actual problem. Practitioners don't care which analyst firm created a framework or what a competing vendor's positioning is — they care whether their SIEM rules fire and whether their controls would stop an attack. When copy spends more words on vendors, categories, or frameworks than on the practitioner's job to be done, it has lost its audience. Name analysts and frameworks sparingly, and only when they add credibility to a practitioner-relevant point — not to signal market awareness.

Don't invent new named concepts. Coining a term — "Agentic Detection Mesh" or "Contextual Risk Normalization Engine," etc. — feels like it adds precision. It usually adds friction. Practitioners are already navigating more vocabulary than any one vendor should impose on them. If the concept can be expressed in plain language, express it in plain language. Ask: is this a new term because the thing genuinely doesn't have a name, or because naming it makes us sound more sophisticated? If it's the latter, don't name it. Treat the introduction of any new proper noun or capitalized concept as a yellow flag requiring justification.

When in doubt, ask: would a skeptical CISO roll their eyes at this? Would a red or blue teamer sigh heavily? That's the internal test. If the answer is yes, rewrite it.

Personas – their Pains, Priorities, Ambitions, Anxieties

	Rational Challenges Business and Operational Realities	**Psychological Motivators** Human Drivers & Internal Narratives
Practitioners Managers, Engineers, Analysts, Architects	● Manual work: tedious, resource-intensive testing & validation then remediation/mitigating ● Prioritization noise: too much is “critical” but lacks context ● Tool frustration: too many, too disconnected, too time-consuming	● Want to get out of grunt work and have meaningful impact or know their work matters ● Looking to level up their skills ● Tired of the hamster wheel and especially want their tools to work for them / work better for them
Leaders VPs, CISOs, CIOs	● Struggle to defend the spend: need to prove security investments are delivering outcomes and worth the $$$ ● Struggle to prove effectiveness: need to assure C-suite & board of threat readiness ● Regulatory pressure: volume & velocity of compliance changes, now exacerbated by growing 3rd-party and supply chain risk	● Want to be a business enabler with a real, not token, seat at the table ● Looking for confidence to speak on risk, readiness, and resilience without relying on assumptions ● Concerned about credibility and taking the fall when breaches do occur

Appendix

Competitors' Categories

Autonomous Exposure Validation	→ Picus
Automated Security Validation	→ Frost and Sullivan
Adversarial Exposure Validation	→ Gartner
Continuous Agentic Red Teaming	→ Armadin
Automated AI Red Teaming	→ Mindgard
Autonomous Red Teaming	→ Offensai
Autonomous Purple Teaming	→ Skyhawk
Autonomous Pentesting	→ Horizon3
Autonomous Offensive Security	→ XBOW
Autonomous Offense	→ Tenzai
Agentic Pentesting	→ Vulnetic
Agentic Vulnerability Mgmt	→ Cogent
Agentic Exposure Ops	→ Nagomi
Agentic Offensive Security	→ Terra Security; Ridge Security
AI Pentesting	→ Novee
AI-Native Exposure Management	→ Zafran
AI-Powered Security Validation	→ Pentera
AI-Powered Defense Mesh	→ Tuskira
AI-Powered Threat Exp. Mgmt	→ SafeHill

Supporting data points

Maturity, Readiness, and Confidence

"**92% of organizations struggle with essential resilience-building efforts, such as pressure-testing defenses, understanding emerging threats** and establishing rapid response mechanisms." - *Accenture State of Cybersecurity 2025*

"**84% [of organizations] struggle to develop and operationalize cyber risk strategies** that align with their transformation goals." - *Accenture State of Cybersecurity 2025*

9 out of 10 organizations have seen cybersecurity budgets increase; 1 in 3 saw budget increases of 30% or more in the last two years. Yet despite investment, "**only 34% feel very confident in the resilience of their company's current cybersecurity infrastructure** against attacks." – *Cisco Cybersecurity Readiness Report 2025*

The Maturity Stagnation: Despite heavy investment, only **4% of companies** have reached a "Mature" level of cybersecurity readiness; **70% remain in the "Beginner" or "Formative" categories** – *Cisco Cybersecurity Readiness Report 2025*

"The 2026 Gartner Board Survey shows 90% of [board] directors lack a measure of confidence in cybersecurity value." – where value is the balance of "level of protection" and "cost of cybersecurity." – *2026 Gartner Board Survey*

83% of cybersecurity programs show no measurable results. – Gartner

VM, CVEs, and Exploit/Attack Dynamics

*"**27% of all CVEs published since 1999 were released in just the past two years. 135 new CVEs every day** — a 40% year-on-year increase."* – Cytidel (Nov 2025)

Mean Time To Remediate (MTTR) is 55-72 days, during which time another 7,000-9,000 new CVEs were published. – *Edgescan Vulnerability Statistics Report 2025*

Teams are often focused on the wrong things: vulnerabilities with a **high probability of exploit (EPSS >0.7)** have a higher average MTTR of **109.4 days**, while those with a low probability (EPSS <0.1) are remediated in an average of **75.8 days** – *Edgescan Vulnerability Statistics Report 2025*

Only takes adversaries **5.4 days to exploit a vulnerability** after a new disclosure. – *Fortinet Threat Landscape Report 2025*

The average enterprise only remediates **15% of their vulnerabilities** in a given month. – Empirical Security

Average breakout time dropped to 29 minutes in 2025, representing a 65% increase in speed compared to the 48-minute average recorded in 2024. Breakout times have steadily decreased by roughly 70% since 2021. The fastest observed breakout time in 2025 was a staggering 27 seconds, down from 51 seconds the previous year. – *CrowdStrike Global Threat Report 2026*

Unit 42 ran a simulated attack integrating generative AI at each stage. They found that **AI reduced the time-to-exfiltration to just 25 minutes**—approximately 100 times faster than the traditional real-world median of two days. – *Palo Alto Unit 42 Global Incident Response Report 2026*

Tools and Efficiency

After "AI Automation," **"Tools Rationalization" is the 2nd highest technology priority** – "This approach involves evaluating and consolidating duplicate security tools across platforms to eliminate redundancies and improve efficiency while reducing costs." – *Wipro State of Cybersecurity 2025*

77% say adopting too many security solutions has actually slowed down their ability to detect and respond to incidents – Cisco

70% of companies now maintain a stack of >10 point solutions, and 26% are juggling 30+ different security tools – Cisco

"Operational efficacy" (does it actually work?) has become the single most important factor for buyers – Futurum Research

AI and AI-driven Risk

"Most (81%) CIOs and technology executives indicated their enterprise will have deployed agentic AI by mid-2027, and 17% have already rolled out the technology" as of January 2026. – Gartner *Emerging Risk Deep Dive: Agentic AI*

90% of organizations lack the maturity to defend against modern AI-driven threats – Accenture

"Two disruptors — the shift from **cyberprotection to cyber resilience**, alongside **urgent needs to adopt and secure AI** — are expanding CISOs' responsibilities to protect and enable their organizations." – *Gartner Top Trends in Cybersecurity for 2026*