HACKTHEBOX

CISO Briefing

Collision Course 2026

50 days in: Three forces reshaping cybersecurity

A research-led analysis for security leaders on the convergence of agentic AI, geopolitical instability, and the talent-budget paradox.

Executive summary

The threat landscape facing enterprise security organizations is accelerating at a rate that outpaces the corresponding growth in defensive investment. This briefing examines three macro forces converging in 2026: the emergence of agentic AI as both an offensive weapon and an unmanaged attack surface; the escalation of nation-state activity targeting civilian infrastructure; and a persistent talent-budget paradox in which organizations increase security spending yet report diminishing confidence in their defensive posture.

1

Agentic AI

The autonomous attack surface

2

Geopolitical instability

Conflict, deregulation, and nation-states

3

The talent-budget paradox

Burning money, starving for people

Individually, each force warrants strategic re-evaluation. In combination, they produce compound risk: autonomous AI compresses the intrusion lifecycle from weeks to minutes, nation-state actors pre-position inside critical infrastructure for years, and the security teams tasked with defending against both remain understaffed and subject to significant attrition. Organizations that prioritise measurable workforce capability - rather than incremental tooling investment - are better positioned to operate effectively in this environment.

Key recommendations

1. **Transition from periodic to continuous assessment**
   Annual tabletop exercises are insufficient against adversaries operating at machine speed. Organizations should adopt ongoing, high-fidelity simulation that maintains pace with the evolving threat landscape.
2. **Rebalance investment toward workforce capability**
   Evidence indicates that the return on a skilled, retained security workforce exceeds incremental tooling expenditure. Organizations should evaluate whether current spend allocation reflects this finding.
3. **Adopt performance-based readiness metrics**
   Security leaders should develop measurable cyber-performance benchmarks that verify operational readiness rather than compliance status alone. Regulatory alignment becomes a by-product of demonstrated capability.
4. **Assume elevated threat posture for CNI-adjacent organizations**
   Organizations with dependencies on critical national infrastructure - energy, water, healthcare, financial services - should plan on the basis that they are already within adversary targeting scope.

Force 1

Agentic AI: The autonomous attack surface

Agentic AI systems have introduced a dual-use dynamic in which the same autonomous capabilities that enhance defensive operations also expand the offensive attack surface. Gartner projects that **40% of enterprise applications** will feature task-specific AI agents by 2026, up from less than 5% in 2025 - yet only **6% of organizations** have an advanced AI security strategy in place. Forrester's headline cybersecurity prediction for 2026 is that an agentic AI deployment will cause a publicly reported breach leading to employee dismissal, citing cascading failure risk as the primary concern. The Verizon 2025 DBIR found that **15% of employees** regularly access generative AI tools on corporate devices, with 72% using personal email accounts - bypassing enterprise data-loss controls entirely.

51 sec
fastest observed breakout time (CrowdStrike 2025)

86%
of organizations breached in 2024 (Fortinet)

4.8M
unfilled cybersecurity roles globally (ISC2)

On the adversary side, the acceleration is measurable. CrowdStrike's 2025 Global Threat Report documents a fastest observed breakout time of **51 seconds**. Palo Alto Networks Unit 42 demonstrated an AI-powered ransomware simulation achieving full compromise-to-exfiltration in 25 minutes - a 100x speed increase over traditional methods. In 20% of Unit 42's incident response cases, data exfiltration occurred within the first hour. ENISA's 2025 Threat Landscape reports that AI-supported phishing now accounts for over 80% of observed social engineering activity.

The capability trajectory extends to autonomous vulnerability discovery. In February 2026, Anthropic's Claude Opus 4.6 autonomously identified **over 500 zero-day vulnerabilities** in open-source repositories, with each flaw validated without human intervention. Gartner predicts that by 2027, AI agents will **reduce the time to exploit account exposures by 50%**, while also forecasting that over **40% of agentic AI projects will be cancelled by end of 2027** due to governance and security failures.

Shadow AI compounds the exposure: **34%** of organizations now report generative AI-related data leaks (Splunk), while Fortinet finds **48%** of IT decision-makers identify insufficient AI expertise as their primary implementation barrier.

IN THE NEWS

AI Model Uncovers 500+ Zero-Days in Open-Source Code

FEB 5, 2026 · AXIOS / FORTUNE

Anthropic's Claude Opus 4.6 autonomously discovered over 500 previously unknown high-severity vulnerabilities in open-source software - including projects fuzzed for millions of CPU-hours. Deployed in a sandboxed VM with no custom instructions, the model parsed Git commit histories, identified unpatched code paths, and wrote proof-of-concept exploits. Every finding was validated by Anthropic's Frontier Red Team or external researchers. Anthropic acknowledged the capability is 'inherently dual-use' and warned that 90-day disclosure windows may not hold against AI-speed discovery.

Axios, Fortune, The Hacker News, Anthropic Frontier Red Team

Key insight

Incident response models designed for human-speed threat timelines require fundamental reassessment. Detection strategies should evolve to identify autonomous agent activity.

Force 2

Geopolitcal instability: The new frontline

Nation-state actors have moved beyond espionage to the systematic pre-positioning of offensive capabilities within civilian infrastructure. Gartner identifies **geopolitical tensions and regulatory volatility** as a top cybersecurity trend for 2026, noting that 'cybersecurity leaders are navigating uncharted territory as these forces converge.' The Verizon 2025 DBIR records a **163% increase** in espionage-related breaches, now accounting for 17% of all incidents. Mandiant's M-Trends 2025 identifies the financial sector as the most targeted industry (17.4%), followed by business services (11.1%) and healthcare (9.3%).

Forrester predicts that **five governments will nationalise or restrict critical telecom infrastructure** in 2026, a direct response to the Salt Typhoon espionage campaign that breached over 600 organizations across 80 countries. Independent polling finds **72% of IT leaders** now consider it plausible that nation-state cyber capabilities could escalate into full-scale cyberwar, with critical infrastructure anticipated as the primary target. Cyble recorded a **51% increase** in hacktivist sightings in 2025, with these groups increasingly targeting industrial control systems.

1,162
cyberattacks on US utilities in 2024 alone

5 yrs
Volt Typhoon inside US infrastructure undetected

1 death
confirmed linked to NHS ransomware (Synnovis 2024)

Persistent infrastructure compromise. The China-attributed Volt Typhoon campaign maintained access to US water, power, and telecommunications networks for over five years prior to detection, pre-positioned to disrupt services at a time of strategic choosing. CISA characterised the approach as 'living off the land,' using legitimate system tools to evade conventional detection methods. Mandiant's M-Trends 2025 confirms that three of the top four most exploited vulnerabilities in 2024 were in network edge devices - VPNs and firewalls - exploited as zero-days before patches were available.

Financial system disruption. Ukraine's HUR intelligence directorate successfully disrupted ATM access, card payment processing, and mobile banking services across more than ten major Russian financial institutions, including Sberbank and Alfa-Bank. The incident demonstrated that targeted disruption of financial services can produce cascading societal effects with minimal kinetic activity.

Water and energy infrastructure targeting. In Denmark, the Z-Pentest group compromised water utility systems, resulting in burst pipes affecting 50 households. In the United States, Iran-linked CyberAv3ngers exploited programmable logic controllers at water utilities, while American Water - the nation's largest water provider serving 14 million customers - was forced to take systems offline. Check Point Research recorded 1,162 cyberattacks on US utilities in 2024, a 70% year-on-year increase.

Healthcare sector impact. The Synnovis ransomware incident affecting NHS services in June 2024 resulted in one confirmed patient death, the disruption of approximately 10,000 appointments, the postponement of 1,710 surgical procedures, and the exposure of 900,000 patient records.

IN THE NEWS

Salt Typhoon Extends to Singapore; Senate Demands Telecom CEOs Testify

FEB 2026 · TECHCRUNCH / AXIOS

Singapore confirmed China-backed Salt Typhoon hackers gained access to critical systems at the country's four largest telecom companies, extending a campaign the FBI attributes to 200+ organizations across 80 countries. Nine US carriers including AT&T, Verizon, and T-Mobile were breached in a multi-year espionage operation that intercepted senior government officials' calls. Senate Commerce Chair Cantwell demanded both CEOs testify, noting neither can prove the hackers have been eradicated. The FBI called it 'one of the most consequential cyber espionage breaches in US history.'

TechCrunch, Axios, US Senate Commerce Committee, FBI/CISA Advisory

Key insight

Every society is three missed meals from chaos. Banks, utilities, and healthcare systems are the new frontline. Organizations with CNI dependencies should assume they are already within.

Force 3

The talent-budget paradox: The resource deadlock

4.8M
unfilled cybersecurity roles globally (ISC2)

71%
SOC analyst burnout rate (Vectra AI)

$550K
additional breach cost with understaffed teams (IBM)

A structural imbalance persists between security investment and defensive outcomes. Gartner forecasts worldwide security spending at **$213 billion in 2025**, rising to **$240 billion in 2026** - yet IANS Research/ Artico Search find that cybersecurity budgets grew only **4% in 2025** (down from 8%), and as a share of IT budgets **fell from 11.9% to 10.9%**, breaking a five-year upward trend. The paradox is clear: absolute spending rises while relative priority declines. The ISC2 Cybersecurity Workforce Study identifies **4.8 million unfilled roles** globally, with the workforce needing to grow by 87% to meet current demand. IBM's Cost of a Data Breach Report finds that understaffed security teams incur an average of **$550,000 in additional breach costs**, while Gartner attributes a **lack of cybersecurity professionals** as the factor responsible for more than 50% of significant security incidents.

ISC2 reports that budget constraints have overtaken talent scarcity as the primary workforce challenge in 2025. Hiring requirements remain rigid - 90% of managers require prior IT experience and 89% mandate certification - yet Fortinet finds organizational willingness to fund certifications has declined from 89% to 73%, eroding the mid-career pipeline. Gartner predicts that by 2027, organizations implementing cybersecurity-specific resilience programmes will experience **50% less burnout-related attrition**. Operationally, SOC teams process an average of 11,000 alerts per day, of which Vectra AI estimates only 19% warrant investigation. Analyst burnout has reached 71%, with average tenure of three to five years.

There is countervailing evidence of defensive progress: Chainalysis reports that ransomware payments declined **35% in 2024** to $813.55 million, driven in part by increased victim refusal to pay - now at **64%** according to Verizon's DBIR. Gartner finds that organizations adopting Continuous Threat Exposure Management (CTEM) are **three times less likely to suffer a breach** by 2026 - the strongest evidence linking continuous methodology to measurable risk reduction.

However, the volume of ransomware incidents continues to grow, with ransomware present in **44%** of all breaches reviewed (Verizon DBIR), up from 32% the prior year. Operational pressure on defensive teams remains acute.

IN THE NEWS

CISA Loses One-Third of Workforce as Threats Escalate

JUNE 2025 · AXIOS / NEXTGOV

CISA - the US government's frontline civilian cyber-defence force - has lost ~1,000 staff through buyouts, deferred resignations, and layoffs, reducing headcount from 3,292 to an estimated 2,324. The FY2026 budget proposes eliminating CISA's Election Security Programme and cutting the DOE cybersecurity office by 19%. ISC2's survey of 16,029 professionals found 47% of cybersecurity workers have experienced cutbacks, with 71% reporting negative workload impact. The reductions come as Salt Typhoon and Volt Typhoon demonstrate nation-state adversaries operating inside US infrastructure.

Axios, Cybersecurity Dive, Nextgov/FCW, ISC2 Workforce Study 2025

Key insight

Outsourcing alone does not resolve the structural deficit. Organizations should consider rebalancing investment from incremental tooling toward measurable workforce.

Implications & recommended actions

Scope 1 Implications for cybersecurity strategy

How these forces reshape the security investment, insurance, and market landscape.

Force	Implication	Action
Agentic AI	The capacity to train, evaluate, and govern AI-augmented security operations will increasingly differentiate organizational resilience.	Establish measurable workforce performance benchmarks against AI-era threat scenarios. Organizations that can demonstrate validated capability will hold a strategic advantage.
Geopolitical instability	Nation-state threat activity requires operational readiness, not theoretical compliance.	Align team capabilities to recognised frameworks (e.g., MITRE ATT&CK, NIST CSF) such that audit readiness is a by-product of demonstrated skill.
Talent-budget paradox	Gartner finds CTEM adopters are 3x less likely to suffer a breach - yet security's share of IT budgets fell from 11.9% to 10.9% (IANS Research).	Adopt measurable cyber-performance metrics - validated proficiency, observable progression - as the primary return-on-investment framework for board-level reporting.

Scope 2 Implications for defensive teams

How these forces affect operations, tooling, staffing, and board-level reporting.

Force	Implication	Action
Agentic AI	At 51-second breakout times, adversary speed exceeds the capacity for on-the-job learning during live incidents.	Implement continuous simulation-based drilling on emerging adversary techniques to maintain response readiness.
Geopolitical instability	State-aligned threat actors are targeting ICS/OT systems. Most defensive teams were recruited for IT environments; OT expertise remains scarce and costly.	Invest in cross-domain training that develops adaptive capability across IT and OT environments, rather than narrow specialisation.
Talent-budget paradox	Budget constraints are now the primary workforce challenge (ISC2). Organizations cannot recruit their way out of the deficit, nor sustain current attrition rates.	Automate routine Tier-1 alert processing; redirect savings toward structured skills development programmes that improve both capability and retention.

Scope 3 Reshaping cybersecurity careers

How these forces affect career trajectories, skills requirements, and professional development pathways.

Force	Implication	Action
Agentic AI	Demand for AI security specialists is accelerating. Professionals who can validate autonomous AI decisions represent the highest-value capability within SOC teams.	Implement continuous simulation-based drilling on emerging adversary techniques to maintain response readiness.
Geopolitical instability	Defence and critical national infrastructure roles carry high demand but elevated burnout risk. Dual US/EU compliance expertise and OT/ICS proficiency command premium compensation.	Structured, role-based development pathways enable accelerated progression through the mid-career experience gap.
Talent-budget paradox	90% of hiring managers require prior experience, yet entry-level positions are contracting - creating a structural barrier to workforce growth.	Demonstrable, validated capability - rather than tenure alone - provides the most effective pathway through credentialing barriers.

Sources

Gartner Top Cybersecurity Trends 2026 · Gartner Security Spending Forecast 2025–26 · Gartner Cybersecurity Predictions 2024–28 · Forrester Predictions 2026: Cybersecurity & Risk · IANS Research / Artico Search Security Budget Benchmark 2025 · CrowdStrike Global Threat Report 2025 · Verizon Data Breach Investigations Report 2025 · Mandiant M-Trends 2025 · Palo Alto Networks Unit 42 Global Incident Response Report 2025 · ENISA Threat Landscape 2025 · Fortinet Global Cybersecurity Skills Gap Report 2025 · Chainalysis Crypto Crime Report 2025 · Splunk State of Security 2025 · Anthropic Frontier Red Team / red.anthropic.com (February 2026) · CISA/FBI Volt Typhoon Advisory 2024 · Check Point Research 2024 · Cyble Hacktivist Threat Report 2025 · NHS England Synnovis Incident Report 2024 · ISC Cybersecurity Workforce Study 2024 · IBM Cost of a Data Breach Report 2024 · Vectra AI State of Threat Detection 2024 · WEF Global Cybersecurity Outlook 2025–26

Self-assessment

CISO 60-second scan

Ten scenarios drawn from this briefing. Score each honestly - the pattern matters more than any single answer.

Scoring: **5** = Fully confident **4** = Largely **3** = Some gaps **2** = Significant gaps **1** = Not at all

#	Scenario	Score (1-5)
1.	Agentic AI An autonomous AI agent exfiltrates data from a misconfigured internal service at 2 a.m. Your SOC has **under 60 seconds** before the payload stages. Could your team detect and contain it without human escalation?	
2.	Agentic AI A departing employee has used a personal-account GenAI tool to summarise client contracts for six months. HR notifies you on their last day. Do you know **what data left** - and can you prove it to the board?	
3.	Geopolitical An advisory warns that a nation-state actor has lived inside your sector's supply chain for **up to five years**. Could your team spot lateral movement via legitimate system tools - or would it look like normal traffic?	
4.	Geopolitical Your payment processor goes offline in a state-backed disruption campaign across multiple countries. Does your **business continuity plan** cover simultaneous cyber and operational disruption - and has it been tested this year?	
5.	Talent-budget Your two most experienced incident responders resign in the same month. The hiring pipeline needs six months. How confident are you in the remaining team's ability to handle a **major incident** during that gap?	
6.	Talent-budget The board asks for evidence - not anecdote - that last year's security spend improved **measurable defensive capability**. Can you produce that evidence today?	
7.	Cross-cutting A critical zero-day drops at 8 p.m. Friday on a network edge device you run. Exploits are already in the wild. Is your team drilled to **patch, monitor, and hunt** across the weekend without burnout-driven errors?	
8.	Cross-cutting Regulators ask for evidence your team can operate against MITRE ATT&CK techniques used by groups targeting your sector. Can you provide **validated, performance-based evidence** - not just a checklist?	
9.	Cross-cutting An adversary pivots from IT into your OT environment. Your security team was hired for enterprise IT. How confident are you in their ability to **detect and respond across both domains** without third-party support?	
10.	Cross-cutting A competitor suffers a major breach through the same technology stack you run. Could you brief the board within 48 hours on **specific exposure, mitigations, and residual risk** - backed by testing evidence?	

Interpretation: **40–50**: Strong posture - validate with live testing. **25–39**: Material gaps - prioritise lowest-scoring areas. **<25**: Significant exposure - urgent capability review recommended.

The strategic case

Cybersecurity as a Business Enabler

The recommendations in this briefing are framed as defensive actions - but their ultimate impact is commercial. Organizations that treat cybersecurity capability as an operational asset, rather than an insurance cost, unlock easurable business value across six dimensions.

Business dimension	Impact of recommendation
Revenue Protection & Continuity	**Continuous assessment** reduces the probability and blast radius of breach. IBM's 2024 data puts the average breach cost at **$4.88M**; organizations with incident response teams and tested plans saved **$2.66M per incident**. Every day of avoided downtime protects revenue, customer trust, and share price.
Insurance & Risk Transfer	Cyber insurers increasingly price premiums on **demonstrated capability**, not just controls checklists. Organizations that can evidence continuous drilling, framework alignment, and performance metrics secure **broader coverage at lower cost** - turning readiness investment into direct P&L benefit.
Regulatory Advantage	Teams aligned to MITRE ATT&CK and NIST produce **audit readiness as a by-product** of operational competence - eliminating the annual compliance scramble. As NIS2, DORA, and SEC disclosure rules tighten, this becomes a structural advantage over competitors still treating compliance as a separate workstream.
Talent Economics	Gartner predicts that resilience programmes cut burnout attrition by **50%**. At an average replacement cost of **$200K+** per security hire (recruitment, onboarding, ramp time, lost productivity), structured development pathways pay for themselves within a single retention cycle - while building deeper institutional capability.
Board Confidence & Investor Signal	Performance-based metrics give boards what compliance dashboards cannot: **evidence that spend translates into capability**. In M&A due diligence, validated security posture accelerates deal timelines and supports valuation. For public companies, demonstrable resilience is increasingly a factor in ESG scoring and institutional investor assessment.
Competitive Differentiation	In regulated sectors - financial services, healthcare, defence, CNI - **security maturity wins contracts**. Enterprise buyers and procurement teams increasingly require evidence of validated capability, not self-attested questionnaires. Organizations that can demonstrate tested, benchmarked readiness convert security investment into **sales pipeline advantage**.

Key insight

The question is not whether your organization can afford to invest in cybersecurity capability. It is whether it can afford the commercial consequences of not doing so.

HACKTHEBOX

Hack The Box is the leading AI-powered cybersecurity readiness and upskilling platform for humans and AI agents, trusted by 1,500+ organizations worldwide, including Fortune 500 enterprises, government agencies, and MSSPs, to build cyber resilience at scale.

Through AI-enhanced intelligence, gamified labs, live-fire simulations, and one of the world's largest cybersecurity communities, Hack The Box enables teams and AI agents to master offensive and defensive skills through real-world scenarios designed for the age of AI.

Founded in 2017, Hack The Box has grown a global community of over 4 million members, enabling organizations to validate resilience, mitigate breach risk, and develop cyber talent.

For more information,
please visit hackthebox.com