CLOUDFLARE

2026

Cloudflare Security Signals Report

Autonomic Resilience

FOREWORD BY MICHELLE ZATLYN

Everything is changing.

AI is moving from pilot to production, autonomous systems are accelerating decision-making, and the digital economy is evolving in real time. For leaders ready to act, that pace of change creates real opportunity.

Resilience has become the new competitive edge. As intelligent systems reshape the digital economy, leaders can design defenses to anticipate change, engineer systems that adapt, and turn volatility into an advantage.

Cloudflare operates one of the world's largest global networks, spanning more than 330 cities in over 120 countries. We protect millions of Internet properties, stop over 230 billion cyberattacks every day, and handle 2.5 billion bot requests daily. From this vantage point, we see both the risks and the opportunities shaping the Internet.

The ***2026 Cloudflare Security Signals Report*** delivers practical insights leaders need today, mapping the forces reshaping the digital landscape, so you can govern intelligent systems, modernize securely, and build resilience from the core.

We're on a mission to help build a better Internet. In 2026, that means helping you operate securely and confidently — at machine speed.

Michelle Zatlyn
Co-founder, President,
and Co-chair, Cloudflare

EXECUTIVE SUMMARY

For today's highly interconnected and automated enterprises, the "absorb shocks and recover" model no longer works.

This approach relies on the naive assumption that we can accurately predict and prepare for every specific disruption. AI systems act autonomously; cloud platforms concentrate critical workloads; supply chains extend deep into opaque ecosystems. In this new reality, security leaders require **autonomic resilience: systems that do more than withstand stress — they regulate, adapt, and recover in real time**.

But while many organizations appear mature, modern, and well-governed, autonomic resilience is not visible in steady state. It is a leadership outcome revealed only under sustained and severe stress.

This report is built on a simple premise: The greatest risks enterprises will face in 2026 do not come from obvious weaknesses. They emerge from hidden fault lines, areas that look sound in normal operations, but fracture when speed, scale, or disruption increase.

Within these chapters, we provide executives with a blueprint to surface these fault lines before they break. Each section offers pointed questions to spark internal debate and uncover hidden fragility within their own organizations. In an era of intelligence, autonomy, and speed, success belongs to leaders who design their enterprises to sense, adapt, and self-correct under stress, while protecting critical outcomes as conditions change.

Six critical fault lines

These fault lines do not stand alone. Pressure in one area can intensify weakness in others.

1 Taming the algorithm: Governing AI at scale

AI programs often appear disciplined, governed, and value-driven. Yet under scrutiny, many leaders cannot clearly explain where AI is running, what data it touches, or who is accountable when outcomes fail. Progress on the surface often masks a visibility and ownership gap that becomes exposed when regulators, customers, or incidents apply pressure.

2 Trust at machine speed: Engineering autonomy

Autonomous systems perform well when conditions are predictable. Under stress, decisions move faster than human oversight, and trust is assumed rather than engineered. This fault line tests whether delegation was deliberate or whether authority quietly shifted to machines without clear boundaries, accountability, or real-time control.

3 Shadow supply chains: Exposing hidden dependencies

Enterprises appear diversified and partner-rich, but depend on layers of third- and fourth-party services they do not fully see. When disruption occurs, the first failure is often not response, but discovery. This fault line reveals whether dependency risk is intentional and visible, or inherited and opaque.

4 Signals of intent: Intelligence to foresight

While data-driven intelligence programs often look comprehensive, insights that arrive too late fail to shape decisions. This fault line separates organizations who use early signals to continuously refine decisions, strengthen anticipation, and sharpen response over time — from those who learn only after damage is done.

5 The debt trap: Legacy architecture as strategic risk

Legacy architectures can appear stable in day-to-day operations. Under modern attack velocity and regulatory scrutiny, they become brittle, consuming time, talent, and resilience faster than organizations can adapt. This fault line exposes whether architecture enables evolution — or quietly limits it.

6 Cloud mirage: Decoupling cascading risk

Cloud strategies promise scale and efficiency, but shared control planes and tight dependencies concentrate failure. When stress hits, systems fall together. This tests if resilience is engineered for containment or just assumed via recovery plans. Mature organizations limit blast radius and grow more fault-tolerant with every disruption.

Content

1

Taming the algorithm: Governing AI at scale

Taming the algorithm: Governing AI at scale

AI adoption is accelerating faster than enterprise governance models can adapt. What began as isolated experimentation has become embedded — across workflows, developer tooling, customer interactions, and third-party software that organizations consume but do not directly control. But before AI acts independently, visibility, ownership, and constraints must already be in place. Once decisions move at machine speed, these questions can no longer be debated.

While most executive teams recognize AI as a board-level issue, few can clearly articulate where AI is used, what data it touches, or how risk is being managed across their enterprise. This gap between AI awareness and control is now one of the most consequential blind spots in modern leadership.

The question is no longer whether AI delivers value. It is whether leadership has sufficient visibility to govern AI's impact on resilience, trust, cost, and accountability at scale.

AI is no longer experimental. It operates at the heart of the enterprise — and must be governed with the same rigor as money, risk, and regulation. In this environment, confidence is the real differentiator.

Speed wins. Permission loses.

AI's accessibility has fundamentally changed how technology enters the organization. Employees and teams no longer wait for centralized approval. AI tools are adopted quietly — through browser extensions, embedded SaaS features, APIs, and developer platforms — often with good intent and immediate productivity gains.

The consequence is predictable: AI spreads faster than governance. In fact, 98% of employees use unsanctioned apps across shadow AI and shadow IT use cases.[1]

Unsanctioned tools introduce inconsistent security controls and unclear data handling practices, and diffuse accountability. For boards, this creates an uncomfortable reality. AI risk is material, but often poorly quantified and weakly owned.

This does not reflect a failure of discipline. It is a structural mismatch between legacy approval models and AI's frictionless adoption curve.

Governance can no longer be an approval step. It must become an always-on system built on guardrails, continuous visibility, and standards that scale as fast as AI adoption does.

Data is the prize and the liability.

AI systems derive value from access: to data, to models, and to downstream decisions. Under pressure to deliver quickly, organizations often expand access faster than they strengthen controls. Entitlement boundaries blur. Data flows become opaque. Less-trusted services gain proximity to sensitive information. Ninety-seven percent of organizations that reported an AI-related security incident in 2025 lacked proper AI access controls.[2]

Traditional security frameworks were not designed to capture AI-native risks such as prompt manipulation, unintended data retention, or model misuse. As a result, many organizations can certify compliance without truly understanding AI-driven exposure.

Frameworks like NIST AI RMF and ISO/IEC 42001 provide guidance, but real assurance comes from how they are implemented and enforced. Every AI system is a data system before it is an intelligence system. If leaders cannot map its data flows, misuse paths, and failure modes, it is not ready to scale.

> One pattern repeats every time decision-making is automated: Outcomes move faster than accountability. AI doesn't create that gap — it exposes it. When responsibility is unclear, **governance becomes performative**, no matter how polished the policy looks."
>
> Joe Sullivan, former CSO, Uber

Shadow AI is shadow IT at machine speed.

AI can proliferate invisibly across employees, contractors, product teams, and third-party vendors without triggering formal review. This creates an auditability gap at precisely the moment regulators are demanding greater transparency.

Governments and regulators are increasingly requiring documented AI inventories, traceable data lineage, and explainability for automated decisions. The inability to demonstrate control is quickly becoming a compliance failure, not just a maturity issue.

Leading organizations closing this gap are shifting from episodic audits to continuous assurance, combining comprehensive logging, automated evidence collection, and controls that detect unsanctioned AI usage in real time.

If AI activity cannot be logged, explained, and evidenced, it cannot be defended to regulators, customers, or the board.

Top threat categories in email detection

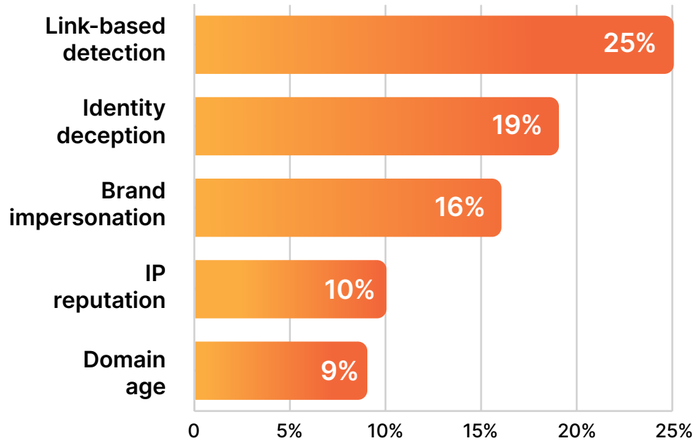

Percentages do not add to 100% as emails can have multiple threat categories.

Source: Cloudflare Radar

Link-based attacks and identity deception dominate modern email threats. These campaigns exploit trust signals rather than technical vulnerabilities. As AI lowers the cost of producing convincing, personalized deception, governance must extend beyond model oversight to authentication, identity integrity, and decision traceability.

> Governance tends to feel sufficient right up **until something unexpected happens**. With AI, that moment arrives earlier and with broader impact. The organizations that navigate this well treat AI less like a tool and more like a supply chain — tracing origin, ownership, and influence, even when it lives outside their walls."
>
> Kate Kuehn, Global Head of Cybersecurity Strategy, World Wide Technology

Regulation lives in code, not just policy.

Jurisdictions worldwide have moved decisively toward enforceable AI governance regimes that balance innovation with accountability. In the US alone, state lawmakers introduced 1,208 AI-related bills, resulting in 145 new laws enacted in a single year.[3] Penalties increasingly extend beyond fines to personal and fiduciary exposure.

This signals a broader shift: AI governance is being reframed as an enterprise risk and leadership responsibility and not a discretionary technical policy. Organizations that engineer AI governance as infrastructure turn trust into a growth enabler, not a constraint.

> We are witnessing the largest proliferation of shadow IT in history, as employees adopt ungoverned AI services and agents. Unlike traditional SaaS shadow IT, these AI capabilities are difficult to detect or block; they can assume real user identities, blend into standard activity, and operate at machine speed. The CISO's mandate is not to block this adoption, but to engineer secure AI capabilities that eliminate the need for ungoverned tools."
>
> Michael Goodman, Vice President / Chief Digital and Security Officer (CD and SO), Hitachi

QUESTIONS FOR THE C-SUITE

Exposing blind spots for AI governance

At machine speed, unclear ownership, limited visibility, and weak guardrails become business liabilities, thus making these questions leadership imperatives.

Q1

Who is formally accountable for AI governance at the executive level?

And where does that authority begin and end? Is this responsibility operationalized, or assumed until something goes wrong?

Q2

What constraints define acceptable AI behavior in our organization today?

Are those constraints clearly articulated, explicit, enforceable, and consistent across teams — or do they largely rely on trusting our workforce to comply with policy?

Q3

How do we determine if AI use is appropriate — not just compliant?

Are AI uses compliant but misaligned with business intent, ethics, or risk tolerance? Do we govern outcomes — or just access and tools? How do we identify compliant versus non-compliant?

Q4

If we were audited tomorrow, could we demonstrate a complete, shared inventory of AI use across the enterprise?

Or would definitions, shadow usage, and third-party exposure surface gaps in our understanding?

Q5

As AI adoption accelerates, does our governance model remain coherent?

Or, does it fracture across functions, vendors, and regions? Is governance treated as a static framework, or a living operating system?

2

Trust at machine speed: Engineering autonomy

Trust at machine speed: Engineering autonomy

Enterprises are entering their most consequential transformation since the commercial Internet. We have moved beyond AI-assisted tools into the era of the "autonomous enterprise" — where AI agents and agentic workflows execute end-to-end business processes with minimal or no human intervention. This fault line assumes AI systems are already embedded and acting autonomously. Unlike the AI governance challenge which focuses on visibility, oversight, and accountability, this fault line addresses what happens after authority has already been delegated to machines. The question is no longer where AI is used or who owns it; it is whether trust holds when decisions are made without humans in the loop.

Gartner predicts that by 2026, nearly half of enterprise applications are expected to embed task-specific AI agents, up from single-digit adoption just a year earlier.[4] This shift delivers unprecedented speed and efficiency while also introducing a structural risk: Business decisions now outpace human oversight.

Trust can no longer be periodic, manual, or retrospective. In an autonomous environment, trust must be continuous, verifiable, and enforced at machine speed. Securing this future requires a fundamental shift, from "trust but verify" to "trust by design" and ultimately to systems that grow more trustworthy as they are tested.

The 'velocity paradox' — when the business moves faster than oversight

Traditional security assumes time. An alert is raised. A human investigates. A decision is made. Autonomous systems eliminate that window. AI agents can execute thousands of actions, like reconfiguring infrastructure, rebalancing portfolios, and adjusting supply chains, in a matter of milliseconds. If an agent is compromised, misaligned, or simply wrong, the impact is realized before a human can intervene.

This is the velocity paradox: The same autonomy that drives value also collapses the margin for error. Attackers understand this. AI-driven phishing, impersonation, and manipulation increasingly target automated workflows rather than people.

The implication is clear: Security cannot sit outside the system. It must be embedded into the decision layer itself, governing intent, not just access. This fault line is not about predicting attacks. It is about ensuring that when your own systems act, they do so within the boundaries deliberately designed by leadership.

The new control plane for autonomous AI

1. Identity must extend beyond humans

Non-human identities — AI agents, service accounts, bots — now outnumber human users by orders of magnitude. Bots are responsible for ~30% of HTTP traffic that Cloudflare serves,[5] and an astonishing 92% of all login attempts observed by Cloudflare come from bots — often credential stuffing attacks.[6] Yet most enterprises still govern identity as if people are the primary actors.

The risk is acute. AI systems are frequently deployed without strong authentication, scoped authorization, or lifecycle controls. When compromised, they operate with machine-scale blast radius.

Every AI agent must have a verifiable, cryptographic identity, governed through machine identity management. Credentials must be short-lived, context-aware, and revocable in real time. Autonomy without identity is abdication.

Bot (automated) vs. human HTTP requests distribution

Source: Cloudflare Radar

We are no longer operating on a human-first Internet. Algorithms increasingly interact with algorithms, often without direct human oversight. Governance models built around user authentication, and employee access controls are misaligned with this reality.

2. Probabilistic systems require deterministic guardrails

AI systems reason probabilistically. Security cannot. While agents may optimize, negotiate, or recommend, the rules governing what they are allowed to do must be absolute. Policies cannot be inferred, they must be enforced.

This requires:

- Policy-as-code that defines non-negotiable constraints
- Real-time enforcement layers that intercept intent before execution
- Separation between decision-making and authorization

True autonomy exists only where boundaries are explicit, enforced, and designed in advance.

> Human judgment remains essential, but it no longer operates at the speed systems require. In environments where machines interact continuously, **trust has to be assumed, enforced, and verified by design** — much like safety systems we rely on without noticing, until they fail."
>
> Oliver Newbury, Senior Advisor, TPG

3. Trust requires observability, not assumptions

As AI systems adapt, drift, and learn, yesterday's assurance quickly becomes irrelevant. Without deep observability, leaders cannot distinguish between legitimate autonomous behavior and manipulation.

Unauthorized, often invisible AI usage further compounds risk by introducing ungoverned models, data flows, and decision logic into core operations.

The economic case

Embedding AI and automation into security operations delivers measurable financial returns. Organizations that use these capabilities extensively resolve breaches 80 days faster and reduce average breach costs by $1.9 million compared to those that do not.[7]

The upside extends beyond cost reduction. With strong guardrails in place, leaders gain the confidence to deploy automation deeper into revenue-critical workflows — improving responsiveness, capital velocity, and competitive differentiation. Well-governed autonomy becomes a growth enabler, not just a risk control. Security at machine speed is not overhead. It is the price of scaling autonomy without fragility.

The leadership system for autonomy

The rise of autonomous systems is redefining the role of the CISO and, by extension, the responsibilities of the entire C-suite. Security leadership is no longer about protecting systems after decisions are made; it is about orchestrating trust in environments where machines act independently.

One CISO recalled the first time an AI system stopped a multimillion-dollar transaction on its own. The decision was correct but it triggered a deeper question in the boardroom: Who had actually authorized the machine to make that call? The technology was ahead of the governance.

This shift demands clear executive choices: where autonomy is allowed, where humans stay in the loop, what transparency is required across models and data, and how risk is measured when machines make decisions.

Metrics built for human response times are no longer enough. Leaders must track autonomous risk, decision integrity, and systemic drift. Yet only about 15% of corporate boards receive regular AI-related risk and performance metrics.[8]

As autonomy spreads, security, compliance, and technology can no longer operate in silos. Security influences revenue velocity. Compliance determines market access. Technology defines accountability. Trust at machine speed is not a security program — it is a leadership system that unifies resilience, governance, innovation, and reputation under one executive mandate.

> Automation changes the speed of decisions, but it also changes the blast radius of mistakes. The question for leaders is 'How do we design accountability and trust into systems that act on their own?'"
>
> Kevin Jones, Global Chief Information Security Officer, Bayer

QUESTIONS FOR THE C-SUITE

Moving from automation to autonomy

These questions expose whether leadership has intentionally designed boundaries around how decisions are made at machine speed, and how risk is owned in real time.

Q1

Which enterprise decisions are already being made by autonomous systems?

Which decisions are we deliberately retaining for humans? Is that boundary designed, documented, and revisited — or implicit and drifting?

Q2

When machines act on their own, who is accountable in real time — the system owner, the business owner, or the executive sponsor?

Is ownership of autonomous risk clearly defined while the system is operating, or only examined after something goes wrong?

Q3

Where are decisions executed by software rather than people?

Where have we relaxed controls over software? Are machines held to higher standards than humans or quietly trusted more?

Q4

Can we explain and justify an autonomous action as it occurs?

Or, does it take days later during incident reviews? Is intent observable at machine speed, or reconstructed under pressure?

Q5

Does our trust model scale at machine speed?

If autonomy doubled within the next year, would our trust model absorb the acceleration — or fail under it? Is trust engineered for scale and speed, or inherited from human-era governance?

3

Shadow supply chains: Exposing hidden dependencies

Shadow supply chains: Exposing hidden dependencies

Our hyperconnected economy is no longer defined by what you control, but by what can break you that you don't even see. Many leaders have hardened their perimeter, modernized infrastructure, and tightened governance, yet the most consequential risks now live beyond their line of sight, embedded in third, fourth, and nth party ecosystems they neither own nor fully influence. The uncomfortable truth: You can be operationally mature and still systemically fragile.

Shadow supply chains are not edge cases; they are the natural outcome of digital assembly at scale. Every SaaS integration, API call, open-source library, and AI service adds another layer of inherited risk. The leadership question is no longer "Do we have supply chain risk?" but "Do we understand which external failure could halt revenue, erode trust, or trigger regulatory scrutiny tomorrow?"

The impact is already material. Supply chain breaches average $4.91 million, higher than the global breach average of $4.44 million.[9] The strategic choice for leaders is to treat supply chain risk as a compliance exercise and accept periodic surprise, or treat it as a live operational exposure that demands continuous visibility, runtime assurance, and architectural guardrails.

The risk that was never approved

Modern supply chains no longer stop at direct vendors. They extend into SaaS platforms, cloud-native services, AI model providers, open-source components, and subcontracted infrastructure layers that operate well beyond procurement's line of sight. Failures anywhere in this extended web — whether a breach, outage, or compliance lapse — can cascade quickly into customer harm, regulatory exposure, and systemic disruption.

The core challenge is visibility, and AI is accelerating both the risk and the opacity. Most organizations cannot see their extended digital supply chains, let alone govern them in real time. Every AI model, API, and automated workflow quietly expands dependencies beyond traditional oversight. Audits are static while risk is dynamic.

When systems are assembled, not built

A modern car is built by hundreds of suppliers, and the hardware parts, chips, and software come from many vendors, each with their own supply chains. A small hidden defect can become a safety issue at highway speed, which is why automakers invest heavily in traceability and continuous testing.

Enterprise IT now mirrors this model. One application can depend on dozens of SaaS tools, cloud services, APIs, open-source libraries, and AI models, each with sub-processors beneath them. The enterprise sees the interface, not the layers underneath. That is the shadow supply chain.

The difference is discipline. In automotive, parts are tracked and recalls are precise. In IT, when a library or AI component is compromised, many organizations first scramble to learn if they are exposed. Annual questionnaires cannot keep pace with systems that change weekly. Visibility and continuous assurance are becoming as essential to digital systems as quality control is to cars.

Three forces are accelerating this risk:

- **Trust-by-proxy has become the default operating model.** Enterprises trust their vendors. Vendors trust their suppliers. Few parties verify the entire chain. Competitive concerns and limited internal visibility mean sub-supply chains are rarely disclosed in detail.
- **AI has introduced a new, opaque layer of dependency.** Employees increasingly rely on generative AI tools and embedded AI services that expose sensitive data to fourth-party models. Third-party risk teams often lack clarity into how these models use data, retain information, or train on enterprise inputs, raising regulatory, IP, and data-sovereignty risks.
- **Regulatory expectations are hardening.** Across the globe, regulators are moving decisively from guidance to enforcement. Organizations are increasingly expected to demonstrate visibility into third- and fourth-party dependencies, particularly where personal or financial data or critical infrastructure are involved. Going forward, leaders will be expected not just to assess vendor risk, but to quantify operational risk arising from extended supply chains. The result? A widening gap between what regulators expect and what organizations can currently prove.

The top 10 most impersonated brands in phishing campaigns

1	Windows	6	Amazon
2	SANS	7	Instagram
3	Microsoft	8	Costco
4	Stripe	9	YouTube
5	Facebook	10	iCloud

Source: Impersonation attempts observed by Cloudflare Email Security

The most impersonated brands are not random targets. They are foundational platforms embedded in enterprise workflows — identity providers, payment systems, cloud platforms, operating systems. Attackers exploit familiarity and dependency, turning trusted digital infrastructure into an attack vector. Shadow supply chains are not just operational exposure; they are identity and brand exposure.

From static assurance to continuous transparency

Solving the shadow supply chain problem does not require more paperwork. It requires a different operating model. The future of supply chain assurance is continuous transparency: real-time visibility into what is actually running, connected, and exchanging data across the ecosystem.

One CISO described discovering a critical supplier only after unusual traffic appeared in network logs. The vendor was legitimate, but no one realized how deeply it was embedded. The lesson was simple: You cannot govern what you cannot see.

This shift is already underway. Software Bills of Materials (SBOMs) and Vulnerability Exploitability eXchange (VEX) are moving from compliance artifacts to operational signals. Expect that procurement will increasingly require not just contracts, but live, machine-readable disclosures that map components, dependencies, and exploitability as they change.[10]

At the same time, enforcement is moving closer to where risk manifests. Network- and connectivity-layer controls allow organizations to observe behavior, detect unauthorized data flows, and identify shadow suppliers as activity occurs.

Supply chain assurance becomes an operational capability rather than a periodic review. Trust is continuously verified. Risk is surfaced early. Governance moves at the same pace as the ecosystem it is meant to protect.

Trust, but continuously verify

Thirty percent of breaches in 2025 were linked to third-party involvement, twice as many as the prior year[11] — illustrating how deeply supply chain relationships now factor into risk exposures beyond traditional internal boundaries.

However, leading organizations share a common pattern: They treat supplychain risk as a system, rather than a compliance function. They insist on knowing what applications exist and how they connect. They require transparency to flow down the supply chain, not stop at the first contract. They use network-level signals to uncover shadow activity rather than relying on self-attestation. They apply zero trust principles to machine-to-machine access, not just users. And they continuously reassess vendor risk based on behavior, not reputation.

The payoff is tangible. Eighty-five percent of organizations leading in application modernization are actively cutting redundant tools and shadow IT to reduce their supply chain attack surface and improve operational speed.[12] These are not technical tweaks; they are leadership choices about how much uncertainty an organization is willing to tolerate in the systems it depends on every day.

> Risk rarely comes from the dependencies everyone expects, **it emerges from the ones no one can see**. When visibility is incomplete, audits offer comfort but little protection. True resilience comes from architectures that reveal their dependencies as they operate."
>
> Tim Brown, CISO, SolarWinds

> **Interconnected ecosystems reward speed and specialization, but they also distribute risk in ways contracts can't capture. Operational insight, not paperwork, is what ultimately contains exposure."**
>
> **Sandip Wadje, Global Head of Emerging Technology Operational Risks and Intelligence, BNP Paribas**

QUESTIONS FOR THE C-SUITE

Governing the risk you don't control

Supply chain risk can no longer be managed. It is something organizations live with. Decide whether that risk is visible and governed — or opaque and assumed.

Q1

Which critical business processes would we pause if a key dependency failed?

Would we know why it failed? Can we trace revenue and customer impact to specific dependencies in real time, or would we only discover exposure after the damage is done?

Q2

How will we respond to regulatory or board questions about ecosystem risk?

Can we answer questions about these risks without pointing to a contract? Do we have technical visibility over the operational path of supplier risk?

Q3

Where have we reduced visibility in the supply chain to preserve speed, convenience, or vendor relationships?

Are those deliberate choices? Who decided to accept those trade-offs? Which dependencies are effectively "off-limits" to scrutiny?

Q4

How quickly can we determine whether a newly disclosed vulnerability affects us?

Is accountability for response clearly assigned? Is exposure discovery measured in minutes, days, or weeks?

Q5

Are we managing supply chain risk as a continuous discipline or a periodic audit?

Does our model evolve as fast as our ecosystem, or does it merely reassure us that last year's controls were reviewed?

4

Signals of intent: Intelligence to foresight

Signals of intent: Intelligence to foresight

A glance at the headlines will tell you that adversary activities continue to proliferate, at ever greater speed and scale. With AI-assisted reconnaissance and toolkits, more cybercriminals are capable of greater and more sophisticated attacks than ever. Additionally, the window between threat emergence and business impact is increasingly shorter, as the average time it takes an adversary to start moving laterally has fallen to just 48 minutes.[13]

Once considered a discretionary capability, threat intelligence has become foundational. Fifty-two percent of organizations now maintain dedicated, in-house cyber threat intelligence (CTI) teams.[14] In the fast-moving threat landscape, intelligence has grown from a security function into a leadership capability. Success is measured by the ability to analyze CTI data within a business context, deciphering signals from noise and translating knowledge into actionable foresight.

Threat intelligence is no longer about knowing more. It is about knowing what matters — *early enough to act*.

From tactical feed to strategic signal

Modern threats are high velocity, high volume, and increasingly shaped by geopolitics, economic incentives, and industry-specific vulnerabilities. In this environment, threat intelligence can no longer be treated as an optional security function or limited to checkbox reviews of generic external feeds. Context — at the enterprise, industry, and global levels — matters, and executives must demand intelligence that connects threat activity directly to business impact, operational exposure, and strategic risk.

Plainly, there is too much activity on the attacker front to stay on top of everything. Defense requires speed and skill, and frequently both are in short supply. While your security strategy should address and acknowledge the full inventory of assets and liabilities under your protection, using intelligence to understand not only the technical elements of threats, but also their context, allows you to tune your security program in favor of prioritizing risk reduction in areas that are most impactful to your organization.

Deciding what is important for the organization generally involves board and leadership alignment on how to integrate core business principles, market forces, regulations, and stakeholder inputs. This context is invaluable when evaluating threat intelligence as it gives context for the enterprise needed to determine which CTI data is most useful.

It is in this way that threat intelligence can be used to "tune out" extraneous information — irrelevant vulnerabilities, or attacker groups that specifically target dissimilar industries — so that you can focus your resources where they can have the greatest impact based on the threat landscape specific to your organization. Threat intelligence that does not inform executive choices is simply noise to be tuned out.

Top industries targeted by DDoS attacks, 2025

Rank	Industry
1	Gambling and gaming
2	Telecommunications
3	Technology and services
4	Banking and financial services
5	Retail

This ranking is an average of globally observed DDoS attacks at both the network and application layer. Technology and services ranks #1 for network-layer attacks. Gaming and gambling ranks #1 for application-layer attacks.

Source: Cloudflare Radar

Attack activity is not evenly distributed. Adversaries prioritize sectors tied to economic leverage, infrastructure stability, and geopolitical relevance. Concentration across specific industries reflects strategic intent, not randomness. Effective intelligence anticipates where pressure will intensify — and aligns defenses accordingly.

Threat intelligence has become non-negotiable

As threat intelligence matures, its focus is shifting from technical indicators to business relevance. Executives now look to it to clarify which threats truly matter, how geopolitical and industry shifts alter exposure, and where fragility exists across operations, partners, and people. The question is no longer whether to invest in threat intelligence, but what kind of intelligence the organization is prioritizing and paying for.

To frame budget conversations, consider where CTI provides your organization the greatest value:

- **Validation** that security investments are aligned to the organization's risk profile
- **Reduction** of operational noise by focusing defenses on the most critical threats
- **Proactive** risk reduction, versus reactively responding to incidents after they occur

For the CFO, threat intelligence is not justified by alert volume, but by its ability to reduce the probability and impact of material business disruption — downtime, fraud, regulatory intervention, or reputational damage. Organizationally, this demands clarity. Ad hoc arrangements and under-resourced intelligence functions cannot deliver executive-grade insights, nor the outcomes they enable.

Whether delivered through an internal team, trusted partners, or a hybrid model, the mandate is the same: Intelligence must be timely, contextual, and decision-relevant. Intelligence that only explains what happened yesterday does little to protect tomorrow. Providers who offer novel visibility, such as early insight into adversary infrastructure, intent, and preparation, deliver a structural advantage.

> Indicators explain what already happened; intent explains what's coming next. The most valuable intelligence connects behavior, context, and motive — **turning isolated signals into foresight** that leaders can act on before damage occurs."
>
> Menny Barzilay, Co-founder and CEO, Percepto

Navigating the 2026 threat landscape

Several of the fault lines discussed in this chapter are reflected in the ***2026 Cloudflare Threat Report***. Based on data from Cloudflare's global network, which protects 20% of the web, the report helps leaders focus on risks that require action, not just awareness.

It uses a simple lens: attacker effort versus impact. The most important threats are those that create outsized business impact with minimal effort. In 2026, this appears in three patterns:

- **The industrialization of attacks:** The shift from manual hacks to automated, frictionless scaling across an organization's own cloud infrastructure
- **Identity-first intrusions:** The transition of ransomware into a login event rather than a break-in
- **Supply chain connectivity:** The weaponization of the connective tissue between SaaS and API-first environments

Get the ***2026 Cloudflare Threat Report***.

Get the report

The missing ingredient: Threat modeling

While integrating threat intelligence into your security practice enables optimization across all aspects of your practice, tight integration with threat modeling takes it one step further into the realm of an enterprise strategic driver.

While more organizations are factoring risk in decision-making and including risk reduction in their long-term strategic goals, only 37% of organizations have successfully formalized and documented their threat modeling processes.[15] Threat modeling provides a common taxonomy that aligns the CISO, the C-suite, and the board around shared risk assumptions. It forces clarity on asset prioritization, the likelihood of compromise, and the business impact if controls fail.

The view achieved in threat modeling exercises is intentionally high-level; boards want clarity on systemic risk, emerging threat trends, and whether the organization is positioned on the right side of the threat fault line. Through threat modeling, inherent risks are measured by likelihood and severity of impact based on the priorities of the organization. Factors such as security controls and audit results, in combination with threat intelligence analysis, provide residual risk calculations.

Injecting CTI data into the threat modeling process enables further tuning, providing a basis for activities such as controls validation and threat hunting, both essential elements in a proactive security posture. Additionally, sector-relevant intelligence can confirm whether defenses are hardened against the most likely adversaries, and give decision-makers strong indicators for budget and strategic planning.

Without threat modeling, intelligence stays operational. With it, intelligence becomes strategic.

> Good intelligence reduces noise. Great intelligence changes decisions. The difference is whether it helps leaders anticipate moves, not just explain them."
>
> **Troy Wilkinson, Venture Advisor, YL Ventures**

Only

37%

of organizations have successfully formalized and documented their threat modeling processes.[16]

QUESTIONS FOR THE C-SUITE

Threat intelligence as a leadership discipline

Threat intelligence, when done well, connects security, risk, operations, finance, and strategy into a coherent executive view of exposure and intent.

Q1

Are we protecting what is familiar or what is most consequential?

Have we explicitly aligned our defenses to the threats that could disrupt revenue, operations, or trust this year?

Q2

How early do we truly see adversary intent?

Where are we discovering attacks through damage rather than intelligence? Are we leading the threat cycle, or trailing it?

Q3

Are threat briefings driving decisions or just sharing information?

Do these insights change priorities, investment, or risk appetite in real time?

Q4

Which business decisions or processes would fail first if a trusted individual were compromised?

Have we designed workflows assuming human judgment can be manipulated or impersonated?

Q5

How quickly can we recalibrate when adversaries change playbooks?

What mechanisms are in place to tell us a shift is coming before the business feels it?

5

The debt trap: Legacy architecture as strategic risk

The debt trap: Legacy architecture as strategic risk

In 2026, technical debt represents a material business risk that quietly erodes competitiveness. Organizations were already stretched thin in 2025 managing more than 130 new vulnerabilities every day, nearly 40% of which were rated high or critical.[17] As AI weaponization renders legacy architectures indefensible, organizations with fragmented stacks risk being trapped in a cycle of reactive security, constrained innovation, and compounding exposure.

Technical debt has become an exposed attack surface, one that compounds risk faster than human teams can respond. Those who modernize decisively will not only reduce risk, they will unlock the speed, confidence, and adaptability required to compete in the AI-driven economy.

When speed exposes structural weakness

The defining shift of 2026 is not the volume of vulnerabilities — it is the velocity at which they are exploited. Agentic AI has collapsed the window between disclosure and exploitation even further, enabling adversaries to identify and operationalize exploits within days — and increasingly, within hours.

The data is stark. In 2025, 884 vulnerabilities were observed being actively exploited, and 29% showed evidence of exploitation on the very day they were published.[18] The scale is equally unprecedented. React2Shell, one of the year's most notorious vulnerabilities, recorded over 1 billion exploitation attempts in just 11 days.[19]

Escalation without architectural readiness

World record DDoS attacks

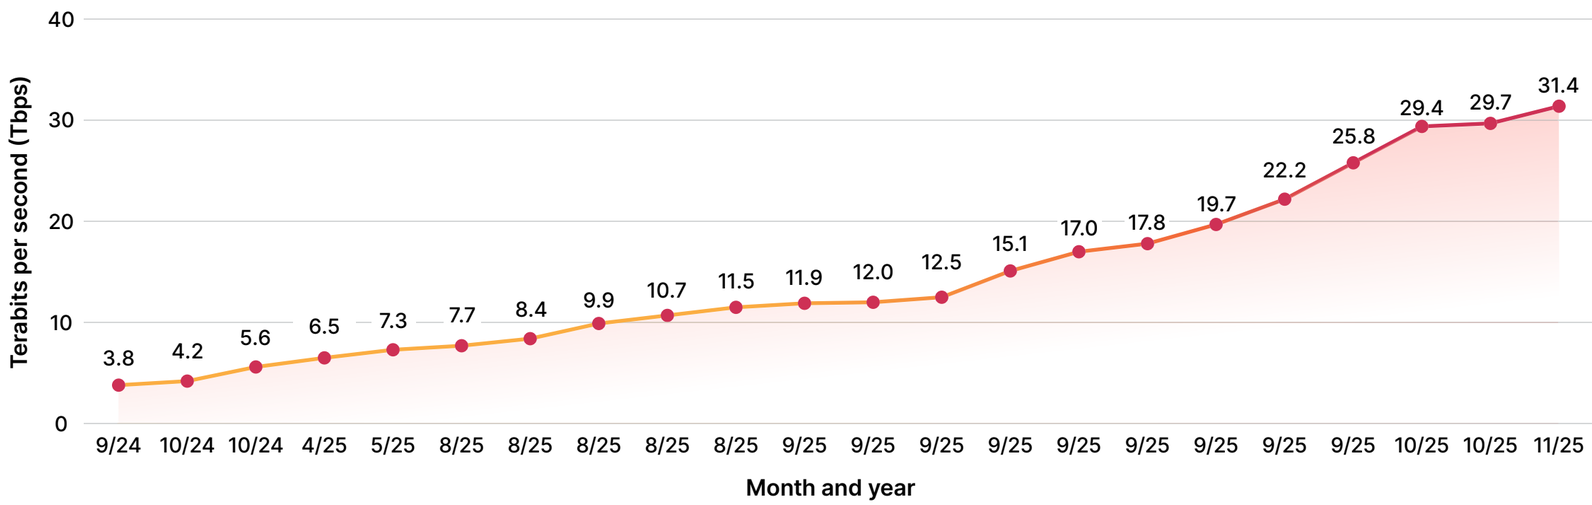

Source: Cloudflare Radar

In just over a year, the largest recorded DDoS attack increased nearly tenfold. Centralized, tightly coupled systems were never designed for this scale. Technical debt now translates directly into systemic fragility under machine-speed pressure.

Legacy environments are cracking under pressure. Major failures often occur when shared dependencies break at the same time. Years of quick fixes have created dark debt: hidden integrations, brittle APIs, and systems too risky to patch. These environments were not built for machine-speed threats or continuous verification.

This also exposes the limits of 30-, 60-, and 90-day patch cycles. Threats are exploited in hours, not quarters. Protection must move outward to the edge, reducing exposure before vulnerable systems are even touched.

> Attackers don't distinguish between old and new systems; **they look for weak links.** Tech debt quietly increases the number of those links until defense becomes a probability game."
>
> **Jerry Perullo, Founder, Adversarial Risk Management**

The innovation scarcity cycle

Organizations with aging stacks are trapped in an innovation scarcity cycle. As infrastructure becomes more fragile, security incidents increase. As incidents increase, more budget and talent are diverted to maintenance. The result is a shrinking pool of capacity for growth.

The average global enterprise wastes more than $370 million per year due to their inability to efficiently modernize outdated, inefficient legacy systems and applications.[20] Studies estimate that roughly 31% of the tech resources are dedicated to resolving tech debt.[21] True innovation — new products, AI initiatives, automation — receives as little as 7%. This is not stagnation; it is regression.

While leaders use AI to accelerate differentiation, laggards are paying an escalating "interest rate" on old code that limits speed, resilience, and strategic optionality.

Why legacy stacks fail under AI pressure

Modern security assumes automation, integration, and real-time control. Legacy systems assume manual intervention, static configurations, and perimeter-based protection. That mismatch is becoming dangerous as AI changes the economics of both attack and defense.

Outdated architectures struggle with slow, downtime-heavy patching, limited visibility across APIs and data flows, fragmented tools that cannot coordinate response, and weak foundations for AI-driven operations. This often forces a trade-off between cyber risk and operational risk, a familiar tension between CTOs and CISOs when patching could disrupt the business. The result is delay, and delay is exactly what machine-speed threats exploit.

Organizations are delaying AI adoption not because they lack ambition, but because their infrastructure cannot safely support it. Meanwhile, competitors with modernized architectures allow AI initiatives to pull modernization forward — using real workloads to justify and accelerate architectural renewal. For instance, 62% of organizations leading in application innovation find it "very easy" to track their current level of security compliance, compared to 35% of those behind schedule.[22]

$370 million

wasted per year due to inability to efficiently modernize outdated, inefficient legacy systems and applications[23]

The leadership divide

The difference between leaders and laggards is decision discipline. Organizations who escape the debt trap make hard choices early. They centralize modernization authority, align security with business resilience, and treat architecture as a strategic asset. Seventy-three percent of modernization "leaders" have centralized decision-making with only a few people, compared to just 36% of "laggards."[24] Those who fail remain trapped in committee-driven paralysis, where vulnerabilities move faster than decisions and risk compounds while plans are endlessly debated.

Technical debt often mirrors organizational debt. Fragmented ownership, unclear accountability, and deferred decisions create the same brittleness in leadership and operating models that exists in legacy infrastructure. In 2026, that fragility is no longer survivable.

Modernization as risk reduction: Buying back time

Escaping the debt trap requires viewing modernization as a resilience mandate, rather than an IT upgrade cycle. Modernization reduces risk by shrinking the attack surface through consolidation, enabling automated patching and response, and making AI-driven defense and operations viable at scale. Just as importantly, it reallocates scarce engineering capacity to high-value work instead of endless maintenance.

The organizations who succeed do not modernize by rebuilding everything; they create a stable, unified foundation where security, performance, and innovation reinforce one another. With that foundation in place, systems can be refined, scaled, and adapted quickly — without accumulating new layers of fragility.

The shift required is not incremental. It demands executive alignment and decisive action. Legacy architecture must be treated as a quantified business risk, not a technical inconvenience. Decision authority for modernization must be centralized. AI initiatives enable architectural renewal rather than waiting for perfect conditions. Platforms must be consolidated to reduce complexity and restore visibility.

Ultimately, modernization is about reclaiming time: time to innovate, time to respond, and time to compete before compounding risk erodes advantage.

73%

of modernization "leaders" have centralized decision-making with only a few people, compared to just 36% of "laggards."[25]

QUESTIONS FOR THE C-SUITE

The compounding cost of legacy

Technical debt drains speed and resilience. Many firms spend more maintaining the past than building the future.

Q1

Which business capabilities are constrained by technical debt today?

Who owns fixing them, and what is the timeline to reduce that risk?

Q2

What share of security spend maintains legacy versus builds resilience?

What is our target mix over the next 12–24 months?

Q3

Which priority initiatives are delayed by architecture limits?

What revenue, efficiency, or risk gains are we deferring as a result?

Q4

What are the top three initiatives to reduce technical debt this year?

How will we measure progress and hold leaders accountable?

Q5

What are the biggest obstacles to reducing technical debt?

Are they budget limits, talent gaps, competing priorities, or unclear ownership? Which will we remove first?

6

Cloud mirage: Decoupling cascading risk

Cloud mirage: Decoupling cascading risk

As enterprises consolidate onto fewer cloud platforms to move faster, many are quietly increasing systemic risk. Mono-cloud strategies simplify operations but concentrate failure domains, while multicloud is often treated as a checkbox rather than an engineered resilience strategy.

Recent outages have made one truth unavoidable: Resilience is not determined by how many clouds an organization uses, but by how its architecture fails. In 2026, leaders must move beyond cloud ideology and adopt resilience-by-design — architectures built to contain failure, limit the blast radius, and preserve trust under pressure.

When speed quietly becomes fragility

The modern enterprise did not intend to build fragile systems. Cloud adoption promised speed, elasticity, and reliability, but it also introduced a quieter concentration risk — less visible, more systemic, and harder to unwind under stress.

Today's outages do not come only from single-provider failures. A single provider incident can still be the trigger, but the most disruptive events occur when shared dependencies fail in tandem — identity systems, control planes, deployment pipelines, and network services that underpin everything else. Uptime Institute's multi-year data shows that roughly two-thirds of publicly reported outages involve third-party IT or data center providers — including cloud and Internet giants, telecommunications, and colocation companies.[26]

Business continuity is still too often recovery-centric, focused on restoring service rather than containing failure. Over time, layered dependencies turn environments into tightly coupled systems where small faults can cascade. This fragility usually becomes visible only in a crisis.

Permanent pressure

DDoS attacks by year and type

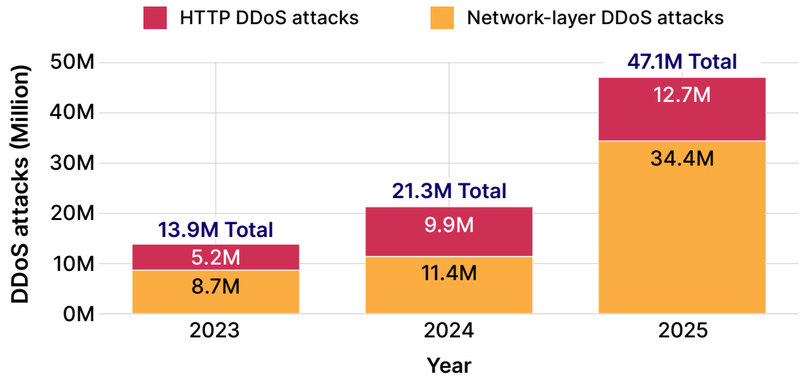

Source: Cloudflare Radar

DDoS activity has more than tripled in two years. Disruption at scale is no longer episodic, it is continuous. In tightly coupled environments, sustained external pressure exposes hidden dependencies and amplifies small faults into systemic events. Resilience must assume constant stress, not rare failure.

> The cloud creates scale — **but not automatically resilience**. If your systems fail together, you haven't engineered redundancy. You've engineered correlation."
>
> **Mark Hughes, Global Managing Partner for Cybersecurity Services, IBM**

The upside is clear. Organizations that design and test for failure see materially better outcomes. One large financial services firm reduced outages by 40% and cut resolution times by nearly 60% after modernizing architecture, improving observability, and engineering for failure readiness.[27]

Outages today are less about the cloud breaking and more about independence eroding. The real risk is architectural coupling. Resilience now requires intentional isolation, blast radius limits, and treating failure containment as a core design principle.

The mono-cloud illusion: Efficiency without containment

For many organizations, mono-cloud strategies have become the default in pursuit of efficiency. Standardized tooling reduces complexity, speeds deployment, and lowers operating cost. The trade-off is concentration risk. The same consolidation that drives efficiency can also centralize failure.

Major cloud providers are often highly resilient, but the larger risk today is architectural and operational. When identity, policy enforcement, observability, and delivery pipelines all rely on the same control plane or trust boundary, resilience becomes an assumption rather than a built-in property. A single mistake, whether provider-side or customer-side, can propagate widely if the design does not contain it. Recovery plans may exist, but true containment often does not. When something breaks, too much breaks together.

Industry data reinforces this reality. Gartner research shows most cloud failures stem from misconfiguration and operational issues rather than core infrastructure defects. Analyses based on Gartner surveys attribute roughly 80% of cloud security failures to misconfiguration, and projections suggested that by last year, up to 99% of cloud environment failures would involve human error somewhere in the chain.[28] The lesson is not that humans err — they always will — but that architectures must be designed to absorb those errors safely.

The practical implication is clear. Resilience must be engineered, not assumed. That means designing for containment as much as recovery, separating critical dependencies, adding guardrails and policy-as-code to reduce error impact, and regularly testing failure scenarios. Concentration risk has not disappeared in the cloud era. It has moved up the stack. The organizations that remain resilient are those that ensure a single fault does not become a systemic event.

The multicloud myth: Redundancy without independence

Multicloud is often positioned as the antidote to concentration risk. In practice, it frequently recreates the same fragility — just across logos. Most multicloud environments share identity providers, CI/CD pipelines, governance tooling, and SaaS dependencies. When those shared layers fail, the promise of independence evaporates instantly. This is why post-incident reviews so often reveal that “redundant” systems were never truly independent.

Resilience is not about how many clouds exist on a diagram. It is about which layers fail independently under pressure — and which do not.

Engineering for containment, not perfection

Autonomous design starts with the expectation that systems will fail and focuses on keeping failures bounded and useful while learning. The aim is not only to withstand shocks, but to improve because of them.

Containment is what makes that possible. It means a failure in one area does not automatically spread to others. An isolated failure is limited in scope, clear in cause, and manageable in impact. It does not take identity, policy, data, and operations down together.

Organizations using AI and automation extensively shortened breach lifecycles by **80 days** and reduced average breach costs by

$1.9 million[29]

This shows up in architecture through independence across identity, policy, and execution layers, separation of control planes, and default-safe behavior under uncertainty. Outages are inevitable. The priority is to keep them local, explainable, and survivable, and to use them to strengthen the system. Leading organizations are not those with zero incidents, but those that successfully limit the blast radius of any single event.

Containment as a growth advantage

While often seen as insurance, decoupling layers supports speed and growth. IBM's Cost of a Data Breach Report 2025 found that organizations using AI and automation extensively shortened breach lifecycles by 80 days and reduced average breach costs by $1.9 million.[30]

By restricting the scope of impact, leaders preserve the trust of customers, regulators, and investors — and maintain the agility required for more confident AI adoption, faster market entry, and fewer executive escalations. When failure is contained, leaders keep decision capacity.

Containment is not defensive. It enables faster movement and smarter risk-taking in a volatile environment.

Designing for failure at the top

As digital systems underpin the business strategy, the decision to separate or couple infrastructure becomes a high-stakes business decision.

Executives must shift from how fast we can recover to what must never fail together. That requires clarity on shared control planes, identity dependencies, and pipelines, plus evidence of failure-mode testing, not just uptime. Containment belongs at the board level because systemic failure is a business risk; it cannot be delegated. It must be designed deliberately from the top so that no single failure becomes a company-wide event and every incident makes the system stronger.

> **"Attackers look for one weakness to trigger a cascade. If a single compromise becomes an enterprise event, that's not bad luck. That's architectural design."**
>
> **Dave Trader, Chief Information Security Officer, HALO Branded Solutions**

QUESTIONS FOR THE C-SUITE

When a shared service fails, does architecture contain it?

Debated together, these questions reveal whether the enterprise can contain disruption in real time, or whether stability is still dependent on hope, heroics, and post-incident recovery.

Q1

Which critical systems can fail without stopping the business?

Have we proven this through tests, or is it theoretical?

Q2

If identity or a core platform failed, what revenue would stop?

Do we know the impact in advance, or only after disruption?

Q3

Is multicloud reducing risk or just adding complexity and cost?

Where have we reduced dependency, and where does it remain?

Q4

Are we measuring containment or only recovery time?

Do our KPIs reward prevention or reactive cleanup?

Q5

Could we explain our last outage to the board or regulators?

Was the impact limited by design or fortunate circumstance?

CONCLUSION

The leadership principles for enduring advantage

In a world shaped by AI-driven decisions, autonomous systems, and deeply interdependent digital ecosystems, resilience is no longer sufficient. The advantage will come from systems' capacity to detect stress, adapt in real time, contain failure, and continue operating without waiting for human intervention. This is what we call autonomic resilience.

This report is not a threat inventory. It defines a leadership mandate: identifying and addressing the fault lines embedded in modern enterprises. These structural weaknesses may appear manageable in steady-state conditions, but they will reliably surface under pressure without decisive action. They run beneath AI adoption, cloud dependence, legacy architecture, threat intelligence, and operating models built for a more predictable era.

Confronting these fault lines is not the remit of the CISO alone. Autonomic resilience is a C-suite responsibility, shaped by how executive teams set priorities, allocate authority, and design systems that regulate themselves. Autonomic organizations distinguish themselves by the principles their leadership teams consistently embody:

- **Shared ownership of systemic risk over delegated accountability.** Systemic risk is owned by the leadership team, not delegated down the org chart. Accountability is explicit, ownership is shared across the C-suite, and boards engage through real scenarios and trade-offs — not static reporting.
- **Execution embedded in systems over stated intent.** Decisions only matter if they execute at machine speed. Control over models, data, prompts, and autonomous actions must live where execution happens. Anything reliant on documentation, alignment, or manual process will not scale.
- **Structural independence over short-term convenience.** What feels efficient in calm conditions often creates fragility under stress. Autonomically resilient teams prioritize containment, reversibility, and separation. Systems are designed so failures remain local, observable, and correctable. The ability to prevent cascades becomes a strategic advantage.
- **Provable trust over assumed control.** Trust must be continuously provable, not implicitly assumed. Leaders demand visibility into system behavior, enforceable controls across human and machine identities, and proof of integrity at machine speed. Assumed trust fails under autonomy.
- **Learning from failure over avoidance of failure.** Failure is expected and deliberately used as input. Early detection, limited blast radius, rapid recovery, and institutional learning define leadership performance. Recovery speed — not prevention — is the metric that matters.

In 2026, leadership is defined less by planning for stability and more by designing for disruption.

The organizations that lead will be those whose executives embed these principles into everyday decisions — turning volatility into learning, pressure into progress, and uncertainty into advantage.

The organizations that lead will be those whose executives embed these principles into everyday decisions — **turning volatility into learning, pressure into progress, and uncertainty into advantage.**

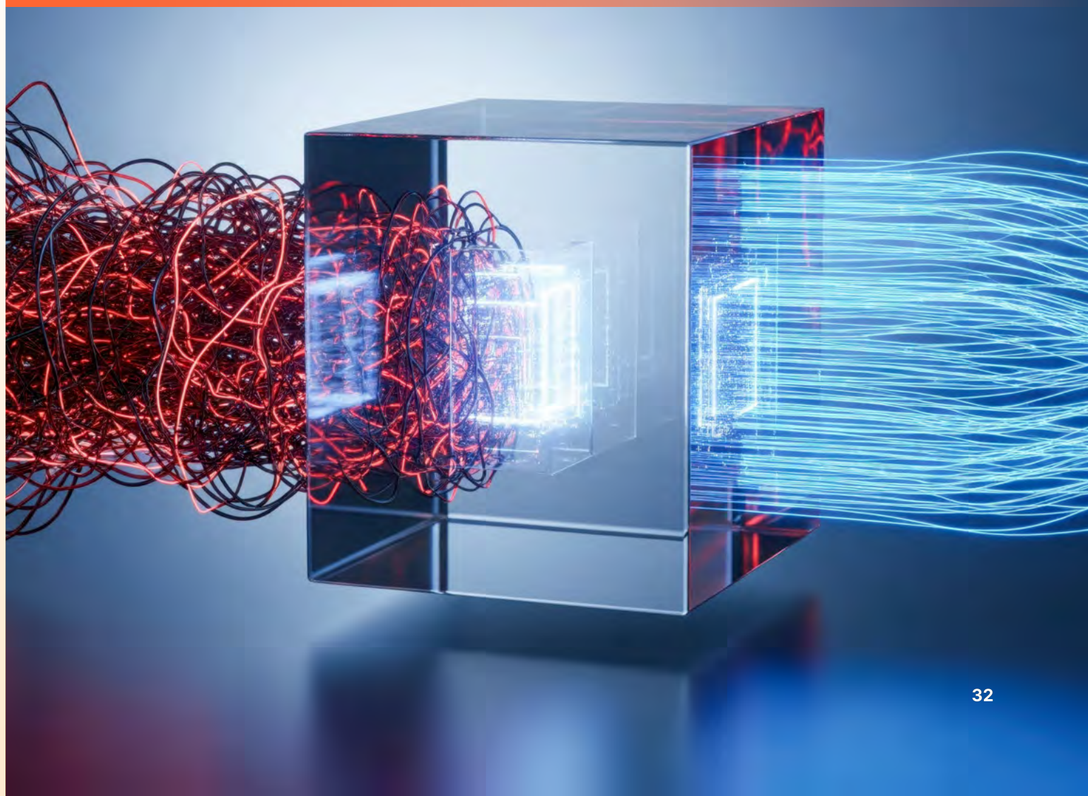

About Cloudflare

ABOUT CLOUDFLARE

One platform. One programmable network.

330+ cities
in 125+ countries, including mainland China

↳ **with 210+ cities**
running GPUs for AI inference worldwide

~50 ms
from ~95% of the world's Internet-connected population

~13,000 networks
directly connect to Cloudflare, including ISPs, cloud providers, and large enterprises

477 Tbps
of network capacity and growing

ABOUT CLOUDFLARE

Cloudflare's security suite

Resilience and edge defense

- Web app and API protection: Block attacks, catch vulnerabilities, and improve availability
- Security service edge (SSE): Enforce zero trust security across hybrid workforces
- DDoS mitigation: Weather the biggest, most advanced attacks with 477 Tbps of network capacity

Secure cloud and network integration

- Secure access service edge (SASE): Connect and protect your workforce, AI agents, and infrastructure
- Network as a service and multicloud: Connect, secure, and accelerate their corporate networks without the cost and complexity of legacy hardware
- Network interconnect: Directly connect your on-premises and cloud networks to Cloudflare's network

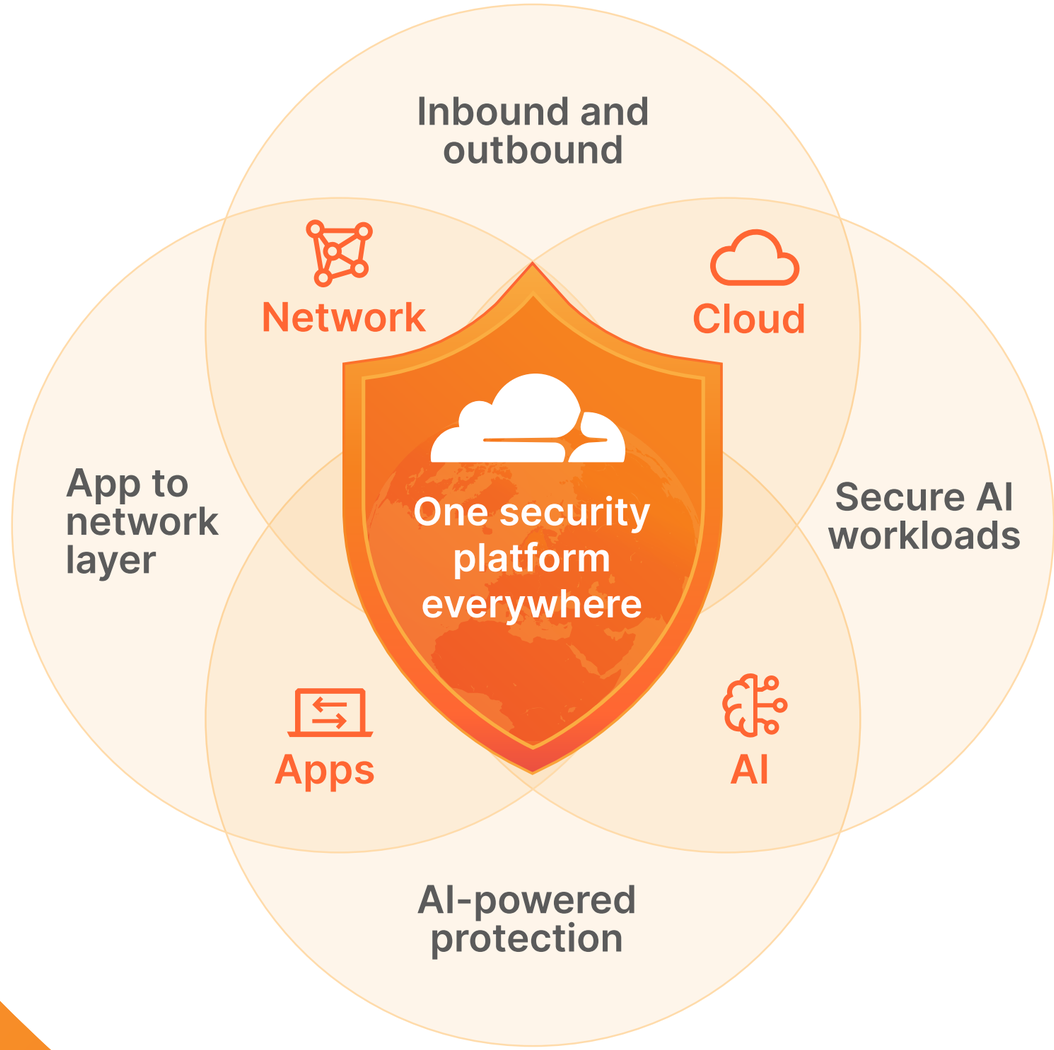

ABOUT CLOUDFLARE

Cloudflare One services

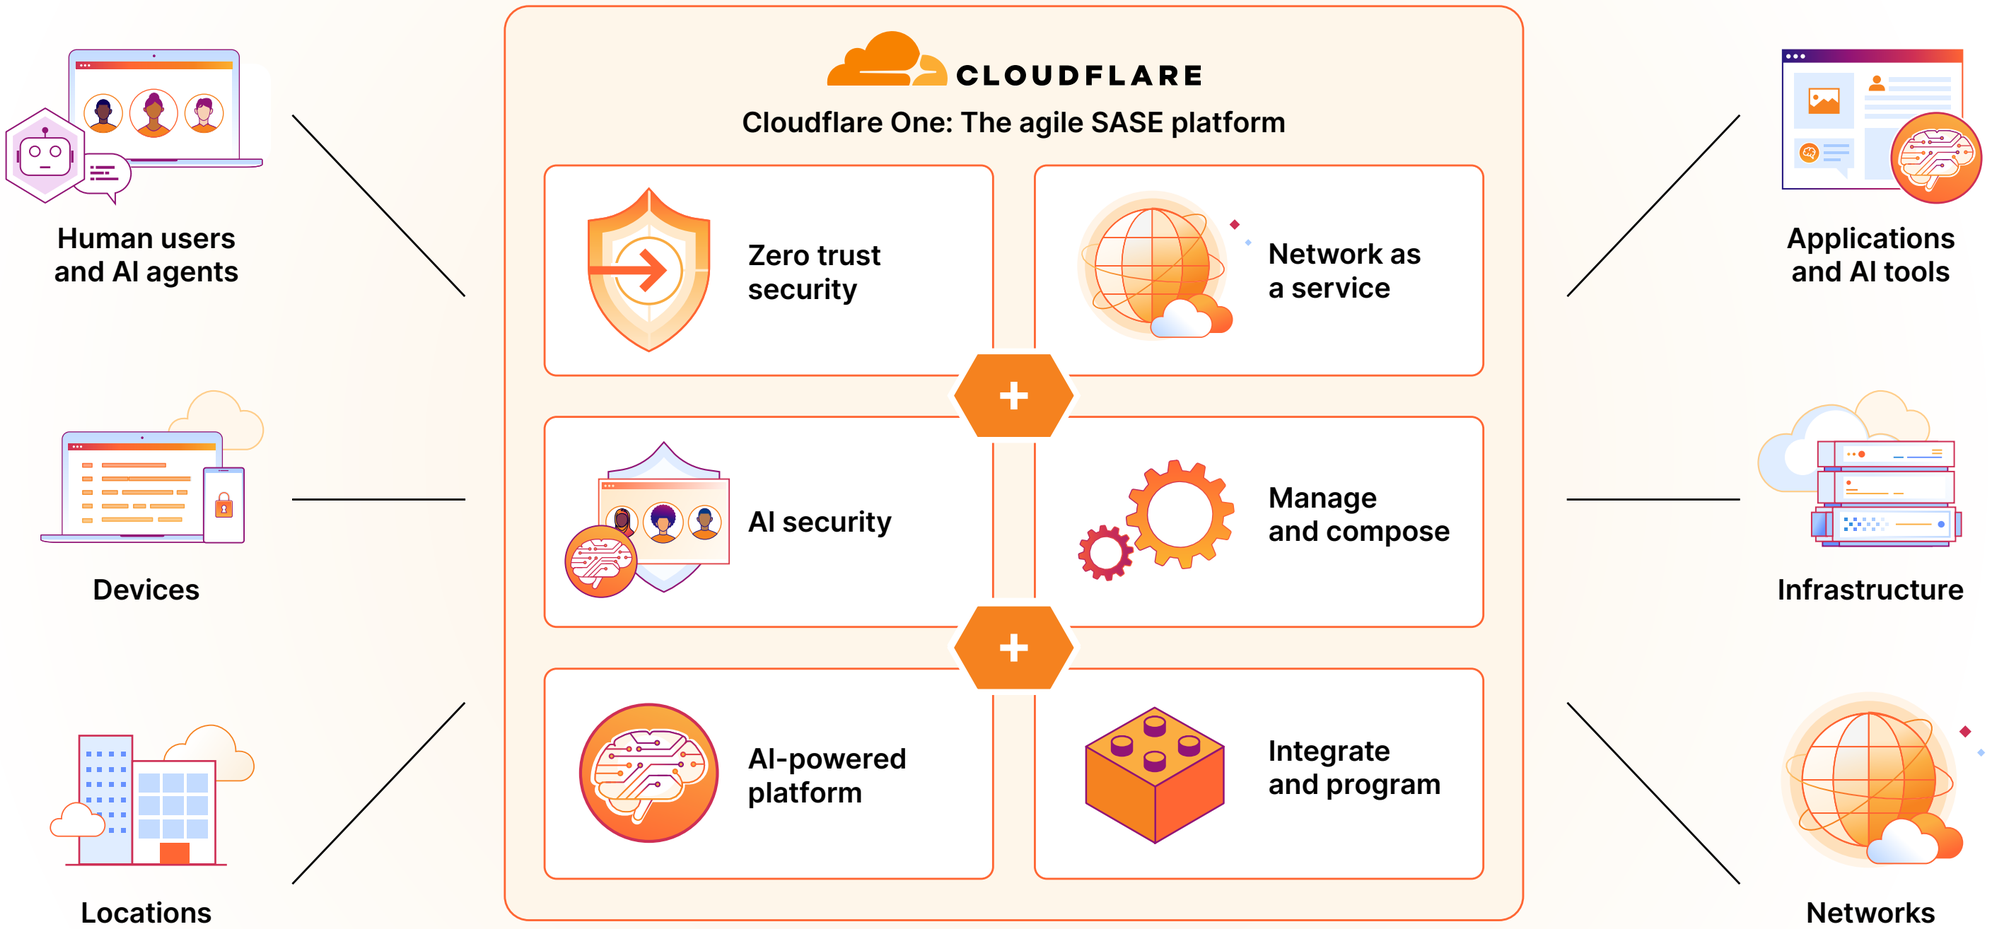

ABOUT CLOUDFLARE

Protect GenAI use and govern AI agents

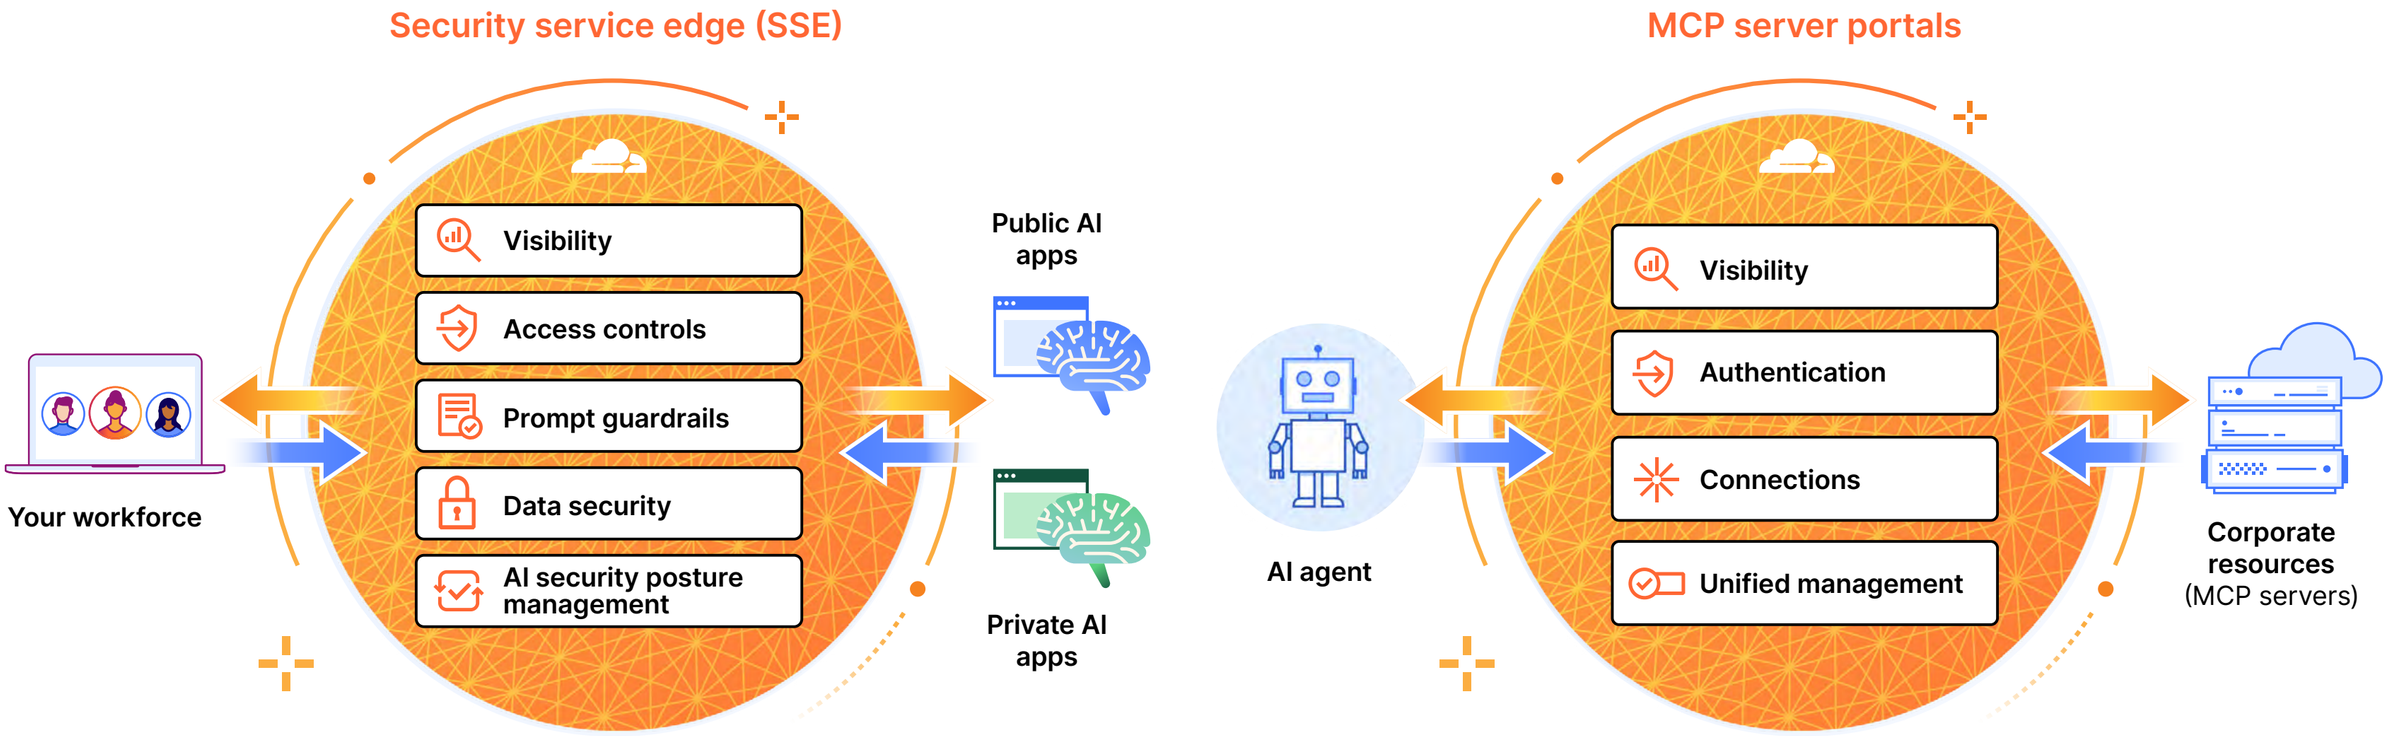

ABOUT CLOUDFLARE

Cloudflare AI services across the full lifecycle

Build AI
with Developer Platform

Build AI on every region on earth, with one click

Workers AI, Vectorize, AI Gateway

Build AI infrastructure, on easy mode

Remote MCP Server, Agents SDK, AI Search, AI Gateway

Protect AI adoption
with AI Security Suite

Secure workforce use of GenAI

SASE platform ***(Cloudflare One)***

Govern AI agents

Cloudflare's ZTNA including ***MCP Server Portals***

Build AI securely

AI Gateway

Protect AI-powered apps

App Security including ***Firewall for AI***

Protect content
with app services

Protect original content online from AI crawlers

Bot Management including AI Crawl Control

AI-powered platform on one global network

Threat detection models · AI agent (Cloudy) · Data loss prevention models

ABOUT CLOUDFLARE

Insights for the modern CxO

Navigating today's threat landscape and rapid technological shifts requires more than operational knowledge — it demands strategic foresight. "The Executive Lens" by Cloudflare is a dedicated resource hub curated specifically for C-suite leaders.

Learn expert-driven insights, actionable frameworks, and exclusive research on critical enterprise topics like cyber resilience, secure AI governance, and global digital transformation.

Explore The Executive Lens today.

Read more

Additional resources

Forrester Total Economic Impact

Meet sophisticated threats and prevent emerging ones. See how Cloudflare helps enterprises use security as a competitive advantage, weathering a complex threat landscape with greater efficiency and predictability.

Read more

Security Signal

Uncover the signal from the noise and focus on today's most important cybersecurity trends. Each episode of Security Signal translates cybersecurity complexities into actionable intelligence for executives at the helm.

Watch now

2026 Cloudflare Threat Report

Understand the 2026 threat landscape defined by a new Measure of Effectiveness (MOE). The report details new risks from state-sponsored pre-positioning, token theft, hyper-volumetric DDoS, and more.

Read more

theNET

Insights across cybersecurity innovation, the threat landscape, and the future of the Internet with executive perspectives on how to solve organizational challenges with technology.

Read more

Contacts

	Global	Americas	EMEA	Asia Pacific	Japan
Market Leadership	 **Mark Anderson** President of Revenue markanderson@cloudflare.com	 **Rick Congdon** Geo Vice President, Americas congdon@cloudflare.com	 **Tony Van den Berge** Geo Vice President, EMEA tonyberg@cloudflare.com	 **Goran Risticevic** Geo Vice President, APAC goran@cloudflare.com	 **Sayoko Matsumoto** Geo Vice President, Japan sayoko@cloudflare.com
Field CXO Team	 **Ramy Houssaini** Chief Cyber Solutions Officer ramy@cloudflare.com	 **Khalid Kark** Field CIO, Americas khalid@cloudflare.com	 **Christian Reilly** Field CIO, EMEA creilly@cloudflare.com	 **Volker Rath** Field CISO volker@cloudflare.com	 **Koichiro Otobe** Field CTO, Japan koichiro@cloudflare.com

CLOUDFLARE

2026

Cloudflare Security Signals Report

Autonomic Resilience

1 888 99 FLARE | enterprise@cloudflare.com | Cloudflare.com

REV:BDES-8897.2026MAR19

Endnotes

1. Jonathan Villa, “Hidden Risks of Shadow AI,” Varonis, www.varonis.com/blog/shadow-ai. Accessed 11 Feb. 2026.
2. IBM, “Cost of a Data Breach Report 2025,” www.ibm.com/reports/data-breach. Accessed 11 Feb. 2026.
3. MultiState, “Artificial Intelligence (AI) Legislation,” www.multistate.ai/artificial-intelligence-ai-legislation. Accessed 11 Feb. 2026.
4. Gartner, “Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up From Less Than 5% in 2025,” 26 Aug. 2025, www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025. Accessed 11 Feb. 2026.
5. Cloudflare Radar, “Bot Traffic,” radar.cloudflare.com/bots?dateRange=12w. Accessed 11 Feb. 2026.
6. Cloudflare Radar, “Application Layer Security,” radar.cloudflare.com/security/application-layer?dateRange=12w. Accessed 11 Feb. 2026.
7. IBM, “Cost of a Data Breach Report 2025.”
8. Lareina Yee, et al., “The AI Reckoning: How Boards Can Evolve,” McKinsey & Company, 24 Oct. 2024, www.mckinsey.com/capabilities/mckinsey-technology/our-insights/the-ai-reckoning-how-boards-can-evolve. Accessed 11 Feb. 2026.
9. IBM, “Cost of a Data Breach Report 2025.”
10. ENISA, “SBOM Analysis - Towards an Implementation Guide.” Dec. 2025, www.enisa.europa.eu/sites/default/files/2025-12/SBOM%20Analysis%20-%20Towards%20an%20Implementation%20Guide_v1.20-Published.pdf. Accessed 11 Feb. 2026.
11. Verizon, “2025 Data Breach Investigations Report (DBIR),” www.verizon.com/business/resources/reports/dbir. Accessed 11 Feb. 2026.
12. Cloudflare, “2026 Cloudflare App Innovation Report,” 2026, www.cloudflare.com/resource/g/app-innovation-report/2026. Accessed 11 Feb. 2026.
13. CrowdStrike, “2025 Global Threat Report,” www.securityweek.com/wp-content/uploads/2025/02/CrowdStrikeGlobalThreatReport2025.pdf. Accessed 18 March 2026.
14. SANS Institute, “SANS 2025 CTI Survey: Cyber Threat Intelligence Survey,” SOCRadar, May 2025, socradar.io/wp-content/uploads/2025/05/SANS-2025-CTI-Cyber_Threat_Intelligence_Survey-SOCRadar.pdf. Accessed 11 Feb. 2026.
15. SANS Institute.
16. SANS Institute.
17. Mohammed Khalil, “Vulnerabilities Statistics 2025: Record CVEs, Zero-Days & Exploits,”DeepStrike, 8 Oct. 2025, deepstrike.io/blog/vulnerability-statistics-2025. Accessed 25 Feb. 2026.
18. VulnCheck, “VulnCheck State of Exploitation 2026,” 21 Jan. 2026, www.vulncheck.com/blog/state-of-exploitation-2026. Accessed 11 Feb. 2026.
19. Cloudflare Global Network Data.
20. Pegasystems, “Average Global Enterprise Wastes More Than $370 Million Every Year Through Technical Debt, Says Research,” 14 Oct. 2025, www.pega.com/about/news/press-releases/average-global-enterprise-wastes-more-370-million-every-year-through. Accessed 11 Feb. 2026.
21. Protiviti, “Global Technology Executive Survey: Tech Debt a Major Burden,” www.protiviti.com/us-en/global-technology-executive-survey-tech-debt-major-burden. Accessed 11 Feb. 2026.
22. Cloudflare, “2026 Cloudflare App Innovation Report.”
23. Pegasystems, “Average Global Enterprise Wastes More Than $370 Million Every Year Through Technical Debt, Says Research.”
24. Cloudflare, “2026 Cloudflare App Innovation Report.”
25. Cloudflare, “2026 Cloudflare App Innovation Report.”
26. Uptime Institute, “Uptime Annual Outage Analysis Report 2025,” 6 May 2025, uptimeinstitute.com/about-ui/press-releases/uptime-announces-annual-outage-analysis-report-2025. Accessed 11 Feb. 2026.
27. Nuno De la Torre, et al., “IT Resilience for the Digital Age,” McKinsey & Company, 20 June 2023, www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/it-resilience-for-the-digital-age. Accessed 11 Feb. 2026.
28. Ashwin Chaudhary, “Managing Cloud Misconfigurations Risks,” Cloud Security Alliance, 14 August 2023, cloudsecurityalliance.org/blog/2023/08/14/managing-cloud-misconfigurations-risks. Accessed 11 Feb. 2026.
29. IBM, “Cost of a Data Breach Report 2025.”
30. IBM, “Cost of a Data Breach Report 2025.”