



Defense in an interconnected, AI-accelerated world

2026 Microsoft Digital Defense Report
Executive summary



Where we are in 2026

Cyber risk is accelerating at rapid speed, driven by the connections among identities, AI systems, cloud services, software supply chains, edge infrastructure, and critical systems.

Empowered by the vast capabilities of AI, attackers are pursuing all available entry points, including points of leverage where one compromise can create effects far beyond the original target, setting up cascading waves of compromise. The methods attackers are using and their targets haven't changed—it's the scale and effectiveness of their attacks that have.

- What's different this year**
- AI is significantly accelerating attacker workflows while creating a new enterprise attack surface.
 - Identity-led intrusions are leading to credential theft that drives follow-on attacks, creating a flywheel of unauthorized access.
 - Open-source software, edge devices, cloud identity, and operational technology are being attacked at an unprecedented rate.
 - The constant pace of attacks means resilience is now leadership discipline, not solely a technical recovery objective.
- In the pages that follow, we walk through the biggest challenges and changes confronting defenders not only this year, but also the next, and what organizations can do to reduce the threat.

63%
of intrusions
involved data theft

5.3 hrs
average time
before exposed cloud
workloads were attacked

93%
of voice
phishing attacks
kept the victim on the line
long enough to begin social
engineering

46M+
business
email
impersonation attacks
were detected

89-95%
of email phishing
attachments
led to a credential theft effort

Source: Microsoft Threat Intelligence



As attackers accelerate with AI, defenders must keep pace

What was once a threat on the horizon has now arrived. AI is changing the economics of attack and defense, with the advantage going to attackers in the near term.

Attackers now have a clearer picture than ever of what they want from AI systems and how to use AI to achieve their traditional goals. Attackers are:

- Attacking AI models (inputs to model building, model data, and production infrastructure) and systems (AI systems themselves, their user data, and their output artifacts)
- Attacking computer systems, people, and information, including using AI to:
 - Find vulnerabilities
 - Generate custom malware
 - Shorten the cycle of lateral movement, secret discovery, and data exfiltration
 - Control and manage entire attack chains
 - Scale phishing and social engineering

The AI target landscape


For sophisticated actors, AI allows unprecedented speed, scale, and customization, reducing the attack chain from days to seconds. For less-sophisticated actors, AI-powered scaling makes accessible the sort of attack persistence that was previously the sole domain of intelligence agencies, and the ability to customize attacks, especially social engineering attacks for phishing and fraud, is likely to increase attack success rates. These trends will only be compounded with increased leveraging of AI agents in the near future.

The AI threat landscape is in constant motion; and is highly likely to be commoditized within a year, and new categories of threats will emerge. While the equilibrium between attackers and defenders will likely ultimately be re-established, we are now in a period where attackers are getting to advantages first, and defenders will need to move sharply to close the gap.

AI is now an essential, not optional, tool for defenders. The ability to automatically and persistently hunt for vulnerabilities and fix them creates a potential world of “persistent defense” that was never before possible at scale. At the same time, to defend against the increasing speed and scale of attacks, defenders will need to rely on people working alongside agentic AI defensive systems, trusting the agents with increasing levels of autonomy to take defensive actions on their behalf. Attacks conducted at machine speed will only be confrontable by defenders moving with the same speed.

Threat intelligence takes on a new urgency in this context. Of all the factors over which defenders have control, inter- and intra-organization signal correlation—the degree to which a defender’s systems share threat intelligence with each other internally or externally—is the primary strategic variable. Each additional source of visibility strengthens the value of every other source, improving detection, accelerating response, and providing broader context for security decisions. Cooperation and collaboration empower defenders and make it harder for attackers to succeed.

Figure 1: The target landscape





Adversaries are attacking identity from all angles

Over time, attacks are conducted in the same ways against the same targets. The most damaging intrusions rarely come from something novel; they depend on a fundamental that was missing, misconfigured, or quietly bypassable.

What has changed in 2026 is the size of the attack surface and the speed at which these attacks arrive. Several forces are intensifying this pressure: more exposed edge infrastructure, identity sprawl, fragile supply chain trust, the human factor, and AI.

There continues to be a fundamental shift by attackers to exploit trusted systems, operational blind spots, and the scale of modern digital ecosystems. They're relying less on custom malware and instead abusing legitimate tools, cloud services, and administrative functions. This makes detection harder because malicious activity often blends in with normal user behavior.

Here are specific trends that we're seeing:

- User execution, valid accounts, social engineering, and phishing represent the majority (73.3%) of initial access attempts—adversaries are targeting humans and identities. Adversary-in-the-middle (AiTM) phishing kits, OAuth app abuse, API exploitation, and password spray attacks are now standard across both criminal and nation-state playbooks, replacing endpoint malware and exploit kits.




- Over half of intrusions lead to follow-on credential theft—once inside, attackers immediately harvest more credentials either for lateral movement or sale on the dark web. A significant portion lead to internal password spray campaigns. This creates a self-reinforcing cycle: one compromised identity fuels the compromise of many more.
- Across China, Iran, North Korea, and Russia, state-sponsored cyber operations continue to evolve toward more persistent, scalable, and strategically aligned access models, with a shared emphasis on long-term positioning inside high-value environments. This means embedding operations within legitimate authentication flows, using account compromise and cloud-based collection, and leveraging high-credibility social engineering.
- Nation-state and ransomware actors keep weaponizing disclosed flaws, especially in internet-facing infrastructure, identity and management appliances, developer tooling, and third-party software. The window between vulnerability disclosure and exploitation is shrinking rapidly, often to just a few days. With AI accelerating vulnerability discovery and scanning, organizations will likely have less time to patch and protect exposed systems.

How attackers leverage identity in the attack chain

With identity in the crosshairs, prevention, detection, and response programs need to be equally strong across endpoint execution controls, identity protection, phishing resistance, and browser defenses. Phishing-resistant multifactor authentication (MFA), disciplined identity hygiene, tiered administration, and strong privileged-access enforcement remain the highest-leverage controls available. Advanced detection systems are also necessary to identify correlated behaviors across hybrid attack chains and automatically isolate devices, disable compromised accounts, revoke sessions, or quarantine cloud resources. By compressing detection-to-response timelines from hours to minutes, automated disruption significantly reduces adversary dwell time and limits attack progression.

Microsoft Incident Response observed a consistent pattern





Adversaries are attacking identity from all angles (continued)

Establishing a strong defense

Protecting human and non-human identities means strengthening identity-centric controls, reducing reliance on reusable credentials, and improving visibility. Organizations should:

- **Prioritize identity-centric controls** by implementing phishing-resistant authentication and reducing reliance on long-lived credentials.
- **Strengthen visibility into identity usage** by monitoring for anomalous access patterns, new permissions, and previously unseen application or API activity.
- **Reduce the attack surface of non-human identities** by enforcing least privilege, rotating credentials regularly, and restricting access to trusted locations.
- **Validate exposure through proactive testing.** Conduct assume-breach exercises focused on cloud services, software-as-a-service (SaaS) platforms, APIs, and workflow automation capabilities, prioritized based on where the highest-value data resides, is accessed, or can be aggregated.

AI and the security operations center (SOC)

AI is reshaping day-to-day SOC operations beyond detection and disruption. Large language models (LLM) and specialized agents can summarize incidents, generate investigative narratives, write log hunting queries, and automate executive reporting. AI-assisted triage can classify incidents, recommend investigative steps, and surface relevant evidence across multiple data sources without requiring analysts to manually pivot between tools. For detection engineering teams, AI accelerates the development, testing, and validation of analytics, reducing the time required to deploy new detections. Operational leaders can use AI-driven reporting and queue analysis to reduce administrative overhead and improve visibility into response performance. This allows skilled defenders to focus on higher-value activities such as threat hunting, hypothesis development, and adversary disruption rather than repetitive operational tasks.

Organizations that successfully integrate infrastructure context, threat intelligence, AI-driven correlation, and automated disruption into a unified detection strategy will be better positioned to keep pace with modern adversaries. The SOCs most likely to succeed are not those with the largest analyst teams or the greatest alert volume, but those that most effectively convert high-scale telemetry into rapid, high-confidence disruption of malicious activity.

Modern problems require modern solutions

For much of the past decade, cybersecurity strategies have been shaped by focusing on distinct risks—for example, ransomware or distributed denial-of-service (DDoS) attacks.

While these threats remain very real, we have entered an era of polycrisis, one defined not by isolated events but by overlapping disruption, leaving little to no breathing room to recover. Geopolitical instability, economic pressure, rapid technological change, and persistent cyber activity interact in ways that increase impact and reduce response time. Digital systems are no longer stressed by a single event; they are tested by overlap, dependency, and cascading failure.

This shift exposes a widening gap between how organizations defend and how risks emerge. Fragmented defenses, optimized for individual threats or compliance requirements, become less successful when crises unfold simultaneously. Today's environment demands a move from fragmentation toward persistent digital resilience. Defending against threats to interconnected systems requires visibility across domains, disciplines, and geographies, as well as an understanding that digital systems are inseparable from societal resilience. This reality challenges traditional approaches to cyber defense such as defense-in-depth, where controls are added in layers, and gives greater focus to resiliency. Resilience by design assumes that failures will occur, that dependencies matter as much as adversaries, and that recovery and continuity are as critical as prevention.



Modern problems require modern solutions (continued)

In a polycrisis environment, digital defense must evolve from guarding individual assets to developing systems that will function under pressure so that services will continue. This requires organizations to design security for disruption, not just compliance, to align cyber resilience with broader continuity planning, and to treat identity, data, and infrastructure as inseparable. It also demands deeper collaboration—between public and private sectors, across borders, and among organizations that share infrastructure and risk—to ensure that dependencies are known and managed.

Major supply chain attacks

Open-source supply chain attacks are becoming more frequent and more damaging as attackers move upstream into the systems organizations use to build and release software. A single stolen credential or compromised workflow can give attackers a path to many downstream organizations at once. Since open-source software underpins critical infrastructure across governments, economies, and essential services, sustaining resilience will require a more active government role to strengthen the institutions, incentives, and guardrails that support it. This means accelerating adoption of AI-enabled secure-by-design tools across the ecosystem, including automated vulnerability detection and stronger supply chain governance, expanding software bills of material (SBOMs), and establishing coordinated vulnerability disclosure.

Threats to critical infrastructure

Critical infrastructure defenders face a fundamental mismatch: adversaries operate continuously, while traditional security programs remain periodic. Undetected exposure does not just risk data loss, but can cascade into physical disruption, service instability, and safety impacts. Nation-state actors are pre-positioning within infrastructure networks to enable future disruption or sabotage of critical operations, while financially motivated actors increasingly exploit unpatched systems targeting availability.


Threat and vulnerability management (TVM) must evolve into a continuous, integrated operational discipline, not a set of tools or periodic activities. Executed correctly, TVM becomes a closed-loop defense system binding asset visibility, vulnerability identification, threat detection, and AI-accelerated insight, each capability reinforcing the others to create a compounding defensive advantage. At the same time, reporting metrics should shift from number of patches deployed to exposure reduced, detection coverage increased, and time-to-mitigate compressed. The outcome is not incremental improvement, but compound resilience. Closing the loop between proactive exposure reduction and reactive threat detection shrinks the attack surface while catching what slips through.

Developing national cyber resilience

Across digitally transforming economies, cyber outcomes are increasingly shaped by both technical controls and how effectively governments coordinate during an incident. This response effectiveness often depends on institutional readiness—clear roles, decision-making authorities, and coordination mechanisms—rather than detection or prevention capabilities. As public sector systems become more interconnected, incidents can trigger cascading effects across ministries and sectors, exposing gaps in real-world response that are not always visible in policy frameworks.

Microsoft’s Advancing Regional Cybersecurity (ARC) initiative engagements have identified three recurring structural challenges: unclear roles across agencies, fragmented escalation pathways, and limited mechanisms for whole-of-government crisis management. These gaps are particularly visible in scenario-based exercises, where coordination often breaks down at the intersection of technical response and policy decision making—for example, around public communication, system prioritization, or cross-sector resource allocation. These findings reinforce a broader trend: resilience is as much a governance challenge as it is a technical one.

The findings also point to several practical readiness measures: establishing a shared incident response lifecycle and playbooks for high-impact scenarios; creating a streamlined “report once” mechanism for incident reporting; linking organizational SOCs to sector-level coordination structures; and defining clear information sharing protocols that enable speed while protecting sensitive data. Regular tabletop exercises and after-action reviews are also essential to operationalize and continuously improve these arrangements.



Ten security priorities for organizations

Cyber risk is now a continuous test of business resilience. The boundaries between state-sponsored activity, cybercrime, and influence operations have continued to erode. This converges into sustained and overlapping campaigns that have tested not only technology estates, but the resilience of organizations and the trust placed in them. For business leaders, this means focusing less on reacting to the last breach and more on how long critical operations can be sustained under continuous cyber stress.

In this environment, cybersecurity is a continuous business condition, defined by persistent adversary pressure, machine-speed operations, and systemic interdependence. The defining question for boards today is not whether disruption will occur, but whether the organization can continue operating, contain harm, recover quickly, and maintain trust when it does.

The response starts with secure foundations: Secure by Design, Secure by Default, and Secure Operations, applied through identity-centric Zero Trust. These foundations make the following priorities durable as AI expands the number of identities, applications, agents, and dependencies organizations must govern.

The organizations that succeed will be those that engineer for endurance: building systems that can absorb disruption, adapt under pressure, and maintain trust even when compromised. For security leaders and boards, building that resilience requires these priorities:

1 Govern cyber risk at leadership level Report identity exposure, patch latency, agent permissions, critical dependencies, dwell time, and recovery readiness.	2 Defend identity as the primary control plane Protect identity by deploying passkeys and phishing-resistant MFA, reducing privilege, governing human and non-human identities, and using rapid invalidation/rotation for tokens and session credentials.	3 Establish governance, visibility, and data controls for AI Maintain visibility into approved AI tools, services, and agentic capabilities while monitoring for unauthorized or unmanaged use. Apply sensitivity-aware access, data protection, and governance to the data those systems and agents can reach.
4 Reduce exposure and edge risk Continuously discover internet-facing assets and prioritize rapid patching of edge devices, enabling centralized and high-fidelity logging, enforcing strong authentication, and integrating network-level visibility with threat intelligence.	5 Focus on cross-signal correlation Correlate endpoint, identity, cloud, application, email, and network signals to detect intrusions earlier. Cross-signal correlation also enables defenders to anticipate attack paths and targets.	6 Enact a continuous threat vulnerability management model Combine asset visibility, vulnerability identification, threat detection, and AI-accelerated insight to create a continuous threat vulnerability management program.
7 Validate exposure via proactive testing Conduct assume-breach exercises focused on cloud services, SaaS platforms, APIs, and workflow automation capabilities. Prioritize testing based on where the highest-value data resides, is accessed, or can be aggregated.	8 Secure software dependencies Harden software installation, maintain a software inventory, validate provenance, monitor packages, and prepare rapid containment.	9 Participate in collective defense Share actionable threat indicators and victim-informed intelligence with trusted partners and law enforcement to support coordinated threat tracking, investigation, and disruption.
10 Prepare for long-horizon technological transitions Build a cryptographic inventory and advance post-quantum planning alongside AI governance.	The defining question The defining question for boards today is not whether disruption will occur, but whether the organization can continue operating, contain harm, recover quickly, and maintain trust when it does.	



2026 Microsoft Digital Defense Report

For more news on cybersecurity, visit:

<https://microsoft.com/corporate-responsibility/cybersecurity>

For more report insights, visit:

<https://microsoft.com/mddr>

**For more news on cybersecurity policy,
follow us on LinkedIn:**

<https://aka.ms/MOILinkedin>

For insights and trends for security leaders, visit:

<https://www.microsoft.com/security/security-insider>

A Microsoft Threat Intelligence report

October 2026