



ARTIFICIAL INTELLIGENCE AND CYBERSECURITY RESILIENCE IN THE UNITED STATES

Machine-Speed Threats, Uneven Defenses, and a 24-Month Path
to Continuity Under Compromise

PUBLICATION TYPE:	Research Paper
AUTHOR:	Bernaldo Domínguez
ROLE:	Founder, MOSAIC Strata / Research Author
PUBLICATION STATUS:	Final Release
VERSION:	1.0
RELEASE DATE:	09/15/2026
EVIDENCE CUTOFF:	09/15/2026

CONTENTS

MOSAIC Strata Research Publication

Abstract	4
Introduction	5
Research Approach and Evidentiary Limits	7
Part I - The Threat Environment Is Moving Faster	8
AI capability is rising, but benchmarks are not reality	10
Attack economics may matter more than autonomous capability	12
Part II - One Threat Environment, Unequal Defenders	13
Enterprise scale creates capability and complexity	14
Healthcare shows what happens when cybersecurity becomes mission failure	14
Healthcare scenario methodology for 2027	16
Critical infrastructure reveals the limits of organization-by-organization defense	17
Water demonstrates both exposure and a measurement problem	17
Part III - The Attack Surface Has Moved	19
The software supply chain is becoming an identity chain	20
AI agents are privileged identities before they are intelligent coworkers	20
Post-quantum migration illustrates why cybersecurity must prepare before crisis	21
Part IV - From Compliance to Resilience	22
Cybersecurity rules need technical renewal cycles	23
Minimum compliance should not be mistaken for security	24
Continuity Under Compromise	24
Author's analytical position	25
Part V - The Next Twenty-Four Months	26
The first 90 days: visibility before technology	27
Months three through six: remove the attacker's easiest leverage	27
Months six through twelve: close the capability gap	27



MOSAIC
STRATA

INTEGRITY • EVIDENCE • TRUST


Months twelve through eighteen: containment must begin approaching attacker speed 28

Months eighteen through twenty-four: stop assuming and start proving 28

Outlook: What Happens If We Do Nothing? 29

Conclusion 30

References 31



Abstract

Artificial intelligence (AI) is changing cybersecurity in the United States, but its most immediate effect is not the replacement of human attackers with autonomous machines. AI is changing the economics, speed, scalability, adaptability, and labor requirements of cyber operations. Reconnaissance, vulnerability research, exploit development, social engineering, credential harvesting, malware modification, stolen-data analysis, and portions of attack orchestration can increasingly be accelerated or delegated to AI systems. At the same time, AI agents are becoming embedded in business, healthcare, government, education, software development, finance, and critical infrastructure, creating a new class of privileged machine identities that must also be secured.

This transition is occurring while conventional cyber risk remains severe. Across adjacent Verizon Data Breach Investigations Report (DBIR) cycles, vulnerability exploitation increased from 20% to 31% of analyzed breaches, ransomware involvement rose from 44% to 48%, and third-party involvement increased from approximately 30% to 48% (Verizon, 2025, 2026). CrowdStrike reported average eCrime breakout time falling from 48 minutes in its 2025 report to 29 minutes in its 2026 report, while malware-free detections rose from 79% to 82% (CrowdStrike, 2025, 2026). Mandiant simultaneously found that internal detection improved from 43% to 52%, demonstrating that defenders are improving even as adversaries become faster and more evasive (Google Cloud, 2026a). FBI Internet Crime Complaint Center (IC3) data show complaints increasing from 859,532 in 2024 to 1,008,597 in 2025, with reported losses rising from more than \$16.6 billion to nearly \$21 billion (Federal Bureau of Investigation [FBI], 2025, 2026a).

The central problem is therefore not a universal lack of cybersecurity capability. The United States possesses exceptional technical expertise. The problem is that capability is distributed unevenly. Large enterprises, global cloud providers, major financial institutions, and national-security organizations may possess extensive defensive resources, while midsize companies, small businesses, nonprofits, healthcare providers, local governments, schools, manufacturers, and utilities may face comparable adversaries with dramatically fewer resources.

This paper argues that the United States is **highly cyber-capable but unevenly cyber-resilient**, and that AI is increasing the strategic cost of that inequality. It proposes an outcome-oriented regulatory model, a recurring technical-review mechanism for cybersecurity rules and standards, an author-defined principle of **Continuity Under Compromise**, and a 24-month resilience program intended to reduce high-leverage attack paths, accelerate containment, expand shared defensive capacity, govern AI agents before their authority becomes entrenched, and require critical organizations to prove that essential functions can continue while compromise is contained and recovery occurs.



Introduction

Twenty-nine minutes.

That was the average eCrime breakout time CrowdStrike reported for activity observed during 2025: the period between an attacker establishing an initial foothold and beginning lateral movement into another system. One reporting cycle earlier, the comparable figure was 48 minutes. The fastest observed breakout in the newer dataset was only 27 seconds (CrowdStrike, 2025, 2026).

Twenty-nine minutes is not simply a cybersecurity statistic. It is a measurement of how far an attacker may already have traveled while an alert is still waiting in a queue.

At a Fortune 100 corporation, those twenty-nine minutes may trigger identity analytics, endpoint isolation, session revocation, cloud correlation, threat-intelligence enrichment, and a global incident-response team that operates continuously. At a midsize manufacturer, the same intrusion may reach a small group responsible simultaneously for endpoint security, cloud infrastructure, vulnerability management, suppliers, applications, networking, and production technology. At a small business or nonprofit, there may be no dedicated cybersecurity team at all.

At a hospital, the consequences can move beyond stolen information. A physician may lose access to imaging. A nurse may lose medication history. Laboratory results, pharmacy workflows, scheduling, admissions, communications, and clinical documentation can suddenly move from digital systems to downtime procedures. At a water utility, the relevant technology may control pumps, treatment processes, remote access, storage, pressure, or chemical systems.

The attacker may use the same vulnerability, the same stolen credential, the same cloud service, and increasingly even the same AI tools. **What changes is the defender.**

That difference lies at the center of the cybersecurity problem examined in this research.

The United States possesses some of the most sophisticated cyber capabilities in the world. Intelligence organizations, military cyber units, federal cybersecurity institutions, global technology companies, cloud providers, telecommunications operators, banks, universities, national laboratories, security vendors, and independent researchers collectively represent extraordinary technical capability.

But the United States is not one network. It is millions of organizations with radically different resources operating through increasingly shared infrastructure.

A large multinational enterprise may maintain dedicated teams for identity engineering, detection engineering, cloud security, application security, threat intelligence, red teaming, incident response, malware analysis, digital forensics, and security architecture. One of its suppliers may employ three technology professionals. A national healthcare organization may run a sophisticated security operations center (SOC) while a rural healthcare provider depends almost entirely on outside vendors. A large manufacturer may employ operational technology (OT) security specialists while a smaller manufacturer uses the same administrator to maintain Microsoft 365, switches, backups, endpoint protection, business applications, and servers. A nonprofit can possess sensitive donor, employee, health, immigration, case-management, or client information while operating on a technology budget that would barely fund one senior security engineer at a major corporation.

Those differences matter because modern organizations do not operate independently. Hospitals depend on telecommunications, electricity, cloud services, identity providers, software companies, laboratories, medical-device manufacturers, pharmacies, insurers, and external service providers. Water utilities depend on communications infrastructure and remote engineering access. Financial institutions depend on cloud infrastructure, payment systems, DNS, telecommunications, identity, and external technology. Large enterprises depend on enormous supplier ecosystems. Small organizations increasingly depend on software-as-a-service platforms and managed



service providers they may lack the capacity to independently assess. Government itself increasingly depends on commercial technology.

A cybersecurity failure can therefore begin inside one organization and become somebody else's emergency.

AI is entering this environment at an unusually consequential moment. In early 2025, public threat-intelligence reporting largely described malicious use of generative AI as an accelerator for research, translation, scripting, reconnaissance, vulnerability investigation, and troubleshooting. By September 2026, Google Threat Intelligence Group was reporting forward-leaning adversaries moving from ordinary prompting toward agentic workflows and automation. In one observed case, a threat actor compromised cloud infrastructure and used AI tooling and agent instructions to plan, build, and execute a mass credential-harvesting campaign in less than six hours (Google Cloud, 2026b).

Anthropic's September 2026 threat-intelligence reporting describes an even broader transition. Across operations disrupted between December 2025 and August 2026, Anthropic observed state-linked actors, financial criminals, and politically motivated operators using multi-agent systems for reconnaissance, exploitation, credential harvesting, malware development, infrastructure management, data collection, and exfiltration. Humans remained involved in strategic decisions such as target selection and review of results, but increasing portions of the technical workload were delegated to machines (Anthropic, 2026).

Human attackers have not disappeared. They increasingly need to perform **less of the work personally**.

If reconnaissance that once consumed an hour requires minutes, more targets can be investigated. If vulnerability analysis accelerates, defenders have less time to remediate. If stolen datasets can be processed automatically, smaller victims become economically viable. If attack tooling can troubleshoot itself, one operator can supervise several campaigns rather than personally executing each technical step.

AI therefore does not need to become an autonomous super-hacker to transform cybersecurity. It only needs to make existing attacks faster, cheaper, more scalable, and easier to adapt.

At the same time, procurement may still take months. Regulatory reform can take years. Municipal budgets operate annually. Medical equipment may remain in service for fifteen or twenty years. Legacy enterprise and government systems can remain operational for decades.

This mismatch produces the central tension of the paper:

Machine-speed threat evolution is colliding with institution-speed defensive change.

The question is no longer whether AI will affect cybersecurity. That process has already begun. The question is whether the organizations society depends upon can become resilient before machine-assisted adversaries fully exploit the distance between America's strongest defenses and its weakest ones.



Research Approach and Evidentiary Limits

This research intentionally treats U.S. cybersecurity more broadly than cybersecurity of the U.S. government. National resilience includes federal, state, and local government, but also Fortune 100 and Fortune 500 enterprises, privately held companies, midsized firms, small and medium-sized businesses (SMBs), nonprofits, healthcare organizations, schools and universities, technology companies, manufacturers, telecommunications providers, financial institutions, defense suppliers, utilities, and other critical-infrastructure operators.

The analysis uses government reports, standards and regulatory material, threat-intelligence reporting, incident-response data, security telemetry, cyber-insurance information, ransomware-economics research, academic work, AI benchmarks, and publicly documented operational case studies available through September 15, 2026.

These sources measure different things. FBI IC3 complaints are not a census of all U.S. cybercrime. Verizon's DBIR is not equivalent to Mandiant's incident-response case population. CrowdStrike observes environments represented in its telemetry. Cyber-insurance claims reflect insured organizations. Threat-intelligence providers may have unusually deep visibility into particular adversaries. AI benchmarks measure performance under defined experimental conditions, not arbitrary real-world networks.

The paper therefore distinguishes four evidentiary categories. **Observed operational activity** means behavior documented during real campaigns or incident investigations. **Demonstrated capability** means performance achieved in a controlled benchmark, cyber range, evaluation, or experiment. **Forecasts and scenarios** describe possible future conditions and are identified as such. **Author analysis** represents interpretation derived from the evidence rather than a measurement reported by an external source.

This distinction is especially important when discussing AI. A model solving a capture-the-flag problem demonstrates technical capability. A threat-intelligence provider observing an agent execute exploitation and credential-harvesting tasks in a real campaign demonstrates operational use. Neither by itself proves that autonomous AI can reliably compromise any enterprise network it encounters.

The objective of this research is not to prove that cybersecurity is becoming universally better or worse. It is to identify how the threat environment is changing, which organizations are most exposed to those changes, and which defensive interventions would provide the greatest measurable increase in resilience.



PART I

The Threat Environment Is Moving Faster

MOSAIC STRATA RESEARCH



The most useful comparison between the 2025 and 2026 reporting cycles is not one statistic. It is the pattern that emerges when several independent datasets are considered together.


TABLE 1
Selected cybersecurity indicators across adjacent reporting cycles

Indicator	Earlier cycle	Later cycle	Interpretation
Verizon vulnerability exploitation	20%	31%	Strong increase in software exploitation as initial access
Verizon ransomware involvement	44%	48%	Continued growth in disruptive/extortion activity
Verizon third-party involvement	~30%	48%	Greater involvement of suppliers, software, and external relationships
CrowdStrike average eCrime breakout time	48 min	29 min	Less time available for detection and containment
CrowdStrike malware-free detections	79%	82%	Greater reliance on identities, legitimate tools, and hands-on-keyboard activity
Mandiant internal detection	43%	52%	Improvement in victim detection capability
Mandiant global median dwell time	11 days	14 days	Increase influenced by long-duration espionage and IT-worker cases
FBI IC3 complaints	859,532	1,008,597	More reported internet-crime complaints
FBI IC3 reported losses	>\$16.6B	nearly \$21B	Higher reported economic loss

Note. Vendor and incident-response datasets represent their respective populations and should not be interpreted as a national census (CrowdStrike, 2025, 2026; FBI, 2025, 2026a; Google Cloud, 2026a; Verizon, 2025, 2026).

Verizon's 2025 DBIR found vulnerability exploitation involved as an initial-access vector in 20% of relevant breaches, ransomware present in 44% of breaches, and third-party involvement in approximately 30%. Its 2026 report found vulnerability exploitation at 31%, making it the leading initial breach path in that report; ransomware reached 48%, and third-party involvement reached 48% (Verizon, 2025, 2026).





The shift complicates a message cybersecurity professionals have repeated for years: that the human user is the weakest link. Human behavior still matters enormously, but employee awareness cannot patch an internet-facing VPN. A phishing exercise cannot fix a vulnerable edge appliance. Training cannot remove an excessive OAuth grant, reduce privilege in a cloud service, validate a third-party software dependency, or produce security telemetry from infrastructure that was never instrumented. The attack surface has become systemic.

CrowdStrike's measurements show the same pressure through speed and stealth. Average eCrime breakout time fell from 48 minutes in the 2025 report to 29 minutes in the 2026 report, while the fastest observed breakout fell from 51 seconds to 27 seconds. Malware-free detections increased from 79% to 82%, and CrowdStrike reported an 89% year-over-year increase in activity by adversaries it classified as AI-enabled (CrowdStrike, 2025, 2026).

Mandiant provides evidence that defenders are improving as well. Organizations first detected their own compromises in 52% of Mandiant's 2025 investigations, compared with 43% during 2024. Global median dwell time nevertheless increased from 11 to 14 days, partly because the case mix included more long-running espionage and North Korean IT-worker incidents; those categories had median dwell times of 122 days (Google Cloud, 2026a).

The resulting story is not simply that cybersecurity is getting worse. **Both attackers and defenders are improving.** The difference is that attacker operations are becoming capable of scaling and adapting at machine-assisted speed while many defensive organizations remain limited by human queues, organizational handoffs, procurement, staffing, and policy.

The economic data reinforces the point. FBI IC3 complaints increased from 859,532 in 2024 to 1,008,597 in 2025. Reported losses rose from more than \$16.6 billion to nearly \$21 billion. The FBI also reported approximately 453,000 cyber-enabled fraud complaints with more than \$17.7 billion in losses during 2025 (FBI, 2025, 2026a). These figures measure reports submitted to IC3, not every cybercrime committed in the United States, but they show the scale at which cybersecurity and financial fraud are converging.

AI capability is rising, but benchmarks are not reality

Concern about AI should be grounded in measurable capability rather than fear.

The 2026 Stanford AI Index reports that the best unguided solve rate on Cybench, a framework containing forty professional-level cybersecurity challenge tasks, rose from approximately 15% in 2024 to approximately 82% in 2025 and 93% in 2026 (Stanford Institute for Human-Centered Artificial Intelligence [Stanford HAI], 2026).




Figure 2. Frontier AI performance on Cybench, 2024-2026. The benchmark measures professional-level cybersecurity challenge tasks; performance should not be interpreted as success against arbitrary enterprise networks. Source: Stanford HAI (2026).

That result is significant. It is also easy to misuse.

A realistic multi-host cyber-range benchmark published in 2026 produced much lower autonomous success. The strongest evaluated system solved 16.1% of web-exploitation tasks and 31.7% of post-exploitation tasks without detailed hints (Liu et al., 2026). The difference between 93% on a structured challenge benchmark and materially lower success in enterprise-like ranges is one of the most important findings in the AI portion of this paper.

AI cybersecurity capability is advancing extraordinarily quickly, but benchmark mastery should not be confused with generalized autonomous intrusion capability.

The same distinction applies to defense. A Cloud Security Alliance study involving more than 140 participants found that analysts using AI assistance completed two tested SOC investigation scenarios 45-61% faster and were 22-29% more accurate than manual counterparts. The study was conducted with an AI investigation product in simulated scenarios, so its findings should be interpreted as evidence for bounded analyst augmentation rather than proof of autonomous SOC capability (Cloud Security Alliance, 2025).

An open-ended 2026 threat-hunting benchmark produced the opposite limit case. Frontier models were given large Windows telemetry datasets and asked to identify malicious activity without guided questions. The strongest system correctly flagged only 3.8% of malicious events on average, and no evaluated model reached the study's threshold for unsupervised SOC deployment (Chona et al., 2026).

The evidence therefore supports AI **augmentation** far more strongly than AI **replacement**. AI can reduce analyst workload, accelerate bounded investigation, enrich alerts, generate queries, correlate evidence, and eventually execute carefully controlled response actions. It cannot yet be responsibly assumed to replace expert human judgment across an open-ended security environment.

Attack economics may matter more than autonomous capability

The most consequential AI effect may ultimately be economic.

Chainalysis estimates that claimed ransomware attacks rose approximately 50% during 2025 even while identified on-chain ransomware payments fell roughly 8% to about \$820 million. Its updated estimate for 2024 was \$892 million. Median ransom payments nevertheless increased 368%, from \$12,738 to \$59,556 (Chainalysis, 2026). Chainalysis notes that payment estimates are revised as new wallet attribution becomes available, so these values should be treated as evolving estimates rather than fixed final totals.

Coalition's insured population shows a related pattern. Initial ransom demands increased 47% in 2025 to an average above \$1 million, while 86% of affected policyholders refused to pay. Overall claims frequency increased 3%, while average claims severity fell 19%. Businesses with more than \$100 million in revenue experienced claims at five times the frequency of smaller businesses in Coalition's population, yet severity for that large-business segment declined, suggesting that larger organizations may have more resources to contain losses once incidents occur (Coalition, 2026).

These datasets are not directly comparable and should not be combined as if they represented the entire economy. Together, however, they suggest that resilience can change attacker profitability. Attackers respond by changing their own unit economics: increasing volume, stealing data rather than relying only on encryption, demanding more from successful victims, purchasing access from brokers, targeting dependencies, and using automation to reduce the cost of each attempted victim.

AI therefore does not need to produce a revolutionary zero-day to transform cybercrime. If it allows one operator to investigate ten organizations where the same operator previously investigated one, the business model has changed.



PART II

One Threat Environment, Unequal Defenders

MOSAIC STRATA RESEARCH



The phrase **private sector** is too broad to describe cybersecurity meaningfully. A Fortune 100 corporation, a 2,000-person manufacturer, a 70-person accounting company, and a five-person nonprofit all exist in the private economy. Their technical resources, attack surfaces, dependencies, tolerance for downtime, and ability to recover can be radically different.

The same is true across government and critical infrastructure. This inequality is where national cyber capability becomes national cyber resilience - or fails to.

Enterprise scale creates capability and complexity

Large global enterprises generally have access to the widest range of cybersecurity capabilities. Their central weakness is often not lack of tools but **complexity**. Multiple identity systems, acquisitions, subsidiaries, contractors, clouds, development platforms, SaaS applications, suppliers, APIs, legacy applications, privileged accounts, and international environments create an enormous number of trust relationships that must be governed simultaneously.

Coalition's 2026 claims analysis offers one indication of the tradeoff. Businesses with more than \$100 million in revenue had claims frequency five times higher than smaller organizations within its policyholder population, while severity for that segment declined year over year. The data cannot be generalized to every large enterprise, but it illustrates a familiar pattern: large organizations can be attractive because their attack surface is extensive, yet they may also have greater resources to contain damage (Coalition, 2026).

Midsized organizations often occupy a more difficult position. They may possess intellectual property, sensitive information, production environments, or supply-chain importance large enough to attract sophisticated adversaries while lacking enough scale to employ specialist teams for every security discipline.

SMBs confront a different reality. In Verizon's 2025 DBIR, ransomware was present in 88% of breaches in the SMB population, compared with 39% of breaches at larger organizations in the same report (Verizon, 2025). This does not mean an SMB has an 88% probability of suffering ransomware. It means ransomware represented a disproportionately large share of SMB breaches in Verizon's dataset.

For many small organizations, identity is effectively the perimeter. Email, payroll, accounting, file storage, communications, customer relationship management, payments, and collaboration may all exist in cloud services. One compromised administrator can therefore expose much of the organization without an attacker ever traversing a traditional corporate network.


Security for this population cannot realistically depend on every business building its own SOC. It must increasingly be delivered through secure defaults, managed identity, strong authentication, automated patching, recoverable backup, managed detection, and accessible incident response.

Nonprofits face many of the same constraints while sometimes handling information about donors, employees, patients, clients, children, immigrants, vulnerable populations, or community members. Their missions do not make them unattractive to criminals; limited resources can make them easier targets. National resilience therefore requires cybersecurity that works at **every economic scale**, not only for organizations able to purchase enterprise security.

Healthcare shows what happens when cybersecurity becomes mission failure

Healthcare deserves special attention because it combines nearly every major cyber-risk category in one environment: legacy systems, cloud computing, highly sensitive information, medical devices, third parties, complex identities, remote support, ransomware, AI adoption, regulatory requirements, and systems that cannot always be taken offline.

The FBI's critical-infrastructure data illustrate the pressure. In 2024, the healthcare and public health sector recorded 238 ransomware complaints and 206 data-breach complaints, for a combined 444 reported events. In 2025, ransomware complaints increased to 460 while data-breach complaints declined to 182, producing 642 combined reported events (FBI, 2025, 2026b).



Ransomware reports therefore increased by approximately 93% in one year in that dataset, while the combined ransomware/data-breach total increased by approximately 45%.

HHS provides a different measurement. Its 2024 report to Congress records 663 breaches affecting at least 500 individuals that occurred during the year. Those incidents affected approximately 242.9 million people. Hacking/IT incidents represented 81% of the large breaches and affected approximately 241.6 million individuals (U.S. Department of Health and Human Services [HHS], 2026a). These are HIPAA breach notifications, not the same population as FBI cyber-threat complaints, and the two series should not be merged.

The scale of the Change Healthcare incident demonstrates why incident counts alone are insufficient. HHS later reported that approximately 192.7 million individuals were affected by that incident (HHS, 2025). A compromise of one hospital is serious; a compromise of an intermediary capable of affecting payment, pharmacy, clearinghouse, or administrative functions across thousands of providers is systemic.

Healthcare cybersecurity is therefore not simply a privacy problem. It is a **continuity-of-care problem**.

When ransomware strikes an ordinary company, email, files, enterprise applications, or customer systems may become unavailable. When the same type of event reaches a hospital, a physician may lose imaging, a nurse may lose access to medication history, a laboratory may revert to manual workflows, pharmacy systems may be impaired, emergency departments may slow, and ambulances may be diverted.

The central healthcare question should consequently become:

If critical technology disappears at 2:00 a.m., can the organization safely treat the patient arriving at 2:01?

That question measures something a compliance checklist cannot.

Healthcare scenario methodology for 2027

A precise 2027 prediction would create false confidence because full-year 2026 FBI data is not available at the evidence cutoff. The manuscript therefore uses a transparent scenario method rather than a probabilistic forecast.

The baseline is the **642 combined healthcare/public health ransomware and data-breach complaints reported by FBI IC3 for 2025**. The 2027 scenario is calculated as:

$$N(2027) = 642 \times (1 + g)^2$$


where g is an illustrative annual growth assumption applied for 2026 and 2027.

TABLE 2

Illustrative healthcare cyber-event scenarios for 2027

Scenario	Annual growth assumption	Calculation	Illustrative 2027 total
Stabilization	0%	642×1.00^2	642
Moderate growth	10%	642×1.10^2	777
Continued escalation	20%	642×1.20^2	925

Note. These values are scenario outputs, not forecasts or confidence intervals. They intentionally use a single FBI series and do not combine partial 2026 HHS breach notifications, attacker-claimed ransomware trackers, or insurer data with the FBI annual complaint baseline.



The value of this exercise is not the exact count. It is to show that even modest sustained growth would keep healthcare under severe pressure, while a flat event count would still leave the sector operating at the elevated 2025 level.

Critical infrastructure reveals the limits of organization-by-organization defense

Healthcare is not alone in being dependent on systems outside its direct control. Water depends on telecommunications. Telecommunications depends on electricity. Healthcare depends on both. Financial services depend on telecommunications, DNS, power, cloud infrastructure, and payment systems. Manufacturing depends on power, logistics, suppliers, communications, and increasingly connected operational environments. Government depends heavily on commercial technology.

The relevant security unit is therefore sometimes not the organization. It is the **dependency chain**.

A utility does not need to be directly compromised for cyber disruption elsewhere to affect operations. Loss of telecommunications can eliminate normal remote access. Loss of identity infrastructure can prevent cloud authentication. A compromised vendor can remove support precisely when it is most needed. National resilience requires organizations to understand these dependencies and test what occurs when they disappear.

Water demonstrates both exposure and a measurement problem

The water and wastewater sector provides one of the clearest examples of unequal defensive capacity.

GAO describes close to 170,000 drinking-water and wastewater systems across the United States and highlights workforce shortages, aging technology, limited resources, increasing IT/OT connectivity, and competition between cybersecurity spending and legally required investments in safe water infrastructure (U.S. Government Accountability Office [GAO], 2026a).


Unlike healthcare, the public evidence does not currently support a clean national year-over-year count of every water-sector cyberattack. That absence matters. **It is difficult to manage national cyber risk when the country cannot confidently measure the incident rate of infrastructure on which virtually every community depends.**

Exposure data provides part of the picture. In 2024, EPA's Office of Inspector General passively assessed 1,062 large drinking-water systems serving more than 193 million people. Ninety-seven systems serving approximately 26.6 million people had critical or high-risk cybersecurity vulnerabilities, and another 211 systems had externally visible medium- or lower-risk exposures (U.S. Environmental Protection Agency Office of Inspector General [EPA OIG], 2024).

EPA separately reported identifying cybersecurity vulnerabilities at 277 water systems during 2025 and working with operators on remediation. Because that activity used a different assessment process and population, the 277 figure must not be compared directly with EPA OIG's 97 high/critical findings as though they were a year-over-year trend (U.S. Environmental Protection Agency [EPA], 2026).


WaterISAC provides a limited but useful directional indicator. Five percent of respondents to its Q2 2025 survey reported at least one cybersecurity incident. In Q4, 8% did; industrial control systems were involved in 34% of reported Q4 cyber incidents, while ransomware accounted for 23%. Respondent populations can vary across quarters, so the percentages are not a national incident rate (WaterISAC, 2026).



MOSAIC
STRATA

A responsible 2027 water-sector outlook should therefore remain qualitative. Absent significant improvements in asset visibility, remote-access security, segmentation, shared defensive capability, and regulatory authority, risk should be expected to remain elevated because the structural constraints documented in 2024-2026 remain unresolved. GAO also reports that EPA identified important gaps in its legal authority to require cybersecurity risk assessments for portions of the sector (GAO, 2026a).

That is a governance problem as much as a technical one.



PART III

The Attack Surface Has Moved

MOSAIC STRATA RESEARCH



The older cybersecurity model imagined an external attacker crossing a firewall into an internal network. Modern environments do not operate that way.

Identity connects SaaS, cloud infrastructure, endpoints, APIs, service accounts, applications, automation, suppliers, and AI agents. An attacker possessing a legitimate token may not require malware. An OAuth grant may bypass the need for a password. A compromised workload identity can possess infrastructure privileges. A SaaS administrator can change business operations without ever entering a traditional LAN.

CrowdStrike's finding that 82% of its 2025 detections were malware-free is therefore more than an endpoint statistic. It illustrates a broader transition toward trusted identities, legitimate tools, cloud services, and normal-looking administration being used by an invalid actor (CrowdStrike, 2026).

The practical perimeter increasingly resembles:

Identity SaaS Cloud Endpoint API Automation Data AI

This changes the SOC's central question from **Was authentication successful?** to **Was the authenticated behavior legitimate?**

The software supply chain is becoming an identity chain

Third-party involvement reached 48% of breaches in Verizon's 2026 DBIR. Verizon's measure is broad and should be interpreted as third-party **involvement**, not proof that a third party caused 48% of breaches (Verizon, 2026).

Modern software can move through a trust chain from a developer workstation to source control, continuous integration/continuous delivery, workload identity, cloud infrastructure, and production. AI adds another layer: coding assistants, generated dependencies, model APIs, agent credentials, automated deployment, connectors, retrieval stores, prompts, and external inference providers.

A production endpoint can therefore be well defended while the system remains vulnerable through the identity and development chain that created it. Supply-chain security is increasingly **identity security plus provenance plus authority**.

AI agents are privileged identities before they are intelligent coworkers

NIST launched its AI Agent Standards Initiative in February 2026 with explicit attention to agent security, interoperability, and identity. A related NCCoE concept paper examines identification, authorization, auditing, non-repudiation, and controls relevant to software and AI agents (National Institute of Standards and Technology [NIST], 2026a, 2026b).

An AI agent may simultaneously behave like a user, application, automation system, administrator, and attack target. It may possess memory, credentials, tools, network access, API access, delegated authority, and communication with other agents.

That leads to a principle that should be adopted before agentic systems become deeply embedded:

Intelligence must never automatically imply authority.

A more capable model should not receive more privilege merely because unrestricted access makes implementation easier. Production agents capable of consequential action should have individual machine identities, accountable owners, limited credentials, explicit tool permissions, restricted network access, auditable action histories, independent policy enforcement, and emergency revocation mechanisms that function outside the agent itself.

Those controls do not require future AI research. They can be implemented now.



Post-quantum migration illustrates why cybersecurity must prepare before crisis

The same principle applies to cryptography.

Executive Order 14412 directs covered federal high-value assets and high-impact systems toward NIST-approved post-quantum cryptography (PQC) for key establishment by December 31, 2030, and digital signatures by December 31, 2031. It also directs federal agencies and sector risk management agencies to support broader migration planning, including assistance to critical infrastructure (The White House, 2026).

The relevant 24-month task is not to complete migration everywhere. It is to establish cryptographic visibility. NIST's PQC migration work emphasizes building and maintaining a cryptographic inventory covering algorithms, protocols, certificates, systems, devices, applications, and data flows (NIST National Cybersecurity Center of Excellence [NCCoE], 2026).

The lesson is broader than quantum computing. Security programs repeatedly suffer when they wait until migration becomes urgent before discovering what they actually depend upon.



PART IV

From Compliance to Resilience

MOSAIC STRATA RESEARCH



The United States already possesses extensive cybersecurity frameworks. NIST Cybersecurity Framework (CSF) 2.0 organizes cybersecurity risk around six continuous functions: Govern, Identify, Protect, Detect, Respond, and Recover (NIST, 2024). CISA's Cross-Sector Cybersecurity Performance Goals (CPGs) intentionally provide a smaller group of high-impact, measurable practices designed to reduce risk and to be usable by organizations with limited resources, including smaller entities (Cybersecurity and Infrastructure Security Agency [CISA], n.d.).

The problem is not that the country has no idea what good cybersecurity looks like. The problem is translating good practice into outcomes across organizations with radically different resources - and ensuring regulatory requirements remain aligned with modern technology.

GAO identified 117 federal cybersecurity regulations established by 37 agencies across nine critical-infrastructure sectors. Eighty of those regulations, approximately 70%, contained at least one category of reporting requirement also found in another regulation (GAO, 2026b). This does not mean cybersecurity regulation should disappear. It means regulation should become **more technically current, more outcome-oriented, and less duplicative**.

Cybersecurity rules need technical renewal cycles

Technology changes faster than law. A requirement can remain legally valid while becoming technically inadequate.

Healthcare provides a powerful example. HHS describes its pending HIPAA Security Rule modernization as the first major update to the Security Rule since 2013. The proposal would make important changes including mandatory multifactor authentication with limited exceptions, encryption of electronic protected health information (ePHI) at rest and in transit with limited exceptions, network segmentation, vulnerability scanning at least every six months, penetration testing at least annually, stronger backup and recovery controls, annual testing of security measures, and procedures to restore designated systems and data within 72 hours (HHS, 2024a, 2024b). The current Security Rule remains in effect while rulemaking continues.

These are not exotic 2026 innovations. Many are now considered fundamental security practices. The lag illustrates the speed mismatch between technology and regulation.

ISO/IEC 27001 should be treated accurately here: it is an international standard, not U.S. law. Its lifecycle nevertheless provides useful precedent for periodic technical review. ISO/IEC 27001:2022 replaced the 2013 edition, and ISO's directives provide for systematic review of International Standards within a maximum five-year interval, with review possible earlier when circumstances warrant it (International Organization for Standardization [ISO], 2022, n.d.).

This paper therefore proposes a similar concept for high-impact cybersecurity regulation. Normal technical-sufficiency reviews should occur approximately every **three years, with five years as the maximum period** before formal reassessment of requirements governing critical or high-consequence environments.

Scheduled review should not be the only mechanism. Major technological changes - agentic AI, major cryptographic breakthroughs, new identity architectures, a catastrophic supply-chain event, new classes of critical-infrastructure compromise, or fundamental changes in attacker tradecraft - should trigger out-of-cycle assessment.

The panel conducting technical review should not consist only of policymakers and lawyers. Legal expertise is essential, but technical adequacy should also be evaluated by practicing professionals in incident response, identity, cloud architecture, security engineering, penetration testing, application security, OT, healthcare technology, digital forensics, AI security, cryptography, recovery, and networking.

Sector expertise matters. Hospital requirements should involve people who understand clinical downtime. Water requirements should involve professionals who understand PLCs, SCADA, pumps, treatment processes, remote access, and plant operations. SMB and nonprofit representation should ensure that regulations do not become technically impressive but economically impossible.

Law should establish durable security outcomes. Regulation should establish measurable obligations. Technical standards should define acceptable implementations and update more rapidly as technology changes. That separates legal durability from technical stagnation.

Minimum compliance should not be mistaken for security

A minimum security requirement exists because anything below it is unacceptable. It should not mean everything above the line is secure.

TABLE 3

Proposed distinction between compliance, capability, and resilience

Level	Meaning
Minimum security floor	Controls below which an organization should not operate
Operational security capability	Ability to detect, contain, and materially restrict an active attacker
Cyber resilience	Ability to continue essential operations and recover when defensive controls fail

A successful audit can show that required controls were observed at a particular time. It does not automatically prove that the organization can withstand an active intrusion.

A hospital should not be considered resilient simply because an incident-response policy exists. It should demonstrate that care continues when major technology is unavailable. A utility should not be considered resilient because a risk assessment was completed. It should demonstrate safe operations when normal remote connectivity disappears. A company should not be considered resilient because it owns backups. It should restore them.

Continuity Under Compromise

The central policy concept proposed by this paper is **Continuity Under Compromise**.

Traditional security programs often organize around preventing unauthorized access. That remains essential, but mature resilience begins with the assumption that at some point a valid credential will be stolen, an endpoint will be compromised, a third party will fail, a vulnerability will be exploited, or an AI agent will behave unexpectedly.

The critical question then becomes what happens afterward. Can the attacker move? Can privilege expand? Can backups be reached? Can OT be affected? Can the organization still operate? Can containment happen while services continue? Can recovery occur without restoring the compromise?

Continuity Under Compromise therefore defines cybersecurity not only as the ability to resist intrusion but as the ability to **perform essential functions while portions of the environment are untrusted, isolated, unavailable, or under active recovery**.

For healthcare, this means maintaining safe clinical operations during EHR, identity, vendor, or connectivity failure. For water, it means maintaining safe treatment during loss of remote access or ordinary IT services. For finance, it means preserving essential customer and payment capabilities while compromised infrastructure is isolated. For business, it means preserving minimum viable operations and restoring systems without paying an attacker or reintroducing the same compromise.

The doctrine becomes:

Prevent where possible. Detect quickly. Contain decisively. Operate safely. Recover cleanly. Learn continuously.

Prevention is Plan A. Continuity during compromise must be Plan B. Recovery without reinfection must be Plan C. An organization with only Plan A is not resilient.



Author's analytical position

The evidence reviewed for this paper leads me to a specific conclusion: the United States does not primarily suffer from a lack of cybersecurity knowledge. It increasingly suffers from inconsistent execution of what is already known.

We know privileged identities need stronger protection. We know exposed systems need rapid remediation. We know vendor access should be controlled. We know default credentials are dangerous. We know backups require isolation and restoration testing. We know OT should not be unnecessarily exposed to the internet. We know identity has become central to cloud security. We know security teams need visibility.

Yet the same weaknesses repeatedly appear in incident investigations.

In my assessment, cybersecurity has therefore become increasingly an **execution problem**. AI makes that execution gap more dangerous. An attacker can adopt a new model immediately. A hospital may need years to replace a clinical system. An adversary can automate reconnaissance tonight. A municipality may need the next budget cycle to hire one person. A criminal group can experiment continuously. Regulation can remain materially unchanged for a decade.

I also believe cybersecurity policy has historically placed too much emphasis on proving that controls and procedures exist and too little emphasis on proving that the organization survives when those controls fail. Compliance, standards, audits, and regulation all matter. None should be confused with resilience.

The question should increasingly move from **Did you satisfy the cybersecurity requirement?** to **Can you prove that the organization continues its essential mission when the requirement is tested by a real attacker?**



PART V

The Next Twenty-Four Months

MOSAIC STRATA RESEARCH



The United States cannot eliminate cyber risk by September 2028. It can materially change the attacker's economics and substantially reduce the consequences of successful compromise.

The next 24 months should therefore be treated as an acceleration window rather than an attempt to solve every security problem. The program should build on existing frameworks such as NIST CSF 2.0 and CISA's prioritized CPG model rather than create another parallel control universe (CISA, n.d.; NIST, 2024).

TABLE 4

Proposed 24-month cyber-resilience acceleration program

Period	Primary objective	Target outcome
0-90 days	Establish visibility	Critical assets, identities, vendors, dependencies, AI agents, recovery systems, OT paths, and major cryptographic dependencies identified
3-6 months	Collapse high-leverage attack paths	Strong privileged authentication, reduced standing privilege, protected remote access, faster edge remediation, isolated recovery
6-12 months	Close capability inequality	Shared public-sector defense expands; SMB managed security improves; AI agents receive formal governance
12-18 months	Reduce response latency	High-confidence reversible containment moves toward machine speed; cross-sector exercises expand
18-24 months	Prove resilience	Critical organizations demonstrate Continuity Under Compromise and clean recovery

The first 90 days: visibility before technology

Before another major security purchase, organizations should be able to identify which systems are internet-facing, which identities possess critical privilege, which vendors connect remotely, which cloud and SaaS integrations hold persistent authority, which service and workload accounts can alter infrastructure, which AI agents can take consequential action, which systems control recovery, which OT environments are externally reachable, and which outside services are necessary for minimum viable operations.

The objective is not perfect inventory. It is visibility into systems whose compromise could produce disproportionate consequence.

Months three through six: remove the attacker's easiest leverage

Privileged and remote administrative access should move toward phishing-resistant authentication wherever feasible. Standing administrative privilege should be reduced. Internet-facing management interfaces should be eliminated unless operationally necessary. Vendor access should be strongly authenticated, monitored, and constrained.

Vulnerability management should prioritize **active exploitation plus exposure plus consequence**, not simply numerical severity. An internet-facing VPN gateway with active exploitation is not equivalent to an internal low-value workstation vulnerability that happens to possess a similar CVSS score.

Recovery environments should also be isolated from primary production identity. A compromised enterprise domain should not automatically provide authority to destroy the systems required to rebuild the enterprise.

By the end of this phase, a stolen password should be significantly less likely to equal ownership of the organization.

Months six through twelve: close the capability gap

This phase addresses the inequality at the center of the research.



A municipality with twelve thousand residents should not need its own threat-intelligence team, but it should have access to one. States and regional organizations should expand shared SOC, endpoint detection and response, vulnerability-management, identity, threat-hunting, incident-response, backup-validation, and OT-support capability for public institutions unable to maintain these services independently.

Private SMBs need the equivalent through managed services, secure SaaS platforms, strong default configurations, cloud-provider controls, and accessible incident response. Healthcare needs specialized regional capability for small hospitals, clinics, and providers that cannot recreate the security structure of a national hospital network.

AI governance must also mature during this phase. Every production agent capable of consequential action should have a registered owner, unique identity, enumerated tool permissions, scoped credentials, durable audit history, and emergency revocation mechanism.

Months twelve through eighteen: containment must begin approaching attacker speed

If attackers can move laterally within 29 minutes, high-confidence incidents cannot routinely remain in analyst queues for several hours.

The solution is not unrestricted autonomous defense. It is **bounded automated containment**.

High-confidence, reversible actions such as revoking a suspicious session, suspending an exposed OAuth grant, isolating a compromised endpoint, forcing reauthentication, quarantining malicious email, or temporarily disabling a confirmed compromised credential can increasingly occur automatically under carefully defined policy.

The automation should preserve evidence and notify humans immediately. Higher-consequence actions - disabling medical equipment, shutting down industrial systems, deleting cloud resources, or materially interrupting business operations - should require substantially stronger controls and human authorization.

The principle is simple: **machines can stop the bleeding quickly; humans should retain authority over consequential treatment**.

Months eighteen through twenty-four: stop assuming and start proving

The final phase changes the definition of success.

A hospital should demonstrate that critical care continues while selected systems are unavailable. A utility should demonstrate safe operation during telecommunications degradation. A business should restore production systems from protected backup. An enterprise should prove identity recovery. A municipality should demonstrate how external incident-response support arrives when internal personnel are overwhelmed. An AI environment should demonstrate that a malfunctioning or compromised agent can be terminated independently of itself.

The metric is no longer **We have a recovery plan**. It becomes **We recovered**.

By September 2028, success should not mean zero breaches. It should mean a different consequence profile: a stolen password does not automatically become privileged access; a compromised endpoint does not automatically become enterprise compromise; enterprise compromise does not automatically destroy backup; a vendor compromise does not automatically create unrestricted production access; an AI agent cannot exceed independently enforced authority; a small organization does not face a sophisticated adversary alone; a hospital keeps treating patients; a utility continues operating safely; and verified malicious activity is contained in minutes rather than waiting until morning.

Outlook: What Happens If We Do Nothing?

Long-term forecasting in cybersecurity is unreliable when presented as certainty. A conditional scenario is more useful.

Suppose that through 2028 and 2030, AI vulnerability research continues improving while remediation still takes weeks. Agentic systems become commonplace but are frequently deployed with broad credentials. Cloud identity connects more systems. Third-party dependencies expand. Software pipelines incorporate more generated code and automated deployment. OT becomes more connected. Smaller organizations remain unable to hire specialized talent.

Attackers automate reconnaissance, credential validation, exploit attempts, malware modification, and stolen-data processing across many victims simultaneously. Defenders adopt AI too, but primarily to summarize alerts rather than redesign architecture, identity, recovery, and containment.

What fails first?

Possibly not the technology. Possibly **institutional capacity**.

Incident-response teams become saturated. Local governments cannot obtain assistance quickly enough. Hospitals exceed safe downtime tolerances. A compromised technology supplier affects hundreds of organizations simultaneously. Telecommunications disruption removes dependencies other sectors assumed would always be available. Regulatory duplication consumes staff during the exact moment engineers are needed elsewhere. Recovery systems exist but have never been meaningfully tested.

The resulting failure may not resemble a cinematic AI cyberwar. It may look much more ordinary - and therefore be more dangerous: a persistent stream of machine-assisted attacks repeatedly overwhelms organizations that society depends on because defensive institutions cannot move at the same speed.

That outcome is not inevitable. Most of its enabling conditions are already visible.



Conclusion

The United States does not face a future in which AI suddenly creates cybersecurity risk where none existed before. It faces something more difficult: a future in which AI increasingly discovers, scales, automates, and exploits weaknesses that have been visible for years.

The exposed systems are already here. The permanent privileges are already here. The forgotten vendor accounts are already here. The fragmented identities are already here. The legacy applications are already here. The aging medical systems are already here. The third-party dependencies are already here. The hospitals that cannot simply shut technology down are already here. The utilities with limited cybersecurity personnel are already here. The nonprofits operating on minimal budgets are already here. The midsize companies important enough to attack but too small to employ every security specialty are already here.

AI changes the equation because an adversary increasingly requires less time, less labor, and in some cases less specialized expertise to discover and exploit those weaknesses.

That makes the next twenty-four months unusually important.

The United States does not need to eliminate every vulnerability by September 2028. It cannot. Every small company does not need a twenty-four-hour SOC. Every municipality does not need its own threat-intelligence organization. Every hospital does not need a malware laboratory. AI should not receive unrestricted authority over critical systems simply because adversaries are becoming faster.

What the country needs is far more achievable. It needs fewer exposed paths into critical systems and fewer permanent privileges once an attacker enters. It needs identity systems designed around the assumption that valid credentials can be stolen. It needs infrastructure telemetry where attackers increasingly hide. It needs healthcare designed around continuation of care rather than confidentiality alone. It needs water and other resource-constrained infrastructure supported with both requirements and implementation capacity. It needs small organizations to consume real cybersecurity capability without recreating large-enterprise security departments. It needs AI agents governed as privileged machine identities before their authority becomes invisible and permanent. It needs cybersecurity laws and standards reviewed at a technical pace compatible with the technology they regulate. And it needs recovery systems that are **proven rather than assumed**.

Cybersecurity should therefore no longer be judged only by whether an organization prevents intrusion. It should also be judged by whether the organization can detect the attacker, restrict what the attacker can reach, continue its essential mission while containment occurs, and recover without recreating the compromise.

America already possesses extraordinary cyber capability. The task now is making that capability reach the large enterprise, the midsize manufacturer, the family business, the nonprofit, the school, the hospital, the clinic, the telecommunications provider, the water utility, the power operator, and the local government **before the attacker does**.

Because national resilience will never be determined solely by how difficult it is to compromise America's strongest network. It will be determined by what happens to the systems America cannot afford to lose **after somebody gets in**.

Preventing compromise is security.

Continuing the mission while compromised is resilience.

Twenty-nine minutes is the warning.

Twenty-four months is the opportunity.

The question is whether we use it.



References

- Anthropic. (2026, September 10). Detecting and countering misuse of AI: September 2026. <https://www.anthropic.com/threat-intelligence-report-september-2026>
- Chainalysis. (2026, February 26). Crypto ransomware: 2026 Crypto Crime Report. <https://www.chainalysis.com/blog/crypto-ransomware-2026/>
- Chona, A., Kozlov, I., & Kumar, A. (2026). Cyber Defense Benchmark: Agentic threat hunting evaluation for LLMs in SecOps [Preprint]. arXiv. <https://arxiv.org/abs/2604.19533>
- Cloud Security Alliance. (2025, October 6). Beyond the hype: A benchmark study of AI agents in the SOC. <https://cloudsecurityalliance.org/artifacts/a-benchmark-study-of-ai-agents-in-the-soc>
- Coalition. (2026, March 5). 2026 Cyber Claims Report. <https://www.coalitioninc.com/claims-report/2026>
- CrowdStrike. (2025, February 27). 2025 Global Threat Report: The rise of the enterprising adversary. <https://www.crowdstrike.com/en-us/blog/crowdstrike-2025-global-threat-report-findings/>
- CrowdStrike. (2026, February 24). 2026 Global Threat Report: AI accelerates adversaries and reshapes the attack surface. <https://www.crowdstrike.com/en-us/press-releases/2026-crowdstrike-global-threat-report/>
- Cybersecurity and Infrastructure Security Agency. (n.d.). Cross-Sector Cybersecurity Performance Goals. <https://www.cisa.gov/cross-sector-cybersecurity-performance-goals>
- Federal Bureau of Investigation. (2025). 2024 Internet Crime Report. Internet Crime Complaint Center. https://www.ic3.gov/AnnualReport/Reports/2024_IC3Report.pdf
- Federal Bureau of Investigation. (2026a, April 6). Cryptocurrency and AI scams bilk Americans of billions. <https://www.fbi.gov/news/press-releases/cryptocurrency-and-ai-scams-bilk-americans-of-billions>
- Federal Bureau of Investigation. (2026b). 2025 Internet Crime Report. Internet Crime Complaint Center. https://www.fbi.gov/file-repository/2025_ic3report.pdf
- Google Cloud. (2026a, March 23). M-Trends 2026: Data, insights, and strategies from the frontlines. <https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026>
- Google Cloud. (2026b, September 8). GTIG AI Threat Tracker: From prompting to autonomy - the evolution of adversarial AI. <https://cloud.google.com/blog/topics/threat-intelligence/from-prompting-to-autonomy-the-evolution-of-adversarial-ai>
- International Organization for Standardization. (2022). ISO/IEC 27001:2022 - Information security, cybersecurity and privacy protection - Information security management systems - Requirements. <https://www.iso.org/standard/27001>
- International Organization for Standardization. (n.d.). Review of International Standards (confirmation, revision, withdrawal). https://www.iso.org/sites/ConsumersStandards/1_standards.html
- Liu, F., Dai, J., Fan, Y., Mai, W., Li, Z., Chen, B., Zhang, J., Lou, Z., Xiang, B., Zhang, Q., Pan, X., Hong, G., Zhang, Y., & Yang, M. (2026). AgentCyberRange: Benchmarking frontier AI systems in realistic cyber ranges [Preprint]. arXiv. <https://arxiv.org/abs/2606.14295>
- National Institute of Standards and Technology. (2024). The NIST Cybersecurity Framework (CSF) 2.0. https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=957258
- National Institute of Standards and Technology. (2026a, February 17). Announcing the AI Agent Standards Initiative for interoperable and secure innovation. <https://www.nist.gov/news-events/news/2026/02/announcing-ai-agent-standards-initiative-interoperable-and-secure>
- National Institute of Standards and Technology. (2026b, February 5). Accelerating the adoption of software and artificial intelligence agent identity and authorization [Concept paper]. <https://csrc.nist.gov/pubs/other/2026/02/05/accelerating-the-adoption-of-software-and-ai-agent/ipd>
- NIST National Cybersecurity Center of Excellence. (2026). Frequently asked questions about post-quantum cryptography: Migration to post-quantum cryptography. <https://pages.nist.gov/nccoe-migration-post-quantum-cryptography/>
- Stanford Institute for Human-Centered Artificial Intelligence. (2026). AI Index Report 2026. Stanford University. https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf
- The White House. (2026, June 22). Executive Order 14412: Securing the nation against advanced cryptographic attacks. <https://www.whitehouse.gov/presidential-actions/2026/06/securing-the-nation-against-advanced-cryptographic-attacks/>
- U.S. Department of Health and Human Services. (2024a, December 27). HIPAA Security Rule Notice of Proposed Rulemaking to strengthen cybersecurity for electronic protected health information: Fact sheet. <https://www.hhs.gov/hipaa/for-professionals/security/hipaa-security-rule-nprm/factsheet/index.html>
- U.S. Department of Health and Human Services. (2024b, December 27). HIPAA Security Rule NPRM. <https://www.hhs.gov/hipaa/for-professionals/security/hipaa-security-rule-nprm/index.html>

- U.S. Department of Health and Human Services. (2025). Change Healthcare cybersecurity incident frequently asked questions. <https://www.hhs.gov/hipaa/for-professionals/special-topics/change-healthcare-cybersecurity-incident-frequently-asked-questions/index.html>
- U.S. Department of Health and Human Services. (2026a). Annual report to Congress on breaches of unsecured protected health information for calendar year 2024. <https://www.hhs.gov/sites/default/files/breach-report-to-congress-2024.pdf>
- U.S. Environmental Protection Agency. (2026, February 6). EPA actions help safeguard water systems from cyberattacks. <https://www.epa.gov/newsreleases/epa-actions-help-safeguard-water-systems-cyberattacks>
- U.S. Environmental Protection Agency Office of Inspector General. (2024, November 13). Management implication report: Cybersecurity concerns related to drinking water systems (Report No. 25-N-0004). <https://www.epa.gov/office-inspector-general/report-management-implication-report-cybersecurity-concerns-related>
- U.S. Government Accountability Office. (2026a, May 21). Critical infrastructure protection: Actions needed to address persistent cybersecurity threats to the water and wastewater sector (GAO-26-109159). <https://www.gao.gov/products/gao-26-109159>
- U.S. Government Accountability Office. (2026b, July 22). Cybersecurity regulations: Multiple sectors are subject to potentially duplicative reporting requirements (GAO-26-108606). <https://www.gao.gov/products/gao-26-108606>
- Verizon. (2025). 2025 Data Breach Investigations Report. <https://www.verizon.com/business/resources/T99/reports/2025-dbir-data-breach-investigations-report.pdf>
- Verizon. (2026). 2026 Data Breach Investigations Report. <https://www.verizon.com/business/resources/reports/dbir/>
- WaterISAC. (2026). 2025 Q4 incident report summary. <https://www.waterisac.org/wp-content/uploads/2026/03/WaterISAC-Exec-Summary-Q4-25-1.pdf>





TWENTY-NINE MINUTES

is the warning.

TWENTY-FOUR MONTHS

is the opportunity.

The question is whether we use it.