EXTRAHOP

WHITE PAPER

# A CISO's Guide to Geopolitics and Cybersecurity 2026 Edition

## Understanding the Cyberwarfare Strategies of Russia, China, Iran, and North Korea

## TABLE OF CONTENTS


**Executive Strategy: The Era of Permanent Gray Zone Conflict . . . 3**

**The New Economics of Cyber Conflict and AI . . . 6**

**Adversary Analysis . . . 10**

Russia . . . 12

China . . . 14

Iran . . . 16

North Korea . . . 17

**Industry Exposure and the Visibility Crisis . . . 19**

**NDR and the Agentic Enterprise . . . 22**

**The Future of Defense: Network, Context, and Machine-Speed Security . . . 24**

**Resources . . . 26**

**Resources, continued . . . 27**


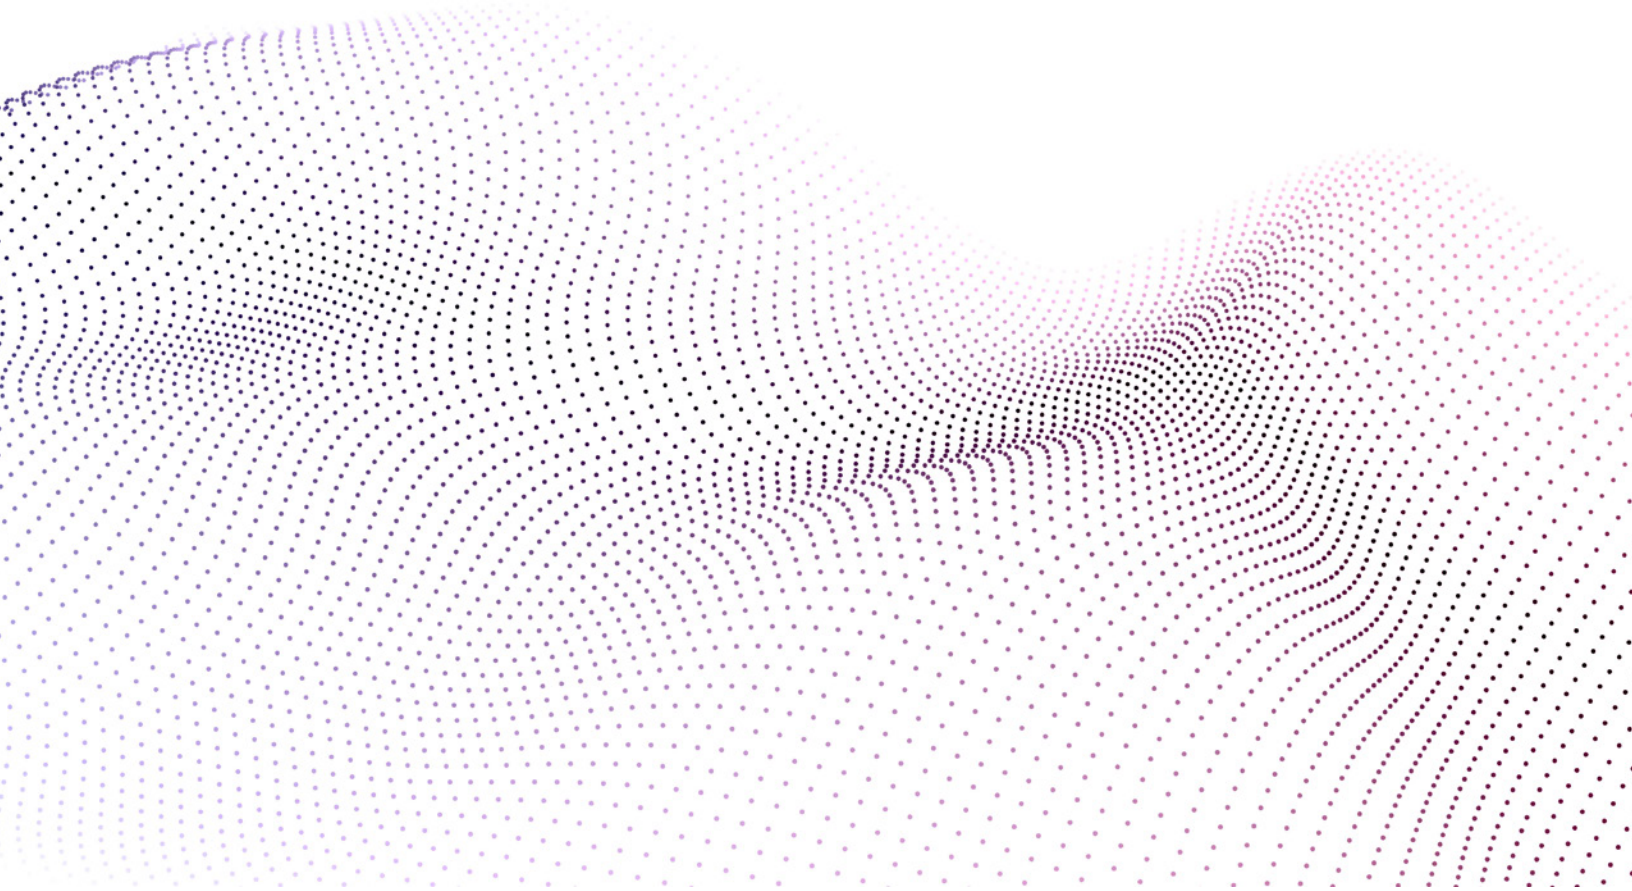

# Executive Strategy: The Era of Permanent Gray Zone Conflict

Cybersecurity has entered a new strategic phase. The defining feature of this phase is not a single breakthrough technology, a single adversary, or a single category of attack. It is the convergence of geopolitical instability, economic coercion, artificial intelligence, and persistent adversary activity into a condition of permanent gray zone conflict. For CISOs and executive teams, this means cyber risk can no longer be treated as a technical problem that sits inside the security organization. It is now a core business risk, a geopolitical exposure, and an operational resilience challenge.

## Four Forces Creating Permanent Gray Zone Conflict

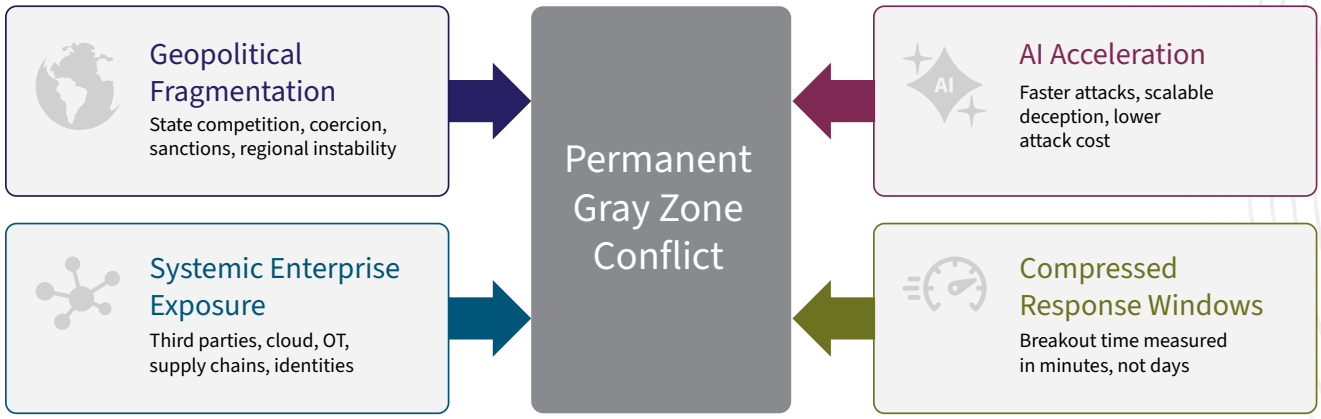


Figure 1: The 2026 threat environment is shaped by converging geopolitical, technological, and operational pressures.
*Concept based on World Economic Forum 2026 and industry threat reporting.*

This report concludes that the organizations best prepared for the 2026 threat environment will be those that recognize three realities.

- First, cyber conflict is becoming continuous rather than episodic. State-aligned actors, criminal groups, and hybrid proxy networks increasingly operate below the threshold of open conflict, using cyber operations to steal, influence, disrupt, and pre-position for future crises.
- Second, artificial intelligence is changing the economics of attack. AI does not eliminate the fundamentals of cyber defense, but it increases the speed, scale, and adaptability of adversary activity.
- Third, traditional security models are under extreme strain because they were built for slower, more bounded, and more observable attacks. In a world where adversaries use legitimate credentials, trusted infrastructure, encrypted traffic, and living-off-the-land techniques, visibility and context become strategic requirements rather than optional security capabilities.

The strategic shift begins with the erosion of the boundary between cyber risk and geopolitical risk. Geopolitical fragmentation is reshaping the cyber landscape by increasing the number of threat actors with incentives to use cyber operations as instruments of influence and coercion. The World Economic Forum's 2026 Global Cybersecurity Outlook identifies accelerating AI adoption, geopolitical fragmentation, and widening cyber inequity as central forces reshaping the global cyber risk environment. For enterprises, the implication is direct: Business operations, supply chains, technology platforms, and data repositories may become targets not only because of their standalone value, but also because of their role in a broader political or economic system.

For CISOs, this changes the meaning of "preparedness." The traditional model of cybersecurity often assumed that the objective was to keep attackers out, detect the exceptions, and respond when an incident occurred. That model is no longer sufficient. The more realistic objective is to operate under the assumption that adversary activity will occur, systems will be exposed, credentials will be compromised, and trusted relationships will be abused. The executive question becomes: How quickly can the organization see what is happening, understand what matters, contain the impact, and continue operating?

## From Incident Response to Continuous Conflict

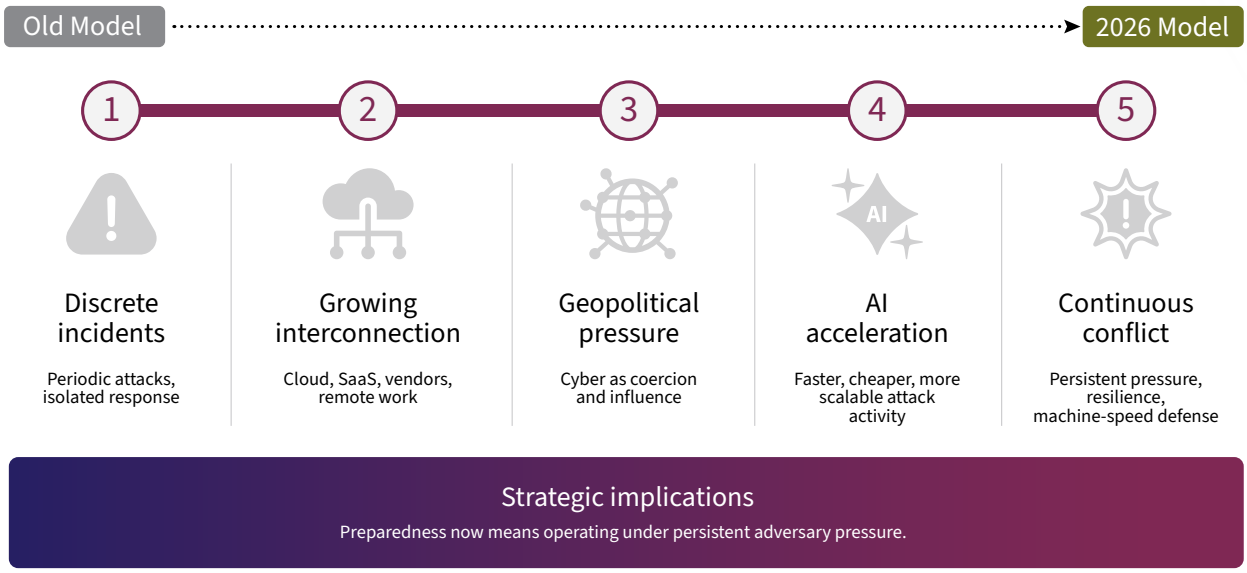


Figure 2: Cybersecurity has shifted from episodic response to continuous operational pressure.

The available breach data supports this shift in mindset. Verizon's 2026 Data Breach Investigations Report analyzed more than 31,000 real-world security incidents, including more than 22,000 confirmed data breaches across 145 countries. Its findings show a threat environment defined by systemic exposure rather than isolated control failure: Exploitation of vulnerabilities became the most common initial access vector for breaches at 31%; credential abuse remained pervasive across breach progression, appearing in 39% of breaches when measured beyond the first known action; ransomware appeared in 48% of all breaches; third-party involvement also reached 48% of total breaches; threat actors were observed using generative AI across stages of attack, including targeting, initial access, vulnerability research, and tool development; and the human element remained present in 62% of breaches.

The lesson is that modern intrusions are rarely reducible to a single failure. They are chains of exposure, identity abuse, misconfiguration, third-party dependency, social engineering, and delayed detection.

The executive strategy, therefore, should not be framed as "prevent every breach." That goal is unrealistic and can lead to brittle security planning. A better strategy is to reduce the probability of compromise where possible, reduce attacker dwell time when compromise occurs, reduce blast radius through segmentation and least privilege, and reduce business impact through resilience planning. The CISO's role is to connect those technical priorities to business outcomes: continuity of operations, protection of intellectual property, trust in digital services, regulatory readiness, and board-level confidence.

This is why visibility becomes the connective tissue of modern defense. Visibility is not merely the ability to collect logs or deploy more tools. It is the ability to understand relationships: Which systems are communicating, which identities are being used, which data is moving, which services are exposed, which vendors are connected, and which behaviors deviate from normal operating patterns. In cloud, hybrid, OT, and unmanaged device environments, network-level visibility becomes especially important because many assets cannot support agents, many attacks use valid credentials, and many signals are only meaningful when viewed in context.

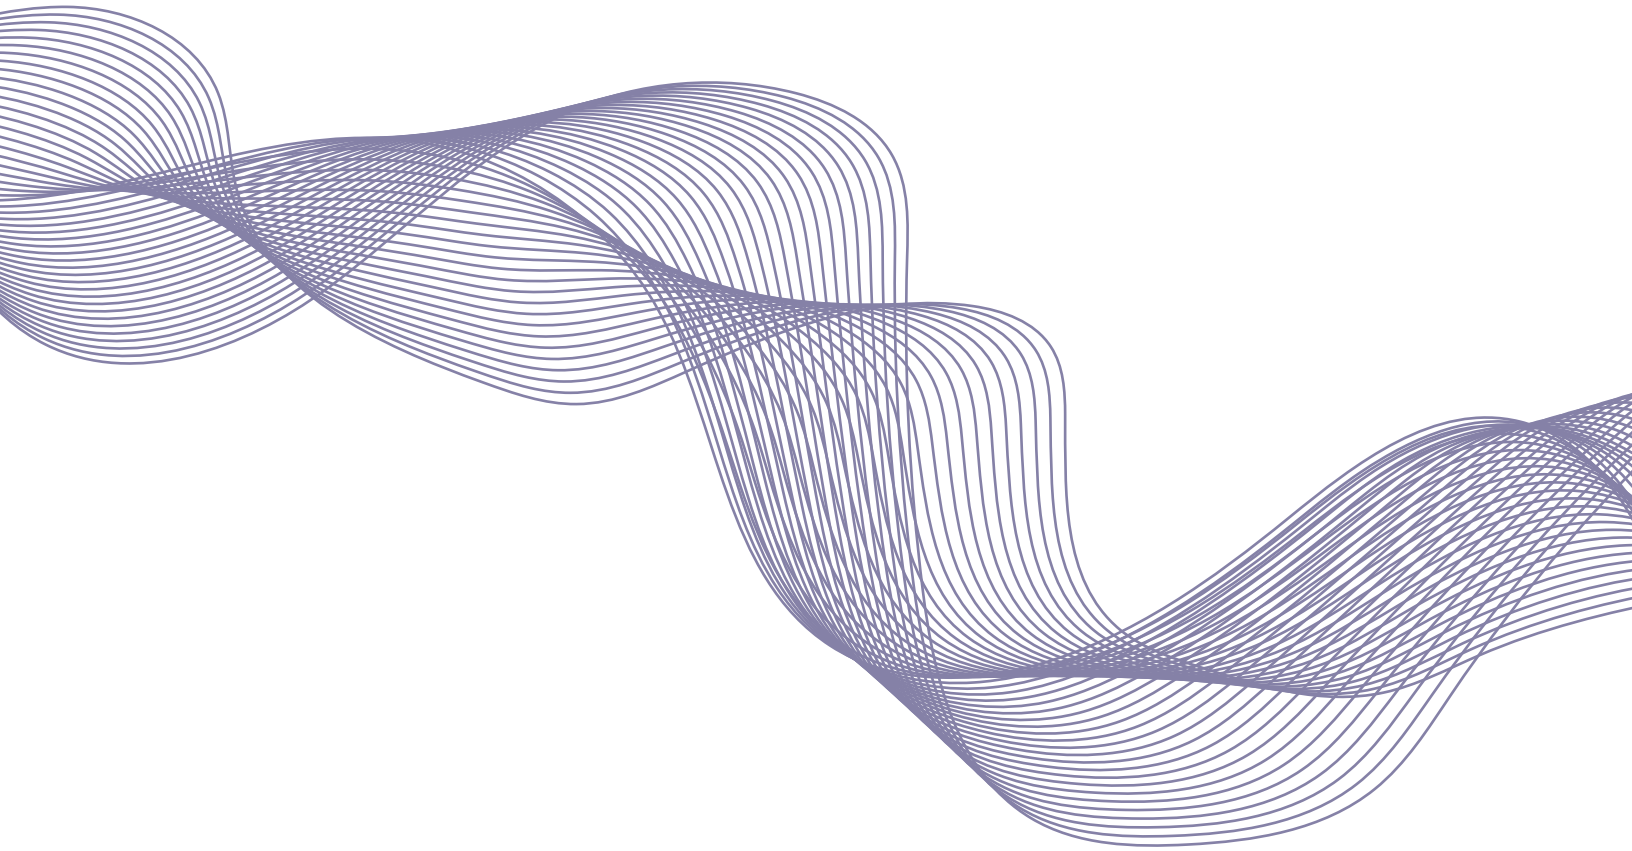

## The New Economics of Cyber Conflict and AI

The economics of cyber conflict have changed. For decades, attackers benefited from a basic asymmetry: It was cheaper to find one weakness than it was for defenders to secure every system, identity, application, vendor connection, and user workflow. In 2026, that asymmetry has widened. Geopolitical competition has increased the strategic value of cyber operations. Criminal markets have professionalized the monetization of access. Artificial intelligence has lowered the cost of reconnaissance, social engineering, vulnerability research, and attack execution. And modern enterprise environments have become too dynamic for security models built around static assets, isolated alerts, and manual response.

The financial incentives remain powerful because the cost of disruption keeps rising. Cisco's Splunk reported that unplanned downtime now costs Global 2000 companies roughly $600 billion annually, with an average impact of about $15,000 per minute and nearly $95 million in lost revenue per organization each year. The same analysis found that downtime can also create reputational harm, customer loss, regulatory exposure, and stock-price effects, which helps explain why extortion remains such an effective pressure tactic. Attackers do not need to encrypt every system to create leverage. They can steal data, threaten disclosure, interrupt operations, pressure customers, or use regulatory exposure as part of the coercion strategy. The business impact can include downtime, legal costs, incident response spending, customer churn, lost revenue, insurance disruption, regulatory scrutiny, and long-term reputational harm.

Third-party exposure compounds that economic risk. As organizations depend more heavily on cloud services, SaaS applications, managed service providers, identity platforms, software vendors, and external data processors, the enterprise perimeter becomes less meaningful. A single compromise can scale across many victims if it occurs at a supplier, identity provider, managed service provider, widely deployed software platform, or cloud integration layer.

The core strategic point is simple: Cyber conflict is becoming faster, cheaper, and more scalable for the attacker, while becoming more complex, interconnected, and consequential for the defender. The CISO's challenge is no longer only to prevent compromise. It is to understand how the economics of attack are changing and to reshape defense around the points where adversary efficiency can be disrupted.

Attackers understand the supply chain. They do not always need to attack the most secure organization directly. They can look for the weaker link that provides access to the same data, the same workflows, or the same operational dependencies. For CISOs, this means vendor risk and third-party monitoring are not procurement exercises. They are a critical part of the core threat model.

AI now accelerates each stage of that supply chain. The World Economic Forum's Global Cybersecurity Outlook 2026 identifies accelerating AI adoption, geopolitical fragmentation, and widening cyber inequity as forces reshaping the global risk landscape. It also warns that attacks are becoming faster, more complex, and more unevenly distributed, increasing pressure on organizations and governments to adapt. For attackers, AI changes the cost curve. Tasks that once required manual effort, specialized language skills, or repeated experimentation can be automated or semi-automated. Reconnaissance can be scaled. Phishing lures can be tailored. Code can be generated or modified. Vulnerabilities can be researched and weaponized in minutes. Stolen data can be sorted and summarized. Defensive telemetry can be studied for gaps. None of this means that AI replaces human attackers entirely. It means AI increases attacker productivity and raises the baseline capability of less sophisticated groups.

## The Attacker Cost Curve is Falling

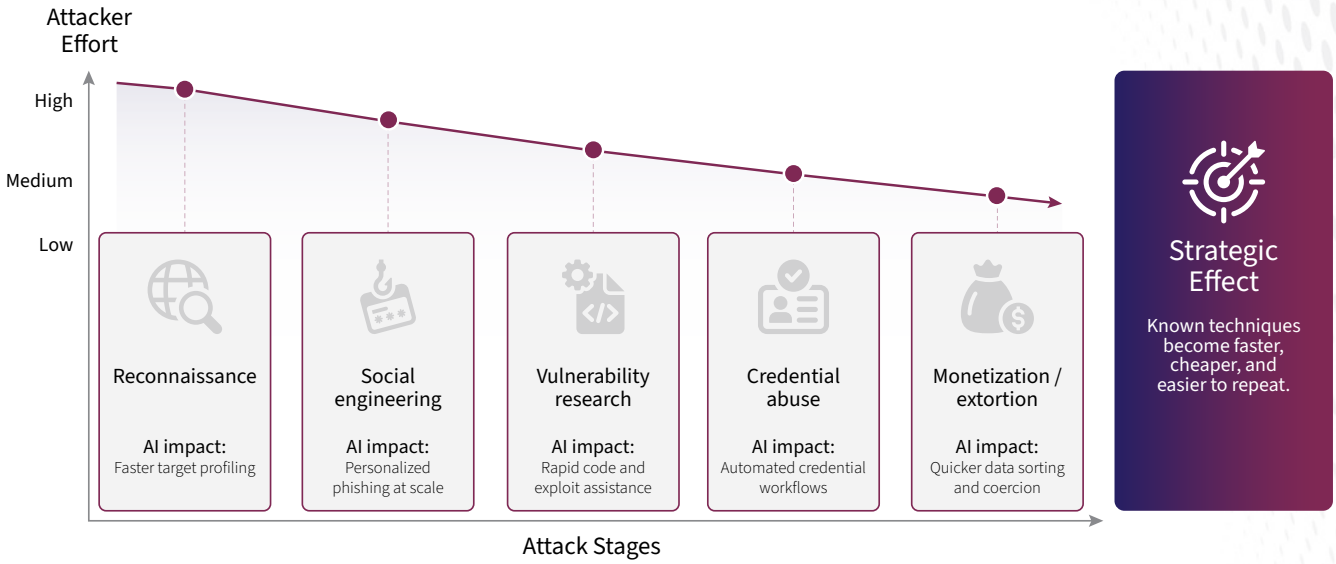


Figure 3: AI lowers the time, skill, and cost required to execute common attack stages.
*Concept based on World Economic Forum 2026 and public reporting on AI-enabled cyber operations.*

The most immediate AI effect is operational acceleration. CrowdStrike's 2026 threat reporting, summarized by multiple technology publications, found that average breakout time fell to 29 minutes in 2025, with the fastest observed breakout taking only 27 seconds. The same reporting described an 89% year-over-year increase in AI-enabled adversary operations and warned that attackers are using AI across reconnaissance, credential theft, evasion, and abuse of trusted identities, SaaS applications, and cloud infrastructure. That matters because breakout time is a practical measure of the defender's shrinking window. If adversaries can move laterally in minutes, security teams cannot depend on manual triage, delayed escalation, or after-the-fact reconstruction as their primary operating model.

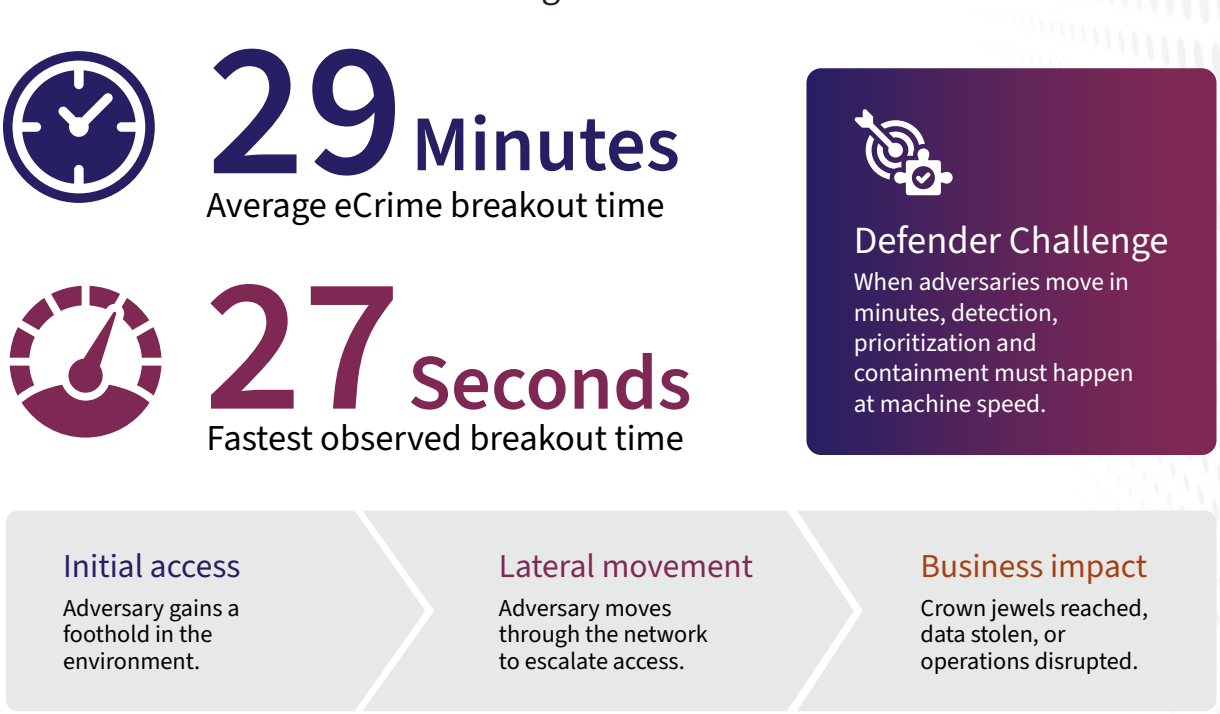


Figure 4: Attack timelines have compressed to the point where manual response alone is not enough.

At the same time, AI is becoming part of the enterprise attack surface. Organizations are integrating AI into development workflows, support operations, analytics platforms, security tools, customer service, productivity suites, and autonomous business processes. These systems often require access to data, applications, APIs, code repositories, tickets, identity systems, and collaboration platforms. That access creates value, but it also creates risk. If an AI system is poorly governed, over-permissioned, vulnerable to prompt injection, connected to sensitive data, or integrated into automated workflows without adequate controls, it can become a path for data exposure or unauthorized action.

For CISOs, this creates a dual mandate. The organization must defend against AI-enabled attackers, and it must secure its own AI adoption. Those are related but distinct problems. Defending against AI-enabled attackers requires faster detection, stronger identity controls, improved vulnerability prioritization, better monitoring of cloud and SaaS behavior, and response automation. Securing enterprise AI requires governance over what models are used, what data they can access, what actions they can take, what plugins or agents they can invoke, how prompts and outputs are logged, and how the organization prevents unauthorized AI use from becoming an uncontrolled data leakage channel.

## AI is Both Attack Surface and Defensive Accelerator

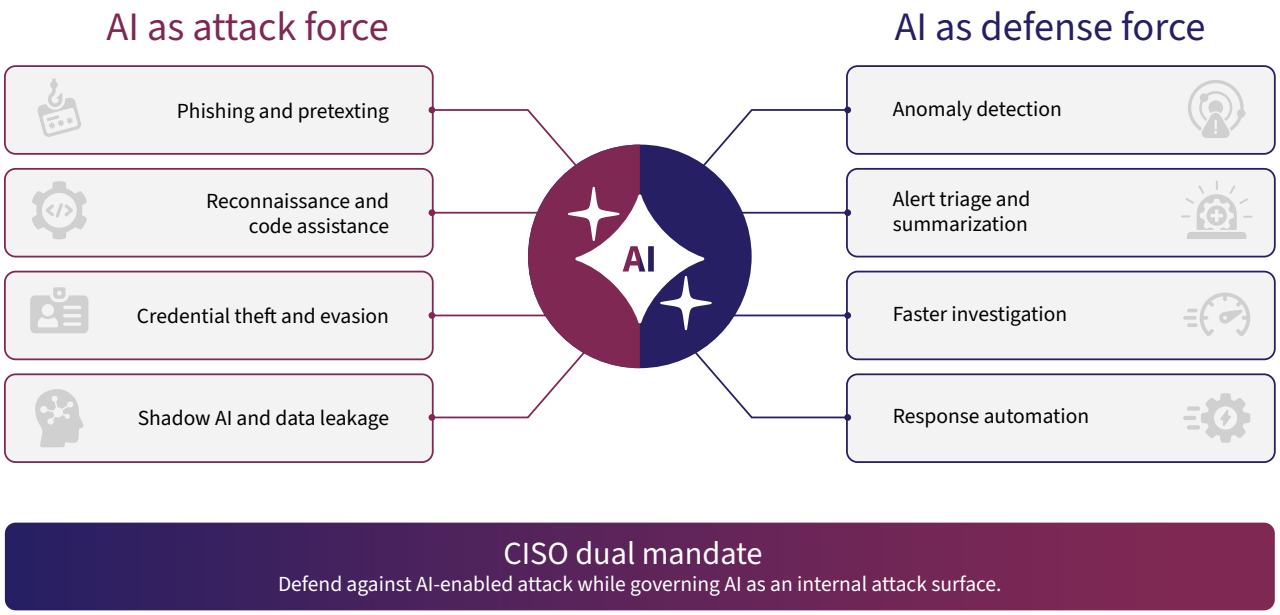


Figure 5: Organizations must defend against AI-enabled attackers while governing their own AI adoption.
*Concept based on World Economic Forum 2026, IBM breach research reporting, and public threat reporting.*

The economic challenge is that defenders must determine intent from behavior. A login, API call, file transfer, PowerShell command, OAuth grant, administrative change, or AI-agent action may be legitimate in one context and malicious in another. The difference is often visible only when signals are correlated across identity, device, network, cloud, SaaS, and application activity. That requirement exposes the weakness of security models built around isolated alerts. The next sections of this report examine adversary behavior and industry exposure in greater detail, but the economic logic begins here: As attackers become faster and more efficient, defenders cannot afford slow interpretation.

# Adversary Analysis

## Overview: Four Strategic Models for Cyber Power

Russia, China, Iran, and North Korea represent four different models of cyber power. They share certain tactical behaviors: phishing, credential theft, exploitation of internet-facing systems, abuse of trusted infrastructure, and efforts to hide inside legitimate network activity. But their strategic objectives differ. Russia uses cyber operations as part of a broader doctrine of disruption, influence, and coercion. China uses cyber operations to support long-term strategic advantage, technological sovereignty, intelligence collection, and pre-positioning. Iran uses cyber operations as a tool of retaliation, signaling, and asymmetric disruption. North Korea uses cyber operations as a state-directed revenue engine, blending espionage, cryptocurrency theft, fake remote-worker schemes, sanctions evasion, and access monetization to generate hard currency and support regime priorities. Its model collapses the line between nation-state activity, cybercrime, and insider risk: A North Korean operation may look like a crypto heist, a fraudulent job applicant, a compromised developer environment, or a financially motivated intrusion, but the strategic purpose is often state funding and sanctions circumvention.

### Four Models of Cyber Power

| | Russia | China | Iran | North Korea |
|---|---|---|---|---|
| **Primary objective** | Disruption and coercion | Persistence and strategic advantage | Retaliation and signaling | Financial extraction and sanctions evasion |
| **Typical style** | Gray zone pressure | Quiet access and pre-positioning | Disruptive operations | Theft, fraud, insider-style access |
| **Main enterprise risk** | Critical services and trust | IP, telecoms, infrastructure | OT and public-facing disruption | Crypto, remote work, development environments |
| **Defensive priority** | Resilience and crisis readiness | Long-term detection and infrastructure hardening | Rapid escalation readiness and OT controls | Identity governance and fraud controls |

Figure 6: Russia, China, Iran, and North Korea pursue different strategic outcomes through cyber operations.

For CISOs, the point of adversary analysis is not to memorize actor names. It is to understand intent. Intent determines target selection, dwell time, operational tempo, and likely business impact. A criminal ransomware group wants leverage and payment. A Russian-aligned actor may want to erode trust, amplify political pressure, or disrupt a supplier during a geopolitical crisis. A Chinese state-aligned actor may want persistent access to intellectual property, telecommunications infrastructure, or critical systems that could be used later. An Iranian actor may seek rapid disruption, public messaging value, or retaliation against sectors associated with U.S., Israeli, Gulf, or Western interests.

The practical implication is that every enterprise now operates inside a larger geopolitical system. A manufacturer may be targeted because it supports defense, logistics, semiconductors, energy, pharmaceuticals, or industrial supply chains. A healthcare organization may be targeted because operational disruption creates public pressure. A financial institution may be targeted because it represents economic leverage. A technology company may be targeted because it holds intellectual property, access to downstream customers, or data that supports national industrial strategy.

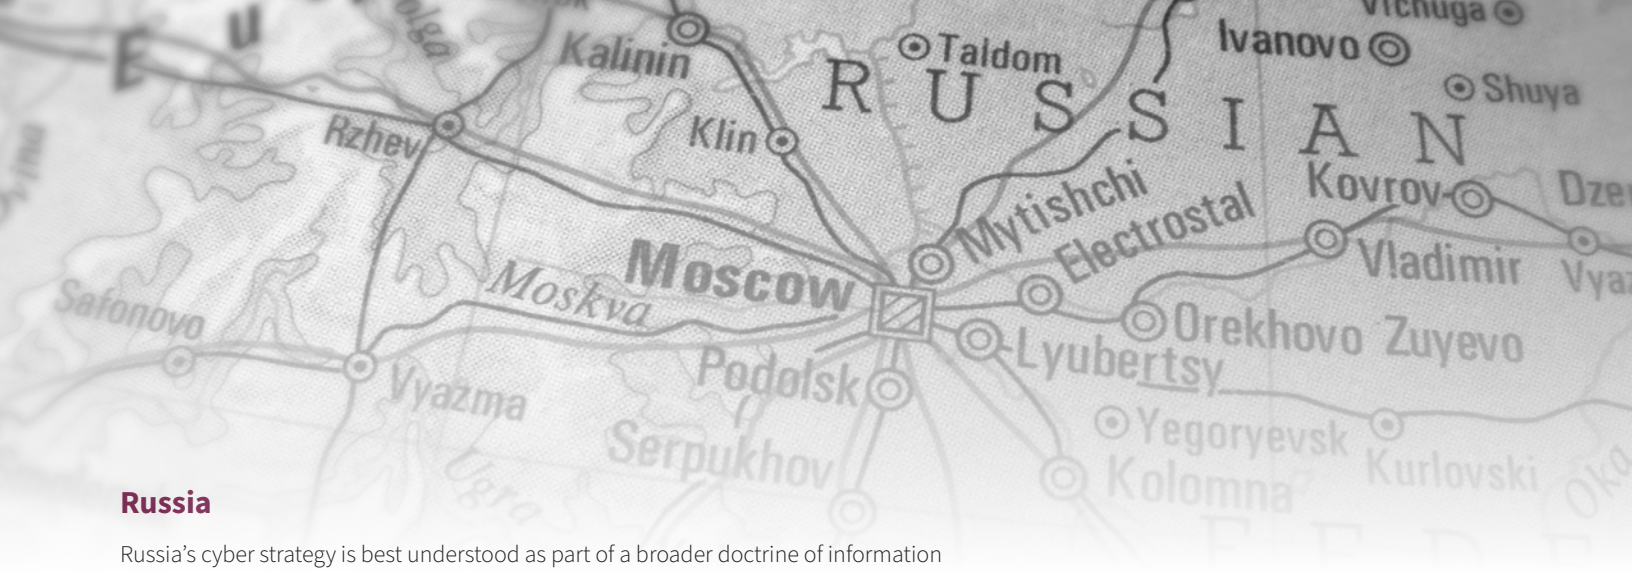


## Russia

Russia's cyber strategy is best understood as part of a broader doctrine of information confrontation. In this model, cyber operations, influence campaigns, espionage, sabotage, and psychological pressure are not separate activities. They are mutually reinforcing tools used to weaken adversaries, shape perception, and create strategic ambiguity. Russia does not need every cyber operation to produce visible destruction. Often, the goal is to impose cost, distract defenders, generate uncertainty, and erode confidence in institutions over time.

Russia's war against Ukraine has sharpened this model. Since 2022, Russian cyber operations have been closely intertwined with military activity, espionage, disinformation, and efforts to degrade the resilience of Ukraine and its supporters. Microsoft's reporting has described an increase in Russian cyber activity against NATO states, with government, research, think tank, and NGO targets prominent among observed victims. This aligns with Russia's strategic requirement: Gather intelligence, monitor policy formation, disrupt Western support for Ukraine, and influence public opinion across allied countries.

The 2026 Signal phishing campaign targeting European officials illustrates how Russia's approach continues to evolve. German authorities suspected Russia-linked actors in a campaign against high-ranking officials, military personnel, and journalists using the Signal messaging app. Reporting described approximately 300 targeted or compromised accounts and noted that victims were tricked by fake Signal support interactions or malicious QR-code workflows that granted attackers access to communications. The significance of this campaign is not merely that Signal was targeted. It is that attackers moved toward the trusted channels used by policymakers, military personnel, and journalists during periods of geopolitical tension.

Russian actors also remain well skilled at living-off-the-land techniques. Rather than relying exclusively on custom malware, they often use legitimate administrative tools, valid credentials, remote access utilities, scripts, and native operating system functions. These techniques reduce the likelihood of detection by tools that focus narrowly on known malware or file-based indicators. They also increase the importance of behavioral detection: unusual authentication patterns, abnormal administrative actions, unexpected data staging, suspicious east-west traffic, or command execution that does not match normal business activity.

The Russian threat model is not limited to espionage. It includes disruption and preparation for disruption. Russian-linked groups have repeatedly shown interest in critical infrastructure, energy systems, transportation, government services, logistics, and industrial environments. In a conflict scenario, these sectors offer coercive value because their disruption affects public confidence and economic continuity. The enterprise lesson is that OT and critical services cannot be treated as isolated engineering environments. They are part of the geopolitical attack surface.

For CISOs, the Russian model demands a resilience-first approach. Organizations should assume that Russian-aligned actors may seek access long before they intend to act. They may use trusted platforms, stolen credentials, unmanaged remote access pathways, or compromised vendors. They may combine cyber activity with information operations, leak threats, or public narratives designed to amplify the impact of the incident. The defensive priority is therefore not only preventing intrusion but also reducing dwell time, detecting abnormal internal behavior, protecting crisis communications, and ensuring that business continuity plans account for coordinated cyber and information pressure.

The executive takeaway is that Russia uses cyber operations to create strategic friction. The immediate technical objective may be credential theft, espionage, disruption, or access staging, but the broader purpose is often to weaken decision-making, impose cost, and exploit uncertainty. Defending against this model requires visibility across identity, network traffic, third-party access, collaboration platforms, and operational systems. It also requires coordination among cybersecurity, legal, communications, physical security, and executive leadership, because Russian-style gray zone operations rarely stay confined to the SOC.

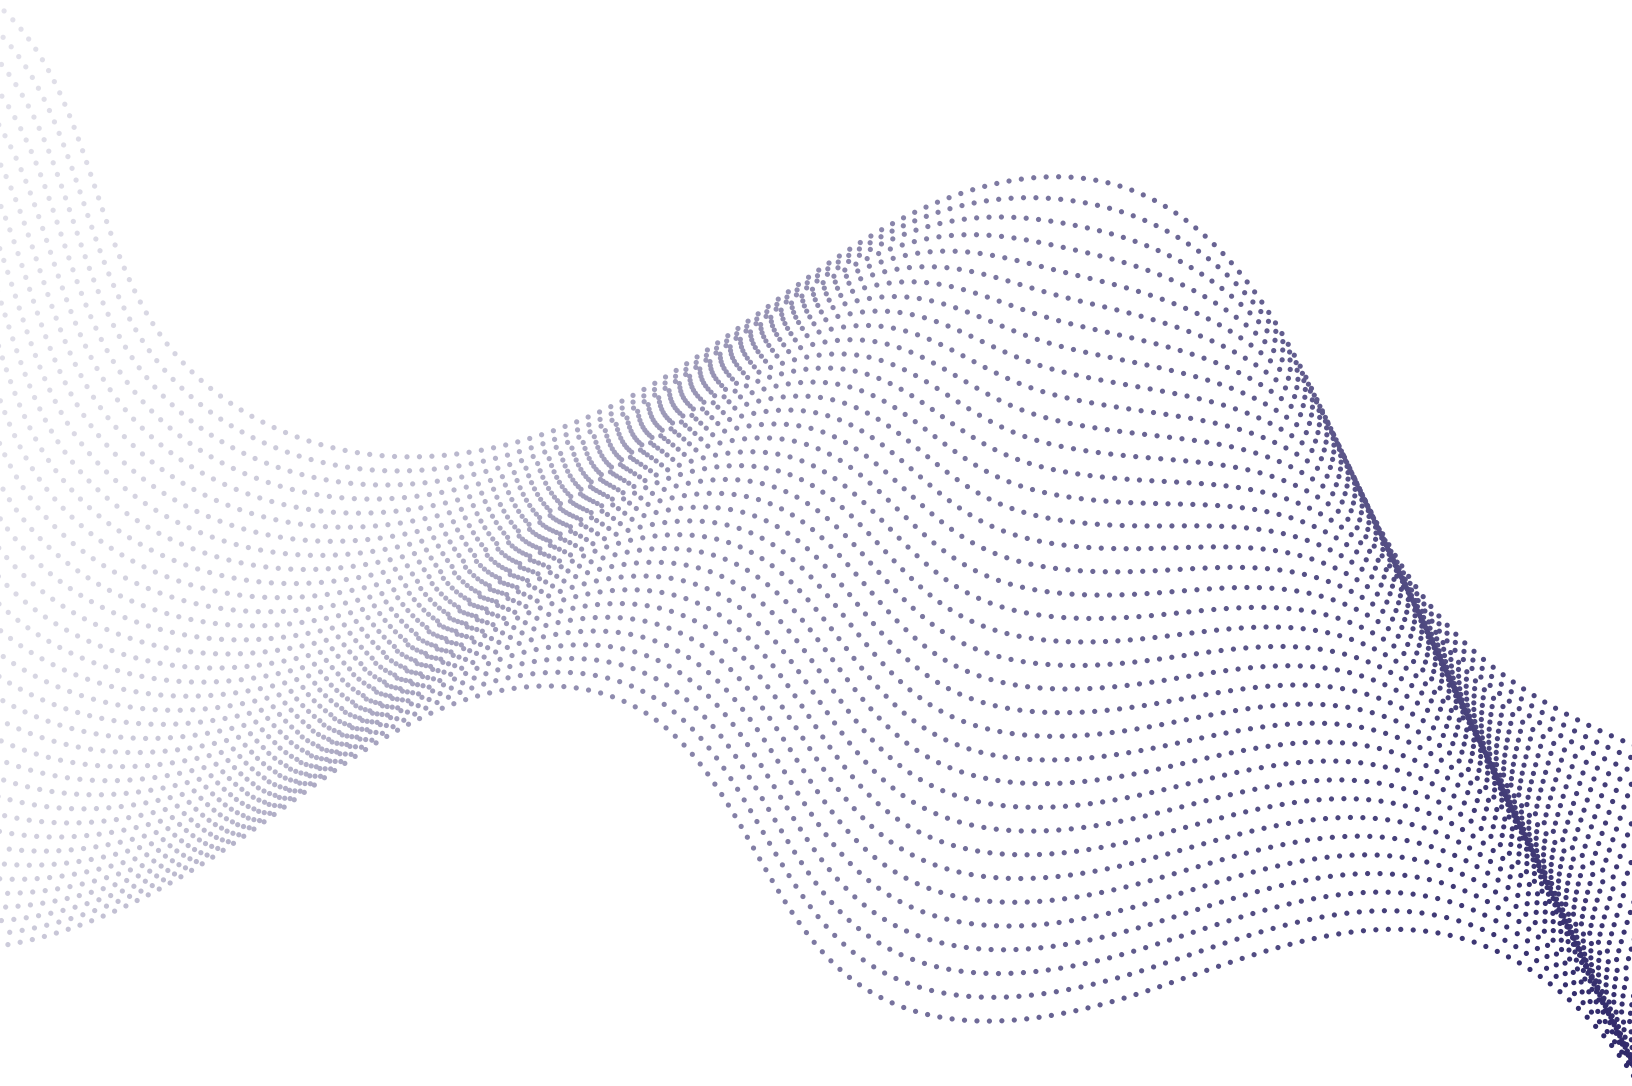


## China

China's cyber strategy is defined by patience, scale, and long-term advantage. Where Russia often emphasizes disruption and influence, China more often emphasizes access, intelligence collection, industrial strategy, and strategic positioning. The objective is not always an immediate impact. In many cases, the objective is to understand a target environment deeply, maintain access quietly, collect data over time, and preserve options for a future crisis.

U.S. and allied agencies have described Volt Typhoon as a PRC state-sponsored actor that targets critical infrastructure and uses living-off-the-land techniques to evade detection. The strategic concern is that such access may support pre-positioning rather than immediate espionage. In other words, the purpose of the intrusion may be to establish the ability to disrupt communications, logistics, water, energy, or transportation during a future geopolitical crisis.

Pre-positioning changes how defenders should interpret quiet activity. A low-and-slow intrusion that does not immediately exfiltrate large volumes of data may still be strategically significant. An attacker mapping network topology, harvesting credentials, testing remote access, or moving carefully through infrastructure may be preparing options rather than pursuing immediate monetization. Traditional incident triage may underestimate this kind of activity if it prioritizes only obvious data theft or malware execution.

### Pre-Positioning vs. Immediate Exploitation

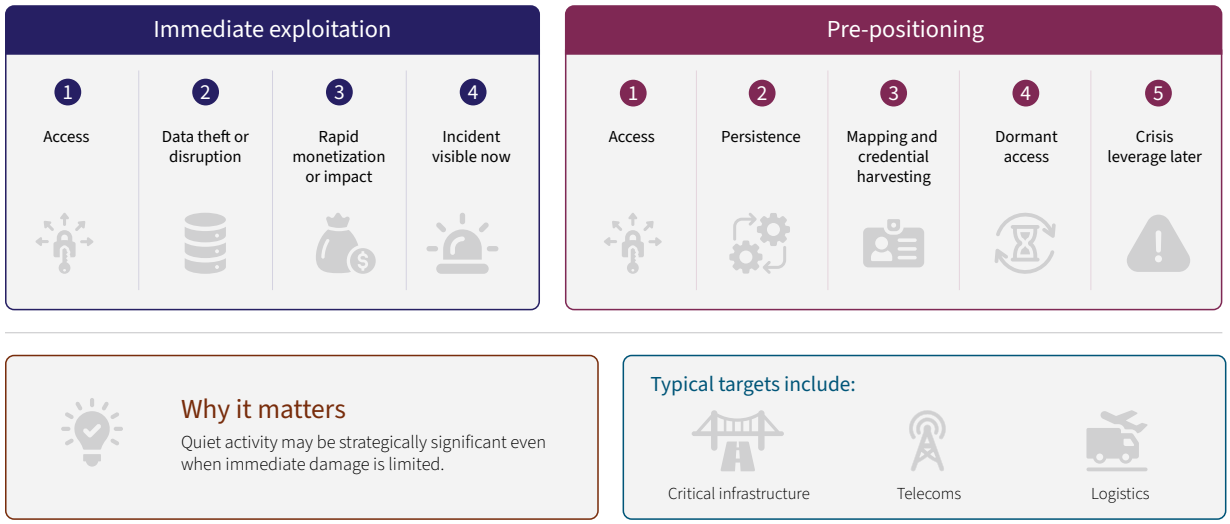


Figure 7: Some adversaries seek long-term crisis leverage rather than immediate disruption or monetization.
*Concept based on public reporting on Chinese state-linked operations.*

Salt Typhoon demonstrates another dimension of China's cyber strategy: access to communications infrastructure. U.S. agencies and public reporting have associated Salt Typhoon with a major telecommunications compromise that exposed sensitive communications-related data and raised significant counterintelligence concerns. Later reporting described the campaign as broader than initially understood, with effects across many countries and sectors. For CISOs, the implication is that telecommunications, network infrastructure, routers, lawful intercept systems, and communications metadata are high-value targets. These systems provide not only data but also strategic visibility into how governments, enterprises, and critical sectors communicate.

Chinese actors have also used compromised edge devices, routers, firewalls, and IoT infrastructure to obscure their operations. A 2026 joint warning reported by multiple outlets described China-nexus actors using routers and IoT infrastructure as covert botnets at scale, including activity associated with multiple "Typhoon" clusters. This reflects a practical reality: Defenders often under-monitor network appliances and edge infrastructure compared with endpoints and servers. Yet these systems are ideal for stealth, persistence, traffic relay, and access to sensitive network flows.

The defensive model must therefore emphasize long-term detection and control. Organizations should monitor for abnormal use of legitimate credentials, unexpected administrative behavior, unusual device-to-device communication, anomalous DNS or encrypted traffic, and persistence mechanisms in network infrastructure. They should harden edge devices, reduce exposed management interfaces, require strong authentication for administrative access, monitor service accounts, and maintain visibility into east-west traffic. The key is not merely to block known Chinese malware. The key is to detect behavior that does not belong.

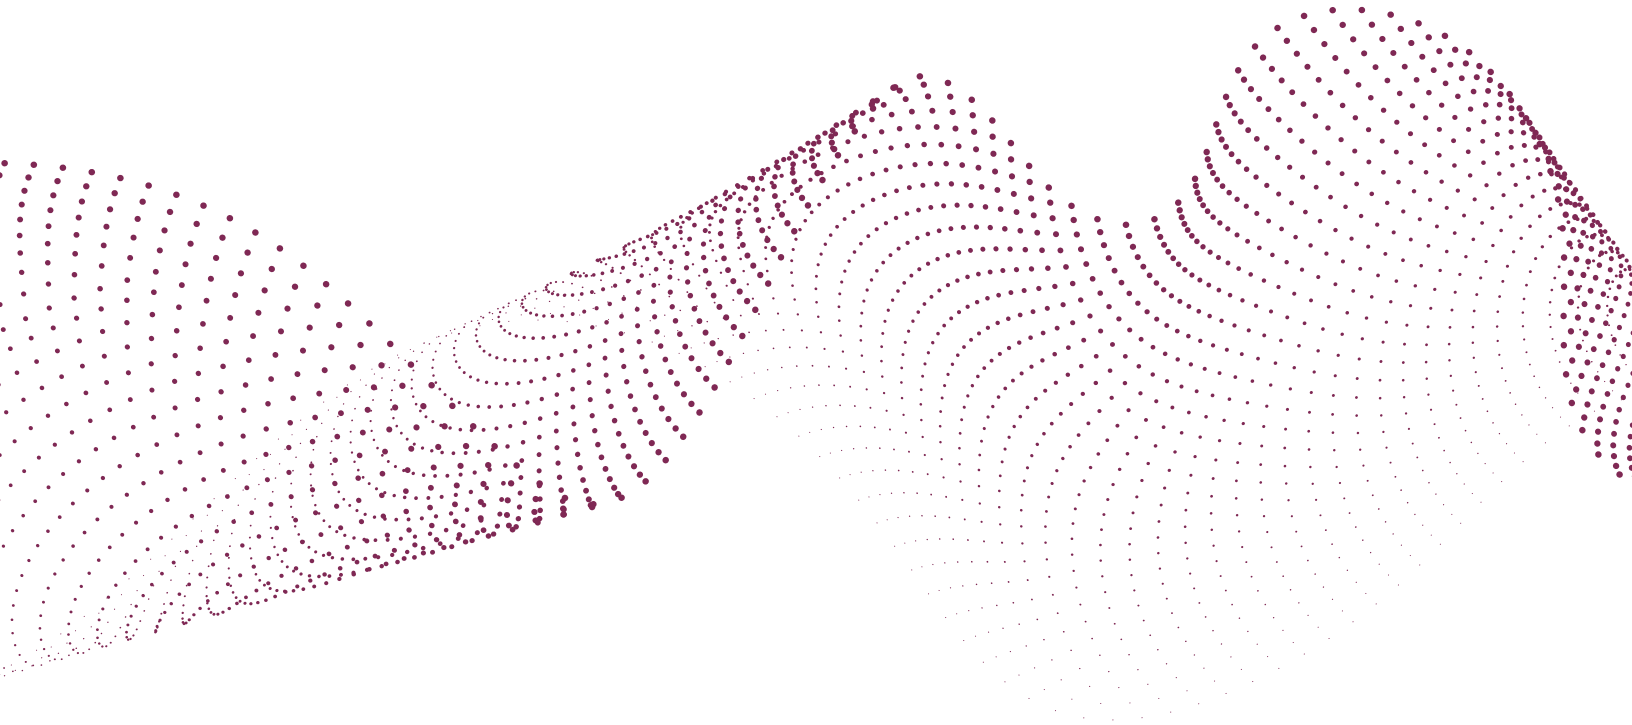


## Iran

Iran's cyber strategy differs from both Russia's and China's. Iran has historically used cyber operations as an asymmetric tool: a way to impose cost, signal resolve, retaliate against adversaries, and project power despite conventional constraints. Iranian operations are often more visible than Chinese espionage campaigns and more tightly connected to regional events than many criminal campaigns. The objective may be disruption, embarrassment, coercion, retaliation, or psychological impact.

In 2026, Iranian-linked activity against critical infrastructure has become a central concern. U.S. agencies issued warnings about Iran-affiliated cyber actors targeting internet-connected programmable logic controllers used in critical infrastructure, particularly in water and energy environments. Public reporting on the advisory emphasized the risk to Rockwell Automation and Allen-Bradley PLCs and urged organizations to shield these systems from direct internet access, review logs, disable unnecessary remote access, and use secure gateways and firewalls. Wired also described Iran-linked hackers targeting U.S. energy and water infrastructure and connected the activity to a broader escalation in destructive and disruptive intent.

The regional context is essential. Iranian cyber activity often increases during periods of military tension, sanctions pressure, diplomatic confrontation, or conflict involving Israel, the United States, Gulf states, or allied interests. Organizations with operations, customers, suppliers, or symbolic value connected to these geopolitical issues may face elevated risk even if they are not direct government targets. This is especially true for energy, defense, transportation, healthcare, finance, telecommunications, and critical infrastructure providers.

Iranian actors have also shown an interest in identity-based attacks. Password spraying, credential harvesting, and abuse of cloud accounts remain attractive because they are cheap, scalable, and effective. Once inside, attackers can search for sensitive data, identify operational dependencies, move laterally, or prepare for disruption. As with Russia and China, this makes identity telemetry and behavioral analytics essential. A valid login is not proof of legitimate activity.

For CISOs, the Iranian model requires readiness for rapid escalation. Unlike a long-term espionage campaign that may remain quiet for months, Iranian-linked operations may be triggered by external events and move quickly toward public impact. Security leaders should maintain geopolitical watchlists tied to their business exposure, review crisis response playbooks when regional tensions rise, and harden systems most likely to produce public disruption if compromised. This includes internet-facing OT, remote access tools, exposed VPNs, identity systems, cloud administration portals, and third-party service connections.

The executive takeaway is that Iran uses cyber operations to create pressure and signal capability. Its campaigns may be less patient than China's and less deeply integrated into global influence operations than Russia's, but they can be disruptive, unpredictable, and highly consequential. The organizations most exposed are those whose operations create public dependency or geopolitical symbolism. Defending against this model requires rapid detection, hardened remote access, strong identity controls, OT visibility, and an incident response posture that assumes cyber and physical consequences may converge.


## North Korea

North Korea represents a fourth model of cyber power: state-directed financial extraction fused with espionage, sanctions evasion, and strategic disruption. While Russia, China, and Iran use cyber operations primarily to coerce, pre-position, collect intelligence, or retaliate, North Korea uses cyber operations as a direct revenue engine for the state. Its cyber program is not only an intelligence capability. It is a mechanism for generating hard currency, evading sanctions, stealing intellectual property, and funding strategic priorities, including its weapons programs.

The most visible part of this model is cryptocurrency theft. North Korea-linked actors, often associated publicly with the Lazarus Group or related clusters, have repeatedly targeted exchanges, decentralized finance platforms, blockchain bridges, and virtual asset firms. The FBI attributed the 2025 Bybit theft, reportedly worth roughly $1.5 billion, to North Korean actors and described the activity as part of the “TraderTraitor” campaign. Reporting on Chainalysis research also found that North Korean hackers stole a record amount of cryptocurrency in 2025, with the Bybit incident representing the largest single event in that pattern. For CISOs, the lesson is that virtual assets, payment infrastructure, treasury systems, and financial workflows are not simply financial targets. They are geopolitical targets.

North Korea’s cyber operations also exploit the remote work economy. U.S. authorities have repeatedly warned that North Korean IT workers use false identities, stolen credentials, laptop farms, and third-party facilitators to obtain jobs at U.S. and global companies. These workers may generate revenue through salaries, but the risk extends beyond payroll fraud. Once inside an organization, they may gain access to source code, credentials, cloud systems, internal documentation, customer data, and production environments. In 2025, U.S. law enforcement announced actions involving North Korean remote IT worker schemes, including seizures of laptops, bank accounts, and infrastructure tied to workers who allegedly gained employment at dozens of companies. This is a different kind of cyber threat: The adversary does not always break in from the outside. Sometimes the adversary applies for a job.

That makes North Korea especially relevant to the agentic enterprise and the modern identity problem. The organization must govern not only human employees, but also contractors, remote workers, service accounts, machine identities, development environments, code repositories, and privileged automation. A fake IT worker with valid credentials can appear legitimate to many controls. A compromised developer workstation can provide access to code and secrets. A malicious insider can use normal workflows to exfiltrate data slowly. A state-sponsored operator can combine employment fraud, identity abuse, cloud access, and data theft in ways that traditional perimeter defenses are poorly designed to detect.

For CISOs, the defensive implications are clear. Organizations should strengthen identity verification during hiring and onboarding, especially for remote technical roles and contractors. They should validate work locations, review device provenance, enforce managed-device requirements, and watch for signs of laptop farms or proxy access. They should apply least privilege to developers and contractors, segment access to sensitive repositories and production environments, rotate secrets, and monitor unusual access patterns across code, cloud, and collaboration systems. Security teams should also coordinate closely with HR, legal, finance, procurement, and engineering because North Korean campaigns often cross organizational boundaries.

Financial controls also matter. Because North Korea's cyber strategy is revenue-driven, companies should treat payment flows, crypto wallets, treasury systems, payroll anomalies, contractor payments, and vendor banking changes as security-relevant signals. Business email compromise, fraudulent invoices, payroll diversion, and credential theft can all support the same strategic objective: converting enterprise access into funds. Security operations should therefore be integrated with fraud, finance, and insider-risk teams.

Network intelligence is particularly valuable against this threat model because North Korean operations often rely on legitimate access that must be interpreted behaviorally. A fake worker, compromised developer, or abused service account may not trigger a traditional malware alert. But the network can show unusual connections, unexpected data movement, abnormal repository access, suspicious remote administration, communication with unfamiliar infrastructure, or behavior inconsistent with a user's normal role. That behavioral evidence helps distinguish legitimate work from adversary activity.

The executive takeaway is that North Korea uses cyber operations as a strategic revenue system. Its campaigns are designed to steal, launder, infiltrate, and monetize. The most exposed organizations are those with digital assets, financial workflows, software development environments, remote technical workforces, privileged contractors, and cloud-native infrastructure. Defending against this model requires strong identity governance, hiring and contractor controls, developer environment monitoring, financial fraud integration, and network-derived behavioral context. In a world where adversaries can enter through a phishing lure, a crypto platform, a software dependency, or a job application, the CISO's mandate must extend beyond conventional intrusion defense to include the full lifecycle of trust.

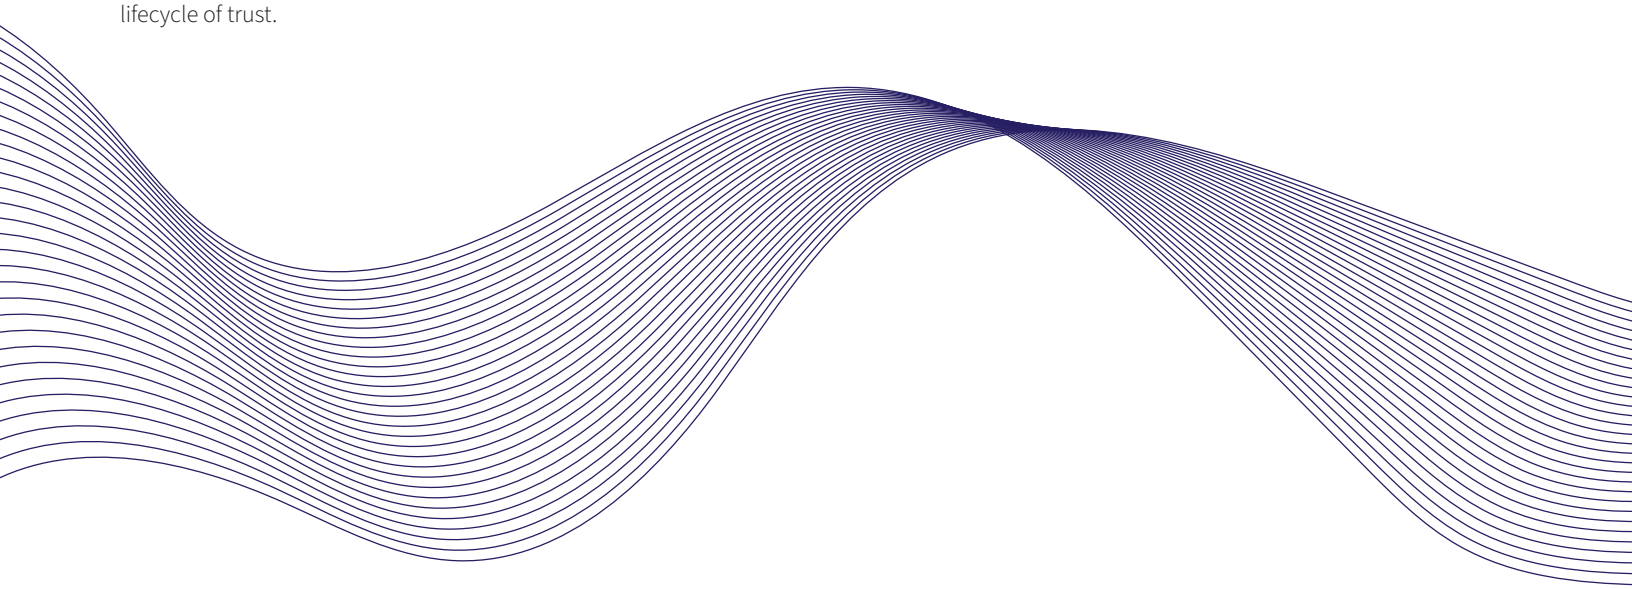

## Industry Exposure and the Visibility Crisis

The risk described in the previous sections does not fall evenly across the economy. Some industries face a higher probability of attack because they hold valuable data. Others face a higher impact because downtime creates immediate operational, safety, or public consequences. Still others are targeted because they sit inside supply chains, support critical infrastructure, or provide access to many downstream organizations. In 2026, the industries most exposed to geopolitical and cyber pressure are not simply those with the largest digital footprints. They are the industries where cyber incidents can become business disruptions, public crises, or strategic leverage.

### Industry Exposure Matrix

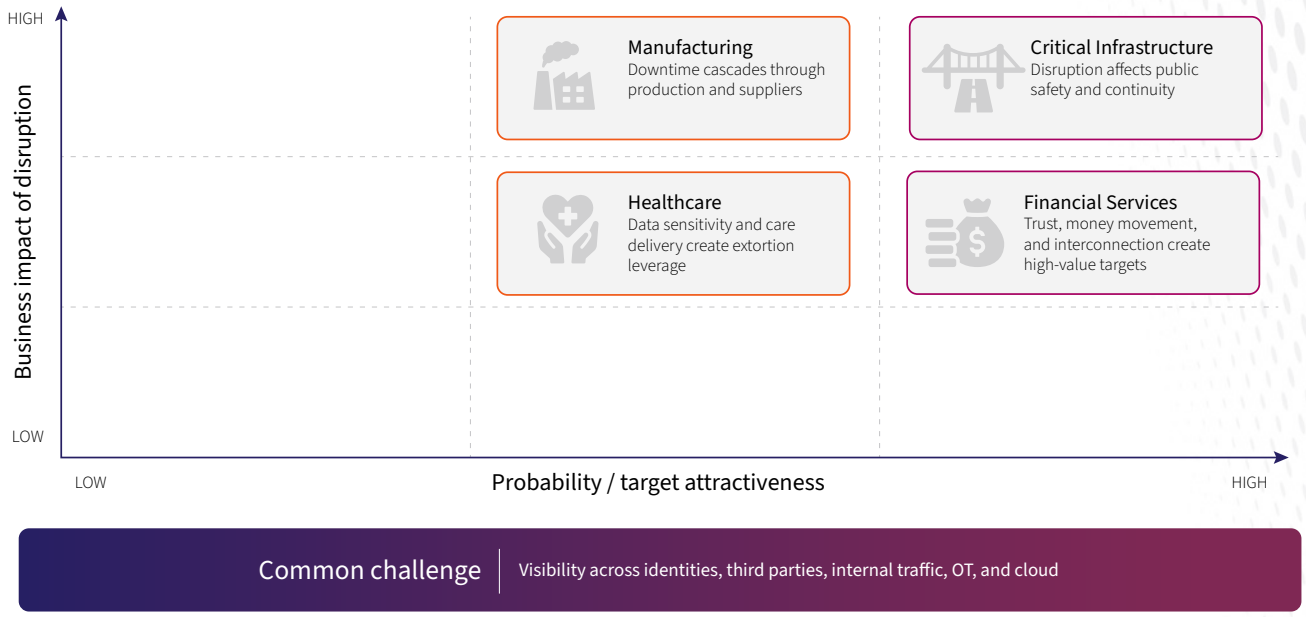


Figure 8: Risk depends on both target attractiveness and the business impact of disruption.

Manufacturing is also exposed because the boundary between IT and OT continues to erode. Production environments increasingly rely on remote monitoring, connected sensors, industrial IoT, predictive maintenance, cloud analytics, and vendor-managed systems. These capabilities improve performance, but they also create more pathways into operational environments. Many OT systems were designed for reliability and availability rather than modern adversarial conditions. They may run legacy software, lack endpoint agents, depend on flat networks, or rely on protocols that provide limited authentication and logging.

The visibility gap is especially acute when OT devices are exposed to the public internet. A 2025 academic analysis of publicly accessible OT found nearly 70,000 exposed OT devices globally, with concentrations in North America and Europe. The study also found that exposed systems often revealed detailed identifying information, including vendors, protocols, and outdated firmware versions with known vulnerabilities. For attackers, this kind of exposure reduces reconnaissance cost and increases the probability of finding a viable target.

## OT Exposure: What Attackers Can See

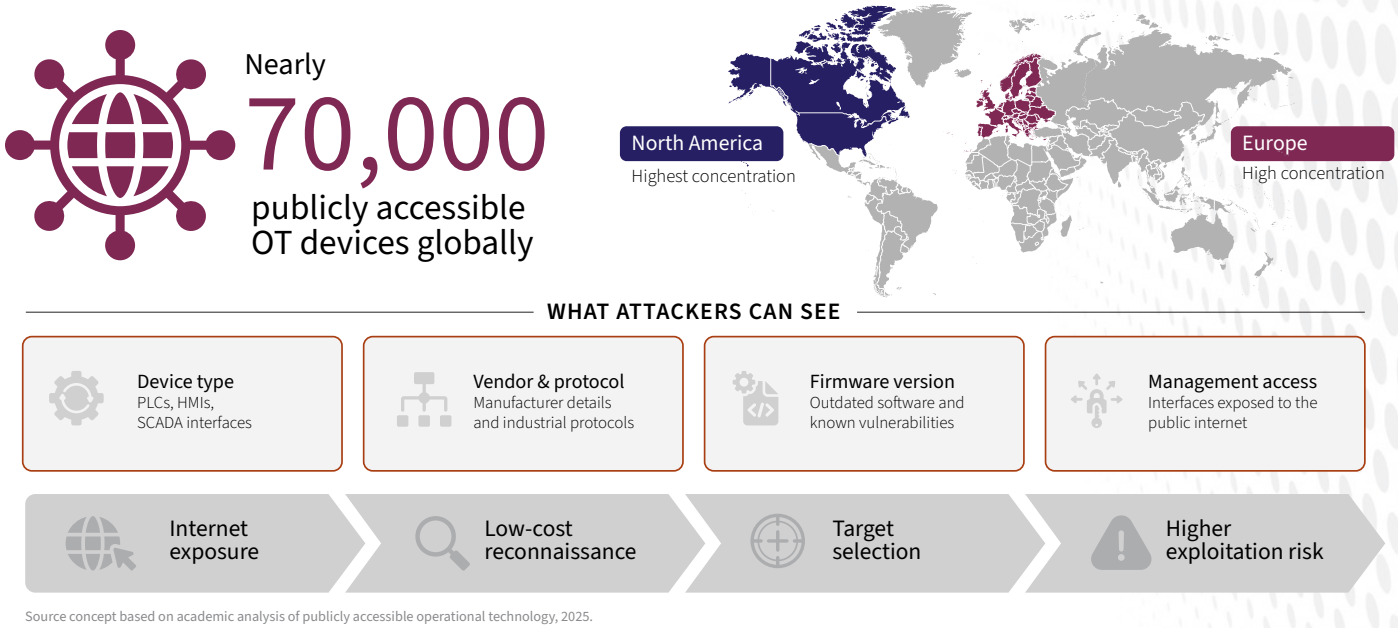


Source concept based on academic analysis of publicly accessible operational technology, 2025.

Figure 9: Publicly exposed OT devices reduce attacker reconnaissance cost and increase the likelihood of exploitation.
*Source concept based on academic analysis of publicly accessible operational technology, 2025.*

For manufacturing CISOs, the implication is that asset inventory is not enough. The organization needs to understand behavior. Which systems normally communicate with PLCs? Which engineering workstations issue management commands? Which vendors connect remotely, and when? Which protocols move between IT and OT? Which devices are unmanaged, unsupported, or using outdated firmware? Without that behavioral context, security teams may not detect the early stages of an intrusion until the business impact is already visible on the factory floor.

The healthcare visibility challenge is complicated by fragmentation. A hospital network may include modern cloud applications, legacy clinical systems, unmanaged medical devices, third-party portals, remote clinicians, contractors, and affiliated practices. Many of these systems are difficult to patch, difficult to monitor, or difficult to take offline for maintenance. Security teams often cannot deploy traditional endpoint tools to every device that matters. In that environment, attackers can exploit identity systems, email, remote access, vendor connections, or unmanaged devices to move through workflows that appear legitimate until viewed in context.

This is why exposed OT and industrial control systems matter so much. In water, energy, and transportation environments, attackers may not need sophisticated malware to create serious risk. Weak remote access, exposed controllers, default credentials, unpatched systems, poor segmentation, or insufficient monitoring may be enough. A review of malicious cybersecurity incidents in the water sector found recurring issues across incidents, including a need for more adaptive, cooperative, and comprehensive cyber defense as threats grew in frequency, diversity, and complexity.

The critical infrastructure visibility crisis is partly technical and partly organizational. These environments often combine corporate IT, OT, field equipment, vendors, regulators, regional operators, and legacy systems. Responsibility may be divided across engineering, operations, safety, IT, cybersecurity, and external service providers. This fragmentation can create blind spots. Security teams may see corporate endpoints but not industrial protocols. Operations teams may see process alarms but not credential abuse. Vendors may manage equipment without providing sufficient telemetry. Executives may see compliance status but not real-time exposure.

The financial sector's visibility challenge is not usually a lack of telemetry. Large financial institutions often have extensive security tools. The problem is correlation and prioritization. Fraud signals, authentication anomalies, endpoint alerts, network flows, SaaS activity, cloud logs, API calls, and transaction data may exist in separate systems with separate owners. Attackers exploit that fragmentation. A credential compromise may first appear as a normal login. A fraudulent transaction may look like customer behavior. Data staging may resemble normal reporting. An API abuse pattern may not be obvious unless correlated with identity, device, and network context.

Financial services are also exposed to third-party and software supply chain risk. The sector depends on payment networks, data aggregators, analytics providers, cloud platforms, core processors, managed service providers, compliance platforms, and customer-facing fintech integrations. Research on software supply chain security notes that supply chains are only as strong as their weakest link and that attacks against vulnerable links can disrupt organizations ranging from major enterprises and government agencies to smaller open-source maintainers. For financial institutions, that means the attack surface includes not only internal systems but also the software and services that connect the institution to the broader economy.

The visibility crisis has several dimensions. The first is asset visibility. Organizations cannot protect systems they do not know exist, and many environments contain unmanaged devices, shadow SaaS applications, forgotten cloud workloads, legacy systems, or vendor-managed assets outside normal security coverage. The second is behavioral visibility. Knowing that an asset exists is not the same as knowing what it normally does. Defenders need baselines for communication patterns, authentication behavior, administrative activity, data movement, and protocol usage.

The third dimension is identity visibility. As attackers rely more heavily on valid credentials, security teams need to understand not just whether a login succeeded, but whether the account behavior makes sense. Is the user accessing unusual systems? Is a service account behaving interactively? Is an administrator using an unexpected path? Is a vendor account active outside normal windows? Without identity context, credential abuse can hide inside ordinary operations.

The fourth dimension is encrypted and east-west traffic visibility. Modern attacks often move inside the network rather than directly across the perimeter. They may use encrypted channels, legitimate remote access tools, cloud APIs, or administrative protocols. If defenders cannot observe internal traffic patterns, lateral movement, data staging, and command-and-control activity may remain hidden until impact occurs. This is particularly important in environments where endpoint coverage is incomplete, such as OT, IoT, medical devices, and network infrastructure.

## NDR and the Agentic Enterprise

The next phase of cybersecurity will not be defined only by better detection. It will be defined by whether organizations can give human teams and AI agents the context they need to reason accurately, act confidently, and operate at machine speed. As attacks accelerate, environments become more distributed, and AI agents begin creating new workflows across the enterprise, security and operations teams need a source of truth that reflects what is actually happening at run time. Network detection and response (NDR) has become central to that requirement.

NDR provides more than visibility into packets or alerts. A modern NDR provides decision-grade intelligence: high-fidelity, network-derived context that shows how assets, identities, applications, services, and data flows behave in real time. That context matters because modern adversaries increasingly avoid obvious malware and operate inside legitimate activity. They use valid credentials, encrypted channels, cloud services, administrative tools, remote access pathways, and third-party connections. They move through the same systems enterprises depend on every day. In that environment, the network becomes the enterprise's most reliable behavioral evidence layer.

This is the strategic role of ExtraHop RevealX™. RevealX turns the network into a source of truth for the modern enterprise by combining network detection and response, network performance management, intrusion detection, and packet forensics into a unified platform. The value is not simply that more telemetry is collected. The value is that telemetry is pre-correlated across asset, identity, behavioral, and threat context, giving both human analysts and AI-driven workflows a more complete understanding of what is happening. In a world where defenders cannot manually assemble every signal in time, pre-correlated context becomes a strategic advantage.

This shift is especially important because the enterprise itself is becoming more agentic. AI agents are being embedded into business workflows, development pipelines, service operations, customer support, security processes, and infrastructure management. These agents can accelerate productivity, but they also create new forms of risk. They may access sensitive systems, invoke APIs, interact with data, automate tasks, or create workflows that are invisible to traditional monitoring. The same problem already exists with machine identities, service accounts, tokens, and automation scripts. As non-human activity grows, organizations need a way to understand not just who is acting, but what is acting, where it is connecting, what data it is touching, and whether its behavior is expected. NDR is uniquely suited to this problem because it observes behavior at run time. Endpoint tools can miss unmanaged assets. Logs may be incomplete or delayed. Cloud and SaaS tools may show only one slice of the environment.

An agentic SOC is a security operations model in which AI agents help analysts triage, investigate, prioritize, and respond to threats. In this model, AI agents should not be fed raw noise. They need enriched, behaviorally grounded evidence that allows them to reconstruct attack chains, understand blast radius, and recommend response actions with confidence. NDR provides that evidence. By correlating network behavior with identity, asset, protocol, and threat context, NDR can help an agent determine whether a credential is being abused, whether an unmanaged device is participating in an intrusion, whether encrypted traffic hides suspicious activity, or whether lateral movement is underway.

The network reveals relationships: which systems communicate, which identities are associated with those communications, which protocols are being used, which transactions are encrypted, which assets are unmanaged, and which activity deviates from normal patterns. This makes NDR a foundational capability for both the agentic SOC and the agentic NOC.

This does not eliminate the human analyst. It elevates the analyst. The agentic SOC should remove the manual drag of repetitive investigation: stitching together logs, searching for related events, collecting asset details, summarizing incident timelines, and identifying likely next steps. Human analysts remain responsible for judgment, escalation, policy, and high-impact response decisions. NDR strengthens that model because it gives both humans and agents the same grounded evidence base. The result is faster investigation, clearer prioritization, and a stronger ability to respond before an intrusion becomes a breach or outage.

The same logic applies to the agentic NOC. Network and IT operations teams face environments that span data centers, cloud platforms, SASE architectures, Kubernetes, SaaS applications, remote users, and encrypted protocols. Performance problems can hide inside dependencies that are difficult to see with traditional monitoring. A user-facing slowdown may originate in the network, the application, the cloud provider, identity infrastructure, DNS, TLS negotiation, an overloaded service, or a misbehaving AI workflow. An agentic NOC uses AI-assisted workflows to detect degradation, explain likely root cause, and recommend remediation before users or business processes are affected.

NDR provides the run-time intelligence the agentic NOC needs. By continuously analyzing transactions, dependencies, protocols, and communication patterns, NDR helps operations teams understand application health across hybrid and cloud-native environments. It can help distinguish performance degradation from security activity, operational misconfiguration from malicious behavior, and normal service dependency from unexpected data movement. This matters because security and performance incidents increasingly overlap. Ransomware, DDoS, credential abuse, AI agent misbehavior, cloud misconfiguration, and third-party outages can all create symptoms that cross SOC and NOC boundaries.

The strategic opportunity is to unify those teams around one source of truth. Historically, SOC and NOC teams often worked from different tools, different telemetry, and different assumptions. That created duplicated spend, slow handoffs, and finger-pointing when incidents crossed domains. The agentic enterprise requires a different model: one enriched context layer that supports both security and operations. With a shared evidence base, the SOC can investigate threats faster, the NOC can resolve performance problems faster, and both teams can understand incidents that span security, availability, identity, cloud, and network behavior.

For ExtraHop, this is the core message: RevealX provides decision-grade intelligence for the agentic enterprise. It gives organizations one all-in-one sensor, one context layer, and one investment that can support security detection, performance assurance, packet forensics, AI governance, and operational resilience. This message is especially powerful because it connects NDR to the broader business problem CISOs and CIOs face: not just finding attacks, but governing complexity, accelerating response, reducing tool sprawl, and enabling trusted automation.

In summary, the agentic enterprise will require more than automation. It will require trusted automation grounded in accurate, real-time, behavioral evidence. ExtraHop's role is to provide that evidence layer: the network-derived context that allows AI agents and human teams to reason accurately, act confidently, and operate at machine speed. In an era of permanent gray zone conflict, NDR becomes the foundation for resilient, agentic security and operations.

# The Future of Defense: Network, Context, and Machine-Speed Security

The future of cybersecurity will be defined by whether defenders can close the gap between the speed of attack and the speed of understanding. The previous sections described a threat environment shaped by persistent gray zone conflict, AI-enabled adversaries, industry-specific exposure, and visibility gaps across complex infrastructure. This final section turns to the defensive model required for that environment.

## NDR and the Agentic Enterprise

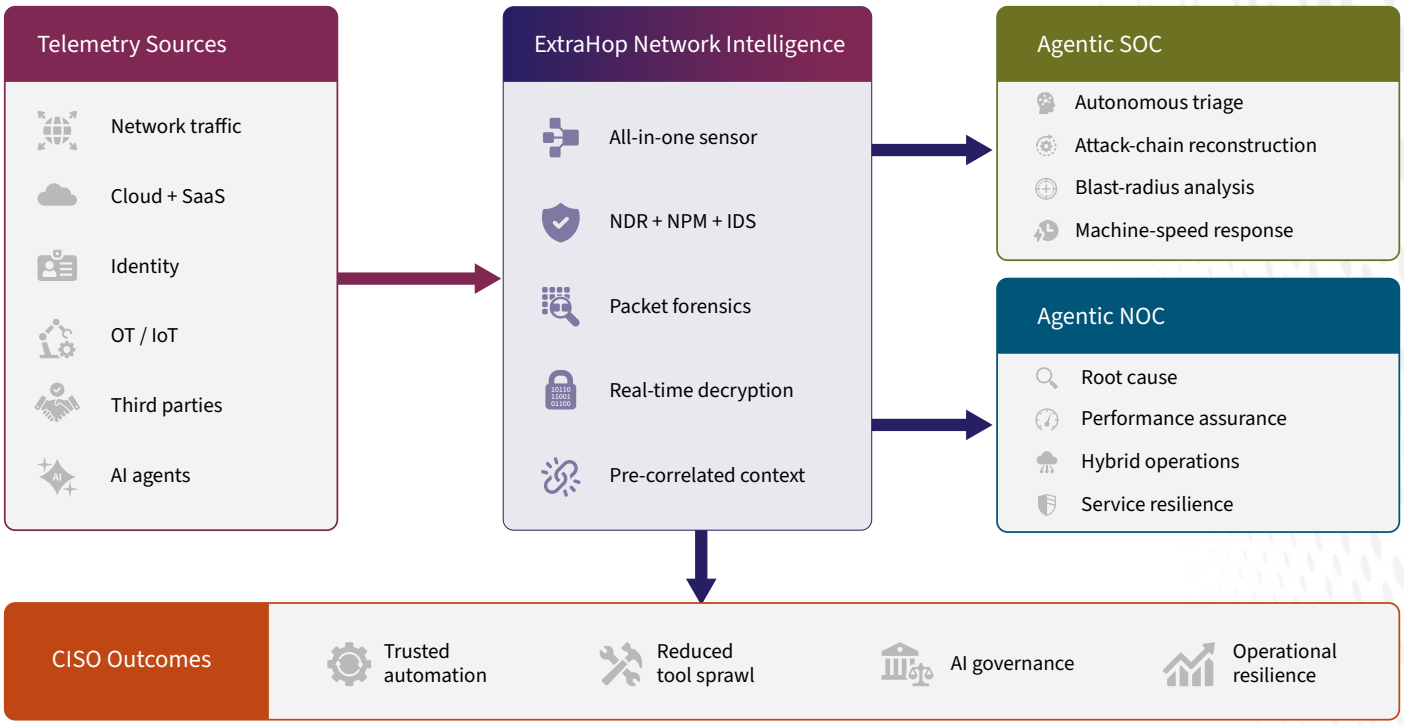


The foundation of this model is visibility. In earlier security architectures, visibility often meant asset inventory, log collection, and perimeter monitoring. Those capabilities still matter, but they are no longer sufficient. Modern environments are distributed, encrypted, hybrid, and dynamic. Workloads appear and disappear. SaaS applications host critical data outside the traditional perimeter. Employees work from unmanaged networks. Vendors connect remotely. OT and IoT devices often cannot support agents. Attackers use valid credentials and legitimate tools. In this environment, visibility must mean the ability to observe behavior across the full digital estate, including the internal traffic that shows how systems actually interact.

The second requirement is context. Visibility without context creates noise. A login event, DNS request, PowerShell command, API call, or file transfer may be legitimate or malicious depending on who performed it, from where, against which system, at what time, and in relation to which other behaviors. Context turns telemetry into judgment. It connects identity to device, device to application, application to data, data to business process, and behavior to risk.

The third requirement is speed. Attack timelines have compressed to the point where manual processes alone cannot keep pace. CrowdStrike's 2026 reporting, as summarized in industry coverage, found that average breakout time fell to 29 minutes and that the fastest observed breakout occurred in 27 seconds. The same reporting described AI-enabled adversary operations as increasing sharply and warned that intrusions increasingly move through trusted identities, SaaS applications, and cloud infrastructure while blending into normal activity. The practical conclusion is unavoidable: If attackers can move laterally in minutes, defenders cannot rely on hours of manual alert triage before action begins.

The future defensive model also requires stronger identity control. As attackers increasingly use valid credentials, identity becomes both a control plane and a detection surface. Every identity should be evaluated in context: user identity, device posture, location, behavior, privilege level, access history, and the sensitivity of the requested resource. Service accounts, machine identities, API tokens, OAuth grants, and AI agents deserve the same scrutiny as human users. In many environments, non-human identities are more privileged, less monitored, and harder to rotate than employee accounts. That makes them attractive targets.

This shift also changes executive reporting. Boards do not need a long list of blocked attacks. They need to understand exposure, resilience, and readiness. Useful metrics include mean time to detect, mean time to contain, coverage of critical assets, percentage of unmanaged devices discovered, third-party privileged access, identity risk, patch exposure for exploited vulnerabilities, backup recoverability, incident response exercise outcomes, and business impact scenarios. Security leaders should report cyber risk in terms of operational continuity, customer trust, regulatory exposure, and strategic dependency.

However, the future of defense is not only a technology investment. It is an organizational change. Security, IT, cloud teams, network teams, engineering, OT, legal, privacy, communications, procurement, and executive leadership must operate from a shared view of risk. Incident response plans must include business owners. Crisis communications must be prepared before a breach. Legal and regulatory obligations must be understood before data is stolen. Vendor access must be governed before it is abused. Backup and recovery must be tested before ransomware hits. Machine-speed defense requires human alignment long before automation is triggered.

The final strategic point is resilience. Prevention remains important, but no credible strategy should assume that prevention will always succeed. The organizations best prepared for the coming environment will be those that can absorb attack pressure, detect adversary activity early, contain movement quickly, maintain essential operations, and recover with confidence. Resilience does not mean accepting compromise. It means designing the organization so that compromise does not automatically become catastrophe.

In summary, the future of defense is built on network visibility, contextual understanding, and machine-speed security. Visibility allows organizations to see what is actually happening across hybrid, cloud, SaaS, OT, IoT, and third-party environments. Context allows them to interpret behavior in relation to identity, asset value, business process, and adversary tactics. Speed allows them to respond before attackers can turn access into impact. AI and automation will be essential, but only when governed, explainable, and integrated into a broader operating model. The goal is not simply to stop more alerts. The goal is to preserve trust, continuity, and control in an era of permanent gray zone conflict.

## Resources


World Economic Forum, Global Cybersecurity Outlook 2026, published January 2026.

Verizon, 2026 Data Breach Investigations Report, cited for incident volume, vulnerability exploitation, credential abuse, ransomware, third-party involvement, generative AI use, and human element findings.

CrowdStrike, 2026 Global Threat Report, cited for breakout time and AI-accelerated adversary activity.

ITPro, reporting on Cisco Splunk's Hidden Costs of Downtime findings, 2026.

TechRadar, reporting on CrowdStrike's 2026 Global Threat Report and average breakout time, AI-enabled adversary operations, cloud activity, and vulnerability exploitation.

Anthropic, "Disrupting the first reported AI-orchestrated cyber espionage campaign," published November 2025.

Microsoft, Microsoft Digital Defense Report 2025, cited for Russian cyber activity against NATO states and prominent targeting of government, research, think tank, and NGO organizations.

AP News, "Germany suspects Russia is behind Signal phishing that targeted top officials," April 2026.

DigiChina, "Forum: Technology in China's 15th Five-Year Plan," cited for 2026–2030 planning context and China's technology priorities.

CISA, "PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure," advisory AA24-038A, cited for Volt Typhoon targeting of U.S. critical infrastructure and living-off-the-land behavior.

Politico, "China-linked hackers stole wiretap data from telcos, FBI and CISA say," November 2024.

Washington Post, "FBI warns Chinese hacking campaign has expanded, reaching 80 countries," August 2025.

TechRadar, "China-nexus cyber actors are turning routers and IoT infrastructure into covert botnets at scale," 2026.

CISA, "Iranian Affiliated Cyber Actors Exploit PLCs Across U.S. Critical Infrastructure," advisory AA26-097A, cited for Iranian-affiliated targeting of internet-connected programmable logic controllers.

Tom's Hardware, reporting on CISA guidance for Rockwell Automation and Allen-Bradley PLC exposure, April 2026.

Wired, "Iran-Linked Hackers Are Sabotaging US Energy and Water Infrastructure," April 2026.

The Guardian, "North Korea behind $1.5bn hack of crypto exchange Bybit, says FBI," February 2025.

AP News, "US sanctions North Korean bankers accused of laundering stolen cryptocurrency," November 2025.

Tom's Hardware, reporting on Chainalysis research that North Korean hackers stole a record amount of cryptocurrency in 2025.

## Resources, continued


Politico, “Hundreds of laptops, bank accounts linked to North Korean fake IT workers scheme seized in major crackdown,” June 2025.

ITPro, reporting on NCC Group research on ransomware targeting industrial organizations, 2026.

Matthew Rodda and Vasilios Mavroudis, “Analysis of Publicly Accessible Operational Technology and Associated Risks,” 2025.

ITPro, reporting on healthcare email security failures and Microsoft 365-related exposure, 2025.

Amin Hassanzadeh et al., “A Review of Cybersecurity Incidents in the Water Sector,” 2020.

Elizabeth Lin et al., “S3C2 Summit 2025-03: Industry Secure Supply Chain Summit,” 2025.

NIST, Cybersecurity Framework 2.0 resource center, cited for the role of cybersecurity as an organizational risk management discipline and for CSF 2.0’s emphasis on governance, profiles, implementation examples, and informative references.

ITPro, reporting on IBM’s 2025 breach research, cited for shadow AI risk, AI-related breach exposure, AI-generated phishing, and deepfake impersonation as attacker techniques.

Axios, reporting on IBM breach-cost research, cited for average breach costs, hybrid complexity, and the role of AI and automation in reducing breach impact.

ITPro, reporting on ExtraHop’s NDR platform and hybrid infrastructure visibility, cited for industry coverage of network detection and response, contextual visibility, and reduced operational complexity.


## ABOUT EXTRAHOP

ExtraHop turns the network—the enterprise’s ultimate source of truth—into actionable insight to power security, performance, and resilience. Delivering superior data by design, we ensure superior defense by default.

The ExtraHop modern network detection and response (NDR) platform provides visibility that thinks, analyzing behavior to intercept evasive risks before they cause damage. We transform network noise into definitive context, enabling security teams to make faster decisions and operate at uncompromised scale.

Whether securing cloud modernization or de-risking AI adoption, ExtraHop gives global enterprises the ground truth they need to thrive.

To learn more, visit extrahop.com or follow us on LinkedIn.

EXTRAHOP®


**info@extrahop.com**
**extrahop.com**