Microsoft

Secure Future Initiative

Security above all else

July 2026 progress report

The Microsoft Secure Future Initiative (SFI) is a continuous commitment to revolutionize the way we design, build, test, and operate our products and services, to achieve the highest security standards.

Contents

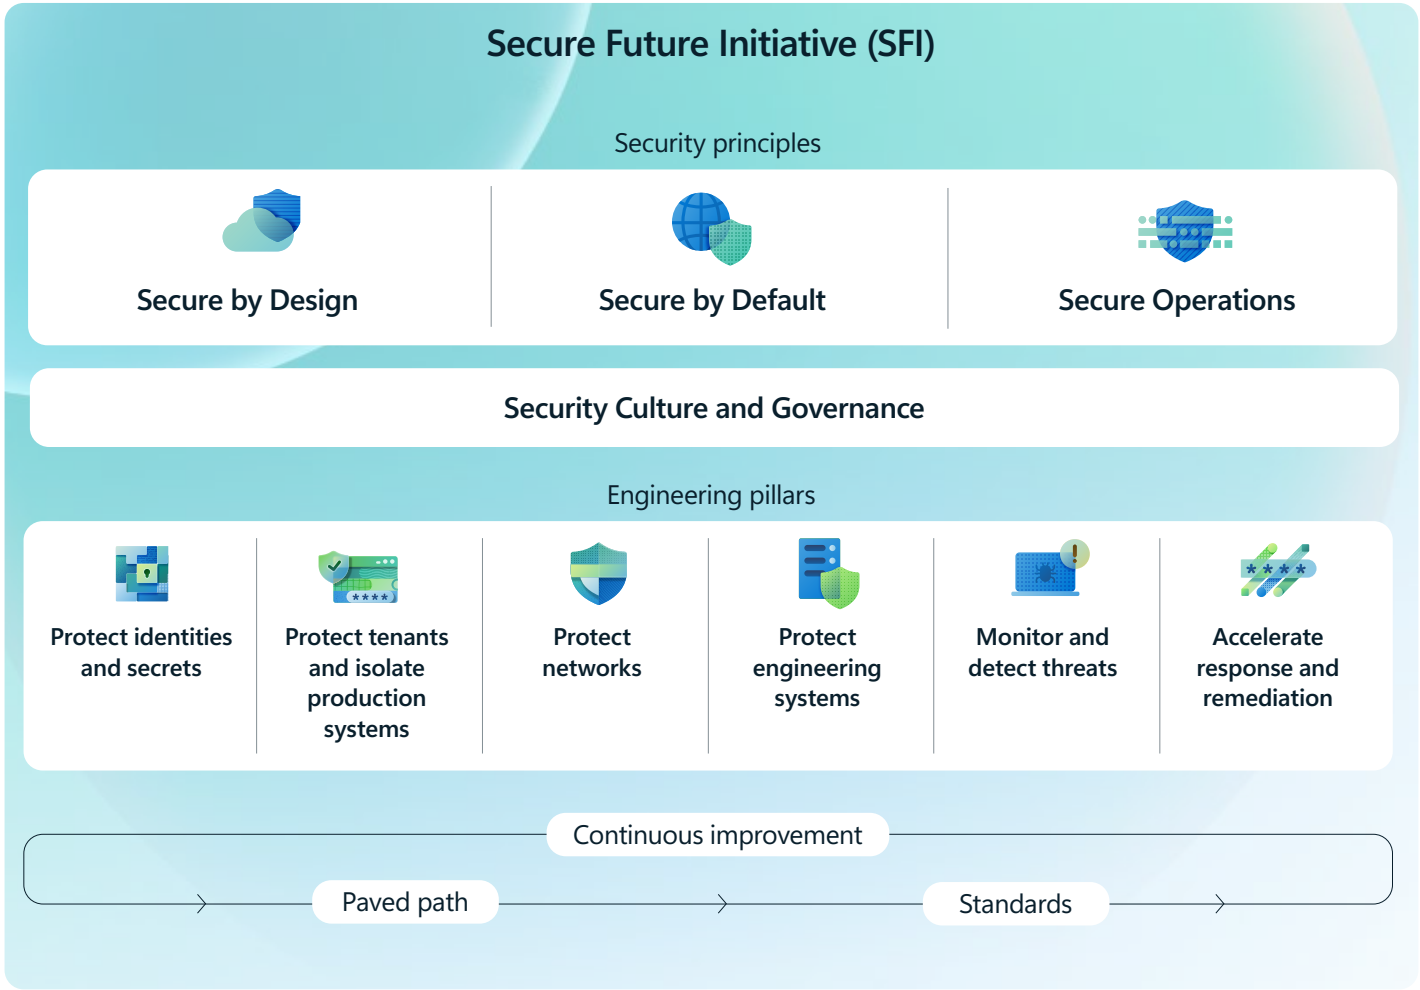

Introduction

Security is never finished.

AI has reshaped both sides of cybersecurity. Attackers use frontier models to discover vulnerabilities, chain attack paths, and scale exploitation faster than manual approaches allow. Every defensive improvement drives attacker adaptation: close one path, attackers move to the next.

The same capabilities are accelerating defense. AI lets us detect, evaluate, and remediate faster than manual review ever could, turning the shift into an advantage. Whether defending against traditional attacks, or AI-led attacks, security is continuous, not a destination. It cannot be treated as a milestone or a point-in-time investment. It is a discipline: designed in, enabled by default, continuously validated, and re-evaluated as the threat landscape evolves.

The Secure Future Initiative (SFI) is how we operationalize that discipline at Microsoft: embedding security across culture, governance, and engineering, and strengthening our ability to adapt as quickly as adversaries do.

SFI stringently applies Zero Trust principles as controls embedded directly into platforms, engineering pipelines, and governance systems. The goal is that every identity is verified, every access is governed, every resource is protected through consistent policy enforcement, and controls are standardized, measured, and enforced across services so that protections persist as environments evolve and new threats emerge.

In this fourth report we share progress made against SFI pillars and objectives through three themes that make it easier to see how individual controls contribute to broader security outcomes.

We also share updates in the areas of culture, governance, security principles, and each engineering pillar. Collectively these updates reflect meaningful progress that has improved security for Microsoft and our customers: three objectives have reached their target state, three are nearing completion, and twelve have made significant progress.

As AI accelerates both attack and defense, we will continue strengthening foundations, advancing proactive defense, and preparing for emerging risks, including the transition to quantum-safe security.

Security is never finished. Every advance in technology creates new opportunities, new risks, and new responsibilities. We look forward to sharing continued progress, lessons learned, and practical guidance as we work together to build a safer digital future.

Secure foundations:

Reduced risk through hardening, secure engineering baselines, asset inventory, segmentation, security boundary isolation, and enforcement-by-default.

Proactive defense:

Used AI, high-quality telemetry and security signals, and behavior-based detection to find real risk earlier, prioritize what matters, and respond decisively.

Future readiness:

Preparing for emerging risks including post-quantum cryptography.

Section 2

Progress made, learnings, & customer guidance

Secure foundations

A sophisticated attacker does not evaluate components in isolation. They evaluate the attack path.

The most consequential security failures we observe across the industry rarely come from a single missing control. They come from environments where identity gaps, unmanaged assets, stale tenants, inconsistent configurations, over-privileged applications, exposed credentials, and unvalidated network trust assumptions sit side by side. Each issue can appear manageable on its own. Together, they create composite attack paths that determined adversaries can chain together. AI makes this more pronounced by increasing the speed, scale, and consistency with which attackers can discover, reason over, and exploit those paths.

SFI was created to improve the security posture of Microsoft, not incrementally, but systemically. We have made significant progress strengthening our foundation: hardening identity controls, defining explicit tenant boundaries, enforcing engineering defaults, and applying continuous validation across layers. This positions us to meet a threat landscape where AI accelerates the speed and sophistication of attack paths. This section describes how SFI established consistent security baselines to ensure that a weakness in one layer does not become a path across the environment.

Key progress

The results below reflect two reinforcing goals: consistent security baselines that reduce variation and drift, and enforced boundaries that narrow what an attacker can reach if any one layer is bypassed.

Identity controls strengthened

Phishing-resistant MFA now protects 99.97% of user/device pairs at Microsoft. Legacy authentication continues to be removed from production paths, and credential lifetimes are shortened through enabling automated rotation of managed identities and enforcing just-in-time access and Privileged Identity Management (JIT/PIM). Guest accounts not managed by an organization's director were reduced by 228,000, with further cleanup in progress. More than 29,000 privileged accounts are being brought under JIT/PIM, and dual-key, time-bound emergency access, reducing reliance on administrative privilege.

Tenant and workload boundaries explicitly enforced

Cross-boundary credential isolation reached 98.7%. More than 3,700 production certificates that were present in productivity environments have been revoked, substantially reducing a historical cross-tenant lateral movement path into production. Tenant inventory and governance are expanding across production, productivity, auxiliary, and ephemeral environments, with the goal of governing tenants from creation through retirement and using drift detection to maintain security baselines. More than 1.4 million unused apps (up from 1.2 million since November) have been decommissioned. Traffic for multi-tenant apps is now being constrained to allowed-tenant lists, and service-to-service authentication into production is being further scoped to approved network boundaries.

Blast radius reduced through segmentation and isolation

More than 732,000 resources have had public access revoked. Network isolation enforcement is continuing scaling across an additional 1 million resources as segmentation expands from high-risk workloads to broad coverage. We have fully retired the remaining classic Azure Service Manager model, removing legacy access mechanisms that could be abused for credential replay or lateral movement.

Engineering defaults hardened to prevent drift

Build-time network isolation prevents 83% of pipelines from access to unapproved package endpoints, Linux packages, and mirrors. Pull-request (PR) policies that block vulnerable packages, enforce the use of approved package feeds, and continuously scan open-source dependencies, prevent vulnerabilities from entering the codebase. Enforced provenance and paved-path templates carry these controls into new deployments, resulting in the most secure path becoming the default path.

This progress forms reinforcing layers of security. Identity feeds into access governance. Access governance feeds tenant and application governance. Tenant and application governance feeds segmentation. Segmentation contains blast radius. Engineering defaults reduce what enters production in the first place, and continuous monitoring makes the resulting posture observable and actionable.

Secure foundations

What we are learning

Foundations are durable only when they are continuously validated, not periodically audited.

Standardized telemetry, enforced baselines, and complete asset and tenant inventories allow drift to be detected and remediated in near real-time. When an incident occurs, defenders can scope affected systems and validate impact quickly because the boundary state is known, not reconstructed.

Attackers adapt across boundaries, not just within them

Continuous validation matters because attackers don't stand still. As Microsoft hardened identity tenants, engineering systems, and production infrastructure, adversary behavior shifted. Closing one part of the attack surface increases pressure on adjacent parts. Hardening in isolation can create the illusion of progress while leaving composite paths intact.

Identity: authentication flow abuse

As phishing-resistant MFA reduced classic credential capture, attackers shifted toward abusing legitimate authentication flows to obtain valid tokens without defeating MFA directly. Observed techniques include OAuth consent phishing, device code flow abuse, adversary-in-the-middle authorization code interception, token and session theft, and service-principal credential abuse. In each case, the attacker's goal is the same: obtain a valid access token meeting MFA requirements, by exploiting the trust that users and systems place in legitimate authentication processes.

Applications and secrets: credential lifecycle as attack surface

As user authentication hardened, application credentials, federated credentials, managed identities, and certificates became more important control points. Durable defense requires governing the full lifecycle of applications, secrets, and service identities.

Tenant and production boundaries: lifecycle gaps as attack surface

As tenant boundaries became more explicit, stale tenants, unmanaged applications, guest sprawl, exposed production credentials, and implicit cross-tenant trust became higher-value targets. Durable defense requires treating tenant and production isolation as lifecycle governance – secure creation, drift detection, policy enforcement, and cleanup – not a one-time boundary decision.

Production and high-impact systems: trust escalation as attack path

For the most critical systems, the question is not whether each control exists. It is whether any relationship allows a lower-trust system to influence a higher-trust system. As perimeter and tenant boundaries harden, these cross-trust relationships become the paths adversaries seek.

Network and infrastructure: exploiting residual trust

As segmentation expands, residual internet access, stale tenants, and implicit trust relationships become higher-value targets. The next phase is not only broader coverage but stronger enforcement, fewer exceptions, and clearer accountability for the trust that remains.

Engineering systems: the software supply chain as attack surface

As governed pipelines and PR-time blocking mature, attackers focus on dependencies, build systems, release pipelines, and non-human identities. The build environment must be treated as a production-grade security boundary.

The pattern is clear – each hardening action changes the adversary's calculus, but does not end it. The compounding effect of SFI is that attackers face a shrinking set of viable paths, while defenders gain better telemetry and detections, stronger defaults, and more precise prioritization for the paths that remain.

Secure foundations

> “AI changes the math for attackers and defenders alike. The same capabilities adversaries use to find weaknesses faster, we use to find and fix them first.

Salim Chawro, Corporate Vice President, Cloud Security and Executive Sponsor, SFI

Customer guidance

Organizations that achieve durable security treat it as a core platform capability, not a project.

- **Secure identity is the primary control plane.** Enforce phishing-resistant MFA, eliminate legacy authentication protocols, and shift to short-lived, automatically rotated credentials backed by hardware key protection.
- **Use multi-admin approval.** Implement multi-admin approval to help ensure that high-impact changes require validation and authorization from multiple administrators, preventing any single individual from making unilateral modifications.
- **Govern application and secret lifecycles.** Require admin consent for high-privilege scopes, enforce publisher verification, restrict user consent to verified apps, decommission unused app registrations on a schedule, and alert on new credentials added to existing app registrations or federated-credential issuers.
- **Restrict abusable authorization flows in conditional access.** Block or scope device code flow, deprecate Resource Owner Password Credential (ROPC), enable token binding and Continuous Access Evaluation (CAE), and apply sign-in risk policies that revoke sessions on impossible travel or risky token use.
- **Make tenant and workload boundaries explicit.** Inventory every tenant and classify it (production, productivity, build, dev/test, partner, ephemeral). Apply secure-by-default provisioning with drift detection, retire stale tenants, and constrain workload identities so service-to-service authentication into production is carefully managed, not implicit. For ephemeral tenants, establish clear ownership, approval chains, and automated lifecycle management so they are reliably decommissioned when no longer needed and new instances can be provisioned on demand.
- **Segment to limit blast radius.** Reduce unnecessary connectivity across network, identity, and application layers. Remove public access from resources that don’t need it, eliminate B2B sprawl and guest accounts not managed by an organization’s director, and replace standing privilege with JIT/PIM.
- **Require a hardened endpoint for production and high-risk access.** Treat the device as a boundary, not just an access point: gate sensitive workloads behind locked-down endpoints.
- **Define security standards and make them default.** Translate risk-prioritized requirements into engineering patterns, templates, and a governance control plane that teams adopt by default. Enforce at PR-time so the same weakness does not reappear in the next deployment.

Proactive defense

Secure foundations reduce the attack surface. Proactive defense builds on that foundation, using continuous evaluation, advanced detection, and AI-driven discovery to find and fix weaknesses quickly.

As the threat landscape evolves, security controls must be continuously evaluated to ensure they remain effective, layered appropriately, and working as intended in production.

Traditional security practices, including code reviews, penetration testing, and red-team exercises, remain essential. At the same time, AI models can discover vulnerabilities, chain exploit paths, and generate proof-of-concept attacks more quickly and across larger environments, helping security teams evaluate systems at speeds that outpace manual review. These capabilities are reshaping cybersecurity, enabling both attackers and defenders to discover more security issues and act on them more rapidly than before.

Microsoft continues to invest in AI-driven security capabilities to secure our internal systems, including a multi-agent AI system that complements traditional approaches by providing broad, proactive assessment of a cloud service's source code, configurations, identities, network connections, and running resources. This system builds on other tools in our security portfolio - such as the Microsoft Security multi-model agentic scanning system (codename MDASH), which scans source code to identify, validate, and prioritize vulnerabilities at scale – and then adds configuration, identity, network, and runtime context to comprehensively assess the service.

Microsoft is applying these capabilities with scale and rigor, integrating AI-driven security evaluation to strengthen our security posture and drive higher levels of protection for our customers.

Key progress

AI-driven proactive defense and vulnerability discovery

Microsoft is using a multi-agent AI system to proactively discover, validate, and prioritize security findings. This system suggests short-term compensating controls and provides durable architectural fixes, with built-in testing to confirm mitigations, which enables broader and more proactive security assessment across complex cloud environments.

Microsoft is also using AI tools to scan code, such as codename MDASH, and is incorporating advanced AI models into our Software Development Lifecycle to identify vulnerabilities and develop mitigations and updates. This allows us to discover more issues more quickly across a broader surface area than previous methods and address them earlier in the lifecycle. As discovery volume grows, scalable response workflows enable the pace of remediation to keep up with the pace of discovery. This shift is visible in increased security update volume in Patch Tuesday releases.

Detection scaling

More than 100 new detections added, bringing the total to 350+, and we are moving from signature-based detection to behavior-based and baseline-driven detection. AI-assisted authoring is increasing throughput and deployment velocity.

Supply-chain defense

More than 550,000 critical and high-risk open-source vulnerabilities remediated; ~3 million container vulnerabilities patched per month via automation.

> "Security foundations only work when they become part of the operating model. By embedding security across people, processes, technology, and governance, organizations can reduce systemic risk, build resilience, and earn trust through consistent execution."

Terrell Cox, Corporate Vice President and Deputy CISO, Customer Security

550,000

critical and high-risk open-source vulnerabilities remediated.

Proactive defense

What we are learning

Traditional defenses remain essential, but alone cannot keep pace with AI-accelerated threats. As attackers use AI and agentic techniques to discover and exploit vulnerabilities faster, defense must shift from reactive controls to proactive, continuous assessment. To meet this need, we built a multi-agent AI system that evaluates service security across the stack; identifying weaknesses, understanding how they can be exploited in practice, and enabling remediation before attackers can act. This is our approach to keep pace with AI-driven adversaries. The system is built on five core capabilities.

Full-service profiling

The system maps each service's architecture (components, data flows, trust boundaries, and risk exposure) to understand the complete attack surface rather than examining code in isolation.

Cross-layer evidence gathering

Specialized agents investigate code repositories, infrastructure and resource configurations, identity policies, network rules, and runtime telemetry simultaneously for a unified posture assessment.

In-depth reasoning

Rather than flagging individual findings alone, the system reasons about how gaps across domains could chain into multi-step attack paths, prioritizing what is exploitable in practice.

Dual-track remediation

For every confirmed vulnerability, the system produces both a compensating control for immediate risk reduction and a durable fix that addresses the root cause, so findings can be surfaced and fixed as early as possible.

Standards-driven continuous learning

As SFI security requirements update to address new threat classes, the system's evaluations automatically evolve, and feedback from service teams sharpen its accuracy with each cycle.

The multi-agent AI system generates findings and recommendations, which are validated and implemented by our security engineering teams. The system's findings have proven to be high quality and actionable, with over 90% of findings confirmed as genuine security issues by our security engineers – enabling proactive actions to improve security posture.

As these capabilities mature and scale across the organization, they feed findings back into secure foundations, creating a continuous loop where discovery strengthens the baselines that prevent future issues. We are also partnering with the broader security community to evaluate models, validate their effectiveness, and encourage the responsible use of frontier AI-driven security tooling across the industry.

› To learn more about Microsoft's multi-agent scanning approach read this blog.

90%+

of findings have been confirmed as genuine issues.

Customer guidance

- **Prioritize rapid patch adoption and resilience.** As AI-driven discovery increases the volume and pace of vulnerability disclosure across the ecosystem, timely patching becomes more critical. Where patches cannot be applied immediately, deploy mitigating controls, segment affected systems, and maintain offline backups to limit exposure.
- **Start with code scanning, then go system-level.** Leverage solutions like codename MDASH to identify, validate, and prioritize vulnerabilities earlier in the development lifecycle using multi-model agentic analysis and proof-based verification, enabling faster and more targeted remediation. Then evaluate how identity, code, configuration, and network relationships interact in production. Prioritize composite attack-paths over isolated findings.
- **Correlate signals across domains.** Aggregate telemetry across identity, endpoints, applications, and cloud so attacks are understood as coordinated actions, not isolated events.
- **Close the loop between detection and engineering.** Convert detection and incident findings back into baseline controls, paved-path templates, and PR-time enforcement so the same weakness doesn't reappear in the next deployment.
- **Monitor supply chain risk actively.** Maintain visibility into third-party dependencies, open-source components, and vendor security posture. Establish processes to evaluate, update, and replace components as vulnerabilities are disclosed.
- **Accelerate remediation with AI-assisted tooling.** Apply AI-driven capabilities to automate dependency updates, open-source package migrations, and vulnerability remediation at scale, reducing the time between discovery and fix while freeing engineering teams to focus on higher-complexity issues.

Future-ready security

Emerging threats are evolving faster than traditional planning cycles. Future-ready security is about anticipating structural risks early and building resilience before those risks become operational incidents.

The most urgent example is the transition to post-quantum cryptography (PQC). Advances in quantum computing and algorithmic techniques are compressing timelines, making it increasingly likely that today's cryptographic protections will not be sufficient against scalable quantum computers. Recent advances have shifted the risk horizon, and Microsoft is bringing its quantum-safe timeline forward so the transition can begin earlier and with greater confidence.

PQC has expanded from an objective under the "Protect Identities and Secrets" pillar to multiple SFI-measured engineering requirements, with defined ownership, milestones, and progress reporting across the company. This shift reflects the scale of the challenge: cryptographic risk is no longer isolated to a single domain, but spans platforms, services, protocols, and dependencies.

SFI will address this challenge by scaling cryptographic modernization across the enterprise – treating PQC not as a discrete upgrade, but as a coordinated, lifecycle transformation across systems, engineering practices, and governance.

› To learn more about the Microsoft approach to PQC, and what organizations can do now, read this blog.

Key progress

- **Accelerating the transition to quantum-safe security.** Through the Quantum Safe Program (QSP), Microsoft is bringing its quantum-safe timeline forward so the transition can begin earlier and with greater confidence.
- **PQC is now an enterprise engineering requirement.** Post-quantum cryptography has been operationalized into SFI with defined requirements, ownership, and measurable milestones across engineering teams, shifting from exploratory readiness to active execution. The goal is to transition products and services to PQC by 2029.
- **Platform integration has begun across core systems.** Early PQC capabilities are being integrated into foundational platforms, including operating systems, TLS implementations, and key management services, establishing the base for broader ecosystem adoption.
- **Windows enables early access to PQC.** Windows has begun shipping post-quantum cryptography capabilities, including updates to core components like Schannel and new platform APIs, demonstrating early, tangible progress toward integrating PQC protections into mainstream Windows security features.

> "Quantum-safe readiness is not a future problem—it's a security priority that needs to start now, giving customers the confidence to adopt new cryptographic standards.

Mark Russinovich, CTO, Deputy CISO, and Technical Fellow, Microsoft Azure

Future-ready security

What we are learning

As PQC and other structural risks scale, execution becomes the primary challenge: ensuring visibility, accountability, and sequencing across deeply interconnected systems.

Cryptographic transition: a coordinated, multi-layer effort

PQC workstreams span the full technology stack, including network protection (TLS 1.3 as a baseline for hybrid and post-quantum key exchange as standards mature), data protection (crypto-agility for data at rest), and trust systems such as code signing, certificates, and key management. Durable transition requires coordinated change across these layers rather than isolated upgrades.

Security is never finished. That conviction is where SFI started, and it is what keeps it moving forward.

Availability is no longer the constraint

Quantum-safe algorithms such as ML-KEM and ML-DSA are already available across major platforms, with broader platform integration (such as Azure Key Vault Managed HSM) underway. The primary challenge is no longer algorithm readiness but complexity, finding where cryptography lives across applications, infrastructure, identities, and certificates, then prioritizing the transition.

Transformation follows a familiar pattern

The post-quantum transition shares a pattern with Y2K (a global technology risk caused by how computers stored dates). Successful transition depends on having a detailed inventory for visibility, risk-based prioritization of critical components (starting with data in transit to address the "harvest now, decrypt later (HNDL)" threat), clear ownership of each action plan, transparent communication across teams with dependencies on each other, and carefully sequenced execution to avoid interruptions and instability.

Customer guidance

- **Inventory your cryptographic dependencies as a living system of record.** Maintain a continuously updated inventory of cryptographic dependencies across applications, services, protocols, and infrastructure. Assign clear ownership and integrate this inventory with risk tracking to help ensure visibility and accountability as PQC transition plans evolve.
- **Prioritize high-risk data flows first (HNDL exposure).** Focus early transition efforts on data in transit and long-lived sensitive data that may be vulnerable to HNDL scenarios. Use risk-based prioritization to sequence migration efforts where delayed action creates future exposure.
- **Design for crypto-agility across the full stack.** Ensure systems can adopt new cryptographic methods without major redesign. Use standardized implementations – such as supported TLS libraries and platform cryptography services (e.g., Schannel and OpenSSL) – to enable fast, coordinated updates across dependent systems.
- **Validate platform readiness and accelerate adoption.** Move to supported platforms and services that are actively integrating quantum-safe capabilities (e.g., updated OS versions, modern libraries, managed HSM services). Treat PQC as a tracked engineering program with defined milestones, reporting, and executive visibility.

Section 3

Culture, Governance, & Security principles

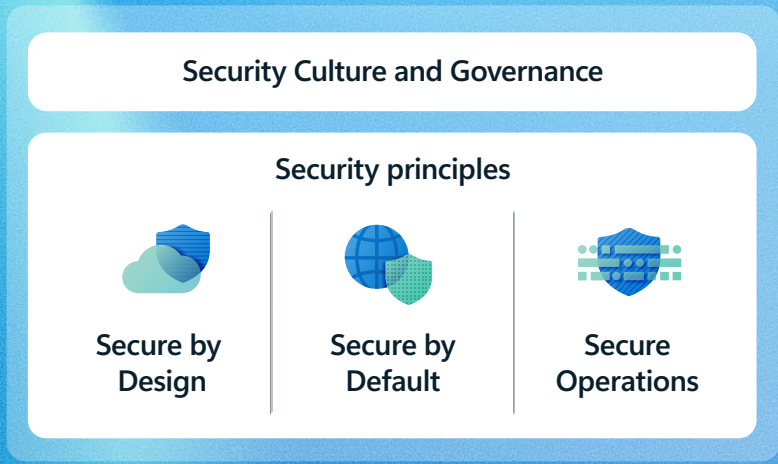

Culture

The progress we have made is the result of daily security work done by our employees. We continue to reinforce a culture that puts security first, strengthening performance expectations, engineering practices, and training and enablement. Security is embedded in how we work, lead, and how we measure impact.

Performance and development

Security continues to be a core responsibility for every employee. Security is integrated into the simplified levels that describe what and how an employee delivers impact. Within the performance management tool, prompts guide employees to describe security contributions, and optional AI features flag when security-related impact is missing. Employee sentiment data shows security is a top five theme in how employees describe their impact, suggesting security ownership is widespread across Microsoft.

Managers are accountable for ensuring their teams include security in their goals and deliver on Microsoft security priorities. Managers will continue to assess security contributions when evaluating impact and recommending rewards, ensuring a strong link between security, impact, and rewards.

Engineering sentiment

This year introduced new sentiment measures to track how security culture is experienced across Microsoft. Employee sentiment continues to reflect that employees prioritize addressing security challenges as part of their role, with average scores remaining stable at 88.

Within Engineering, employees continue to report clarity on how to apply security principles in their work, feeling encouraged to proactively spend time addressing security issues, and recognizing the role of recent SFI improvements in lowering risk for Microsoft and our customers.

Training and enablement

We continue to evolve training and enablement to help ensure that security remains a core part of the employee experience. Trust Code, our mandatory Standards of Business Conduct training, included security content and was completed by over 99% of Microsoft full-time employees.

The latest required security training reached greater than 95% completion, with employees scoring 4.55 out of 5 on "I play a critical role in keeping Microsoft secure."

95%+

of Microsoft FTEs have completed this required training in the series.

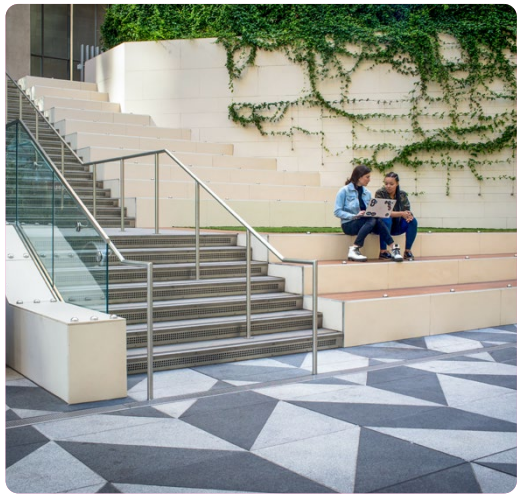

Security is a top five theme in how employees describe their impact, suggesting security ownership is widespread across Microsoft.

Customer guidance

- **Create security examples for day-to-day behaviors.** Give employees concrete examples of security behaviors mapped to their roles so they can translate priorities into clear actions and measurable impact.
- **Create clear definitions for impact and performance expectations.** Establish clear impact levels with security as a key component, giving employees and managers common language to align on expectations throughout the year.
- **Use performance tool prompts to consistently capture security impact.** Use prompts and offer optional AI features within performance tools to flag when security impact is missing, supporting more complete and consistent impact narratives.

Governance

We continue to evolve our governance model to address the threat landscape and growing regulatory complexity. Our governance approach is designed to strengthen accountability, reduce silos, and ensure cybersecurity risk decisions reflect both technical reality and business impact.

AeGIS SFI Command Center

The Artificial Generative Intelligence Safety and Security (AeGIS) SFI Command Center is a new centralized system of record for SFI execution across AI and security-critical services, providing a real-time view of ownership, risk signals, and compliance progress. It aggregates SFI KPI data across teams and surfaces actionable insights on SLA performance, blockers, and remediation trends, enabling leadership visibility and accountability at scale.

Risk frameworks

Microsoft cybersecurity risk management program operates through our Deputy CISO (DCISO) structure, aligning with global regulatory frameworks and strengthening execution through a centralized risk register. DCISOs are now explicitly accountable for reviewing, prioritizing, mitigating, and accepting risks within their domains. The risk register functions as an active execution system: risks are triaged and scored within set timeframes, mitigation plans are required after prioritization, and risk owners provide continuous updates on remediation progress. Through regular Cybersecurity Governance Council reviews, domain-level risk decisions connect directly to enterprise oversight, and DCISO-driven risk decisions are consistently reflected in executive and Board-level risk views.

Cybersecurity policy and diplomacy initiatives

Microsoft cybersecurity policy and diplomacy efforts continue to evolve alongside SFI to help shape the broader policy and governance ecosystem. As cyber risks grow in scale and complexity, we are deepening engagement with governments, regulators, and international institutions to ensure security is embedded in public policy, regulatory design, and international deliberations.

Microsoft is actively working with AI model developers, critical infrastructure operators, open-source communities, and governments around the world to ensure advanced AI strengthens defenses. Our partnership with the UK Laboratory for AI Security Research (LASR), advances AI security and joint research on threats to critical infrastructure and agentic systems.

Another core focus is advancing greater alignment across cybersecurity regulations. Fragmentation across existing and future regulations can create complexity and inefficiencies – diverting limited cybersecurity talent toward compliance, increasing costs, and ultimately undermining collective cyber defense outcomes. Microsoft has worked alongside the broader industry community to build political momentum for regulatory alignment at the Organization for Economic Co-operation and Development (OECD). Formal support from member states continues to build, while Microsoft is contributing industry perspectives on the costs of fragmentation to elevate the OECD as the forum for this work.

The Trusted Tech Alliance (TTA), announced during the Munich Security Conference on February 13, 2026, brings together 16 leading technology companies from Africa, Asia, Europe, and North America around shared commitments to support a trusted global technology stack. The Alliance is grounded in the belief that trust should be based on how technology is designed, developed, secured, and operated – not on the nationality of the provider.

16

leading technology companies brought together by TTA.

Customer guidance

- **Incorporate telemetry and security data into risk management.** Continuously ingest security telemetry (findings, vulnerabilities, threat intelligence) into your Government, Risk, and Compliance (GRC) program, so risk decisions reflect real-time signals, not periodic snapshots. Maintain a documented trail of how AI is used in your vulnerability management process and other critical security areas.
- **Align security metrics to top risks and measure mitigation progress.** Define metrics that map directly to your highest-priority risks, then track mitigation progress over time to validate whether risk is being reduced.
- **Engage with regulatory and industry efforts to reduce cybersecurity compliance fragmentation.** Participate in policy forums and cross-industry alliances working to align cybersecurity regulations across jurisdictions. Reducing regulatory fragmentation strengthens security outcomes and lowers compliance costs.

Security principles

Our approach to security continues to build on the established principles of Secure by Design, Secure by Default, and Secure in Operations, applied consistently across how products are built, delivered, and operated. Below are recent security advancements across Microsoft Azure, Microsoft 365, Windows, Surface, and the broader Microsoft Security portfolio.

Secure by Design

Security comes first when designing any product or service.

Secure by Default

Security protections are enabled and enforced by default, require no extra effort, and are not optional.

Secure Operations

Security controls and monitoring will continuously be improved to meet current and future threats.

Microsoft Azure Microsoft 365 Windows 11 Microsoft Security

Microsoft Azure

Advancements in hardware-backed trust, confidential computing, and distributed infrastructure security continue to strengthen protection for cloud, AI, edge, and disconnected environments through hardware-backed key protection, isolated execution environments, encrypted memory, attestation, hardened configurations, and verifiable security controls.

Innovation highlights:

- **Azure Integrated HSM (generally available):** Strengthens hardware-rooted trust for cloud and AI workloads through hardware-backed cryptographic key protection with increased transparency across published firmware, drivers, and software components. By reducing external key handling and operational complexity, Azure Integrated HSM improves verifiability of security controls and reduces key management risk.
- **Azure Intel® Trust Domain Extensions (TDX) Confidential VMs:** Strengthens confidential computing in Azure through Intel® TDX hardware-based isolated execution environments, encrypted memory, and remote attestation, protecting data in use and enabling sensitive workloads to run with minimal or no application changes while improving workload isolation and trust at scale.
- **Azure Local improvements:** Azure Local helps address the challenge of securing critical workloads across distributed, edge, and disconnected environments where it is difficult to standardize, provision, and continuously validate consistent security baselines. It extends Azure Arc-enabled security baselines, hardened configurations, secure provisioning, encryption at rest, and signed and verified components to local deployments – helping ensure software authenticity and integrity outside always-connected cloud environments. Azure Local Small Form Factor (SFF) deployments (preview) bring these protections to resource-constrained edge environments.

Customer guidance

- **Use hardware-backed protections for sensitive operations** to ensure cryptographic keys and critical processes remain within secure hardware boundaries. Learn more
- **Design for continuous validation at scale** by enabling ongoing verification of trusted execution environments, workload integrity, and security posture through attestation and hardware-backed isolation. Learn more
- **Apply the same secure-by-default practices** used in the cloud across local and disconnected environments by using hardened configurations, secure provisioning, and signed components to maintain consistent, verifiable security controls across distributed and resource-constrained deployments. Learn more

Security principles

Microsoft Azure Microsoft 365 Windows 11 Microsoft Security

Microsoft 365

Added baseline security mode to enable secure-by-default configuration, reducing exposure from legacy configurations across services. Also strengthened security for AI-enabled scenarios through identity-first controls and Microsoft Entra governance for enterprise AI agents, including managed identities, lifecycle management, and role-based access controls.

Innovation highlights:

- **Baseline security mode:** Microsoft 365 baseline security mode enforces secure-by-default configurations by reducing exposure from legacy settings and strengthening protections against emerging AI-enabled threats across Microsoft 365 environments.
- **Agent identity and lifecycle management:** Microsoft Entra managed identities enable scoped authentication, policy enforcement, and lifecycle governance for enterprise AI agents across Microsoft, internally developed, and third-party environments.
- **Least-privileged role enforcement for agents:** Role-based access controls for AI agents help ensure operations remain within authorized permissions, reinforcing least-privilege access and strengthening operational oversight.

Customer guidance

- **Enable baseline security mode** by reviewing recommended settings and applying secure-by-default configuration. Learn more
- **Review enterprise agent activity regularly** to confirm ownership, usage, and risk signals, and apply policy templates and automation to maintain consistent governance. Learn more
- **Assign both an owner and a sponsor to each agent** to balance operational control with clear business accountability. Learn more

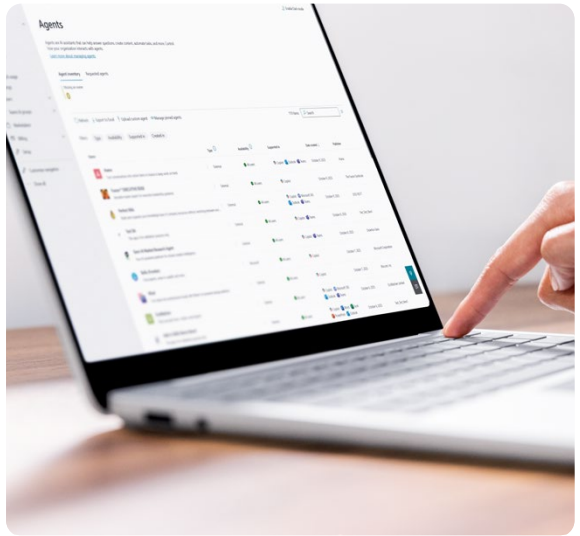

Microsoft Windows and Surface

Advanced Windows and Surface trust foundations by modernizing cryptographic and identity systems, strengthening kernel-level and hardware-rooted integrity enforcement, and introducing runtime transparency and consent controls for applications and AI agents.

Innovation highlights:

- **Modern cryptographic and authentication foundation in Windows:** Advances platform trust by combining post-quantum cryptographic capabilities for certificates, digital signatures, and TLS with the reduction of legacy authentication dependencies through the transition from NTLM to Kerberos-based identity flows. Together, these enhancements reduce attack surface, strengthen protection for data in transit, and establish a more resilient identity and communication foundation aligned with evolving security standards.
- **Windows kernel trust enforcement:** Strengthens platform integrity through a default kernel trust policy in Windows 11 (April 2026), ensuring only drivers signed through the Windows Hardware Compatibility Program (WHCP) can load by default. This improves system resilience by enforcing stricter driver trust and reducing risk from unverified kernel-level components.
- **Windows Baseline Security Mode and User Transparency and Consent:** Defines runtime integrity protections, application transparency requirements, and consent-based access controls for applications and AI agents, ensuring execution occurs within approved permissions and access to sensitive resources is visible to users and administrators.
- **Surface for Business:** New devices incorporate memory-safe Rust-based firmware and secure embedded controllers developed through the Open Device Partnership, alongside expanded Rust-based drivers and a hardware-rooted chain of trust to help reduce firmware and memory vulnerability risk. Select Surface Laptop models include an integrated, IT-managed privacy screen that helps protect sensitive on-screen information and reduce the risk of visual data exposure.

Customer guidance

- **Modernize cryptographic and authentication foundations** to reduce reliance on legacy protocols such as NTLM and strengthen resilience against evolving threats. Learn more
- **Enforce kernel-level integrity** by restricting kernel-mode execution to WHCP-signed drivers through platform trust enforcement controls. Learn more
- **Design for runtime integrity, transparency, and controlled access** by ensuring applications and AI agents operate within defined permissions and visibility into sensitive resource access. Learn more
- **Adopt secure devices** with Open Device Partnership innovations, hardware-rooted security, and integrated privacy controls to strengthen endpoint protection and reduce risk. Learn more

Security principles

Microsoft Azure Microsoft 365 Windows 11 Microsoft Security

Microsoft Security

Enabled AI-led vulnerability discovery and mitigations, AI-ready posture to reduce exposure, and AI-powered solutions to defend at scale. Added Nation State Notifications (NSN) alerts in Microsoft Defender to provide actionable intelligence about breaches attributed to nation-state actors and enabled secure access to threat intelligence across Microsoft's platform and AI systems.

Innovation highlights:

- **New multi-model AI-driven scanning system (codename MDASH):** Strengthens vulnerability management through a multi-model AI-driven scanning capability that discovers, validates, and prioritizes security issues at scale. By enabling earlier identification of critical vulnerabilities in the development lifecycle and enhancing proof-based verification and automated analysis, codename MDASH supports broader and more proactive vulnerability assessment.
- **Microsoft Exposure Management with SecureNow:** Patching, while critical, is not sufficient on its own. We have identified five areas where autonomous AI-driven attacks gain an advantage—patching, open-source software, customer source code, internet-facing assets, and baseline security hygiene. To help customers address these risks, Microsoft Security Exposure Management includes SecureNow, a new integrated experience that combines guidance and actionable capabilities to help organizations assess exposure, prioritize remediation, and strengthen security posture. This capability is available to all customers with a Microsoft Entra ID.
- **NSNs in Microsoft Defender:** Surfaces actionable nation-state, ransomware, and fraud intelligence, including attributed actor context, directly within Microsoft Defender alongside related alerts and incidents, enabling security teams to investigate and respond within existing workflows.
- **Microsoft Entra Tenant Governance:** Unifies multi-tenant environments under a centralized governance model, enabling visibility across tenants, detection of unmanaged or shadow tenants, and enforcement of consistent least-privilege access and security baselines to reduce drift and improve identity control.

Customer guidance

- **Leverage codename MDASH** to identify, validate, and prioritize exploitable vulnerabilities earlier in the development lifecycle using multi-model agentic analysis and proof-based verification, enabling faster and more targeted remediation. Learn more
- **Apply Exposure Management and SecureNow** to continuously validate security posture across identity, endpoint, customer source code, and cloud, applying baseline controls and automated remediation to reduce drift. Learn more
- **Integrate adversary-informed intelligence into security operations** to improve detection, investigation, and response decisions using contextual threat signals. Learn more
- **Strengthen identity governance across multi-tenant environments** by enforcing least-privilege access, maintaining visibility into unmanaged or shadow identity constructs, and reducing configuration drift. Learn more

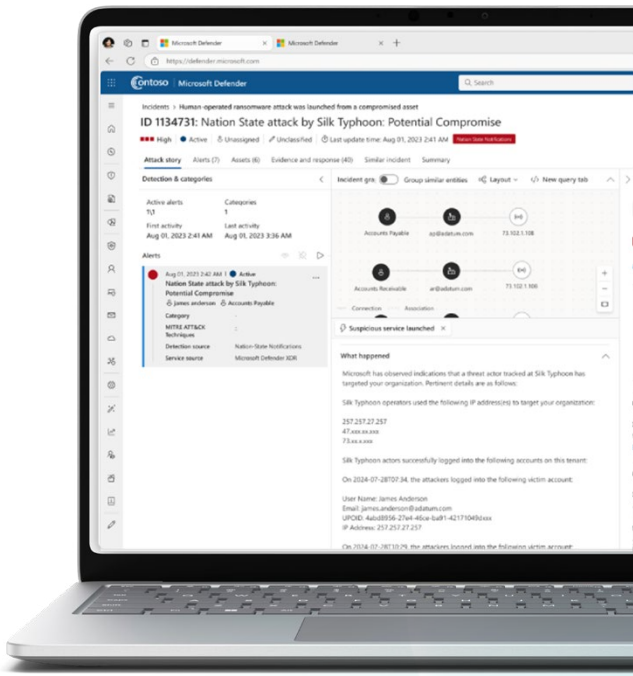

Security principles

Microsoft Implementation Insights

This section highlights how Microsoft applies SFI principles across engineering initiatives to better secure Microsoft and our customers through new capabilities.

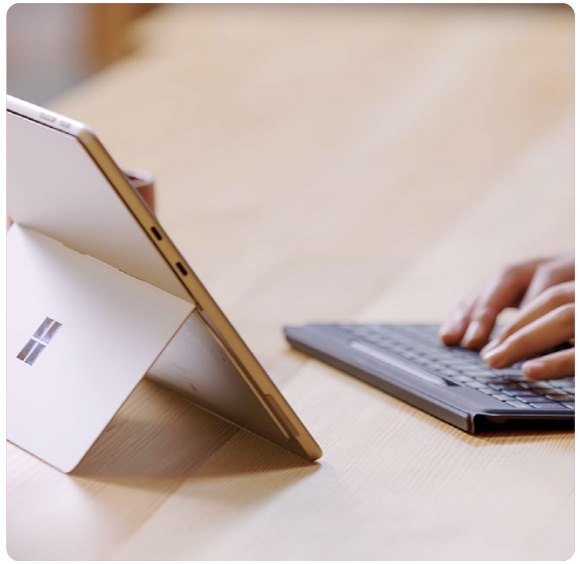

- **Microsoft Entra Tenant Governance:** Centralized tenant governance has been expanded across production and non-production tenant environments to reduce unmanaged tenant sprawl and strengthen governance consistency. Enhanced visibility, least-privilege controls, and consistent governance policies reduce identity blind spots, strengthen oversight, and enforce secure boundaries across tenant estates.
- **Agent identity integration:** Implemented internally through Agent 365 integration with Microsoft Entra ID to assign persistent identities to registered agents, enabling centralized control over authentication, authorization, conditional access, and runtime protections. A unified Agent Registry powered by Microsoft Defender, Entra, and Intune that governs unmanaged local agents across 20+ agent types, with Microsoft Purview Audit logging all activity for full traceability, strengthening visibility and least-privilege control over agent sprawl.
- **Least-privilege governance:** A role-based governance model has been implemented in Microsoft Entra to separate technical administration, business accountability, and oversight, enforcing least-privilege access by restricting agent enablement and access package management to approved roles only.
- **Agentic vulnerability discovery with codename MDASH:** Piloted internally to discover, validate, prioritize, and remediate exploitable risk across our codebases. Orchestrating 100+ specialized AI agents across an ensemble of models and processing more than 100 trillion signals per day, the system prioritizes real exploitability over theoretical noise reaching a 96.55% CyberGym benchmark score (up ~10% in under three weeks), with native Microsoft Defender integration moving vulnerability discovery and remediation from reactive scanning to production-grade defense at scale. Learn more in this blog
- **Windows Hotpatch rebootless updates:** Deployed to accelerate security compliance while eliminating disruptive monthly reboots. Orchestrated through Intune and Windows Autopatch with existing deployment rings preserved, rebootless patching reached 71% compliance within 24 hours versus 35% under standard patching (a 2x improvement), with 80% of devices compliant in under 5 days, closing exposure windows faster with minimal operational disruption.
- **Autopatch Update Readiness (AUR) reporting:** Built within Microsoft's Customer Zero servicing model to improve update visibility and remediation efficiency across large-scale device fleets. Leveraging Windows Autopatch telemetry across 450,000+ devices, it surfaces deployment issues and readiness gaps, reducing readiness analysis time by 40% and lowering operational overhead.
- **Copilot+ PC at enterprise scale:** Established Copilot+ PCs as the default for primary refresh and new hire devices, scaled from a 200-device pilot to 56,800+ globally via zero-touch Autopilot/Intune provisioning and secure-by-default configuration. On-device NPU processing reduces data transfer exposure, returning 2+ hours/week per employee (~3.7 million hours/year) and lifting NSAT by +4, demonstrating secure, AI-enabled endpoint architecture at enterprise scale.

Section 4

Engineering pillars

Each objective represents a significant body of work, most will take years to complete.

The work under each objective evolves continuously as the risk landscape changes, new attack vectors emerge, and foundational work matures.

Three objectives have reached their target state. The remainder are at varying stages of maturity, with the highest-risk, most critical assets prioritized first. Platform engineering practices and AI-driven automation have scaled execution and reduced toil across the portfolio.

Progress legend	Total objectives
100% Complete	3
95–99% Nearing completion	3
66–94% Significant progress	12
33–65% Progress	3
0–32% Initial progress	3
No percentage	4

1. Protect identities and secrets

We are strengthening cryptographic safeguards, reducing exposure from legacy authentication, and have completed the rollout of phishing-resistant, credential-safe defaults at scale. Our work is focused on durable, hardware-backed key protection, modern authentication libraries, and stronger credential lifecycle management. Together, these investments have reduced token-theft risk, improved traceability, and strengthened resilience to identity-based attacks.

1. Protect cryptographic signing keys

PR.AA.04 PR.PS.01

Completed migration to Azure Confidential Computing for 100% of Microsoft Entra ID token signing keys in Azure Public and US Gov cloud, ensuring all cryptographic operations occur within secure execution environments.

2. Adopt standard SDK's for identity

PR.AA.01-03

Microsoft Authentication Library (MSAL) has been updated to support protocol enhancement to disrupt authorization code social engineering campaigns, hardening targeted Microsoft applications against identity-based social engineering attacks, and released to customers. 93% of Entra traffic now flows through the standard SDK, replacing ad-hoc token-handling with a single validated library. Legacy authentication protocols continue to be retired across Office and Microsoft Entra workloads to reduce attack surface.

3. Phishing-resistant MFA

PR.AA.01 PR.AA.03

Phishing-resistant MFA enforcement has reached broad maturity with 99.97% user and device coverage, including full enforcement on macOS and expanded Linux support.

Closing objective: This objective is transitioning to steady-state operations with ongoing monitoring to prevent regression. In response to attackers pivoting past credential theft, defense efforts now focus on OAuth consent phishing, open redirect abuse, and redirect URI hijacking to close emerging identity attack vectors.

4. Safe Secrets Standard

PR.AA.04 PR.DS.01 PR.PS.01

We continue to strengthen credential security across Microsoft by expanding secure-by-default authentication and reducing reliance on long-lived, reusable credentials. New services increasingly use modern identity-based authentication by default, while improvements to credential management, automation, and access controls reduce the risk of credential theft. We are also expanding the use of short-lived, workload-bound credentials that help prevent stolen credentials from being reused outside their intended environment, further improving the security and resilience of Microsoft services.

5. Stateful validation for identity tokens

PR.AA.03 DE.AE.02 DE.AE.07

Microsoft Entra now includes linkable identifiers in access tokens to correlate activity across Microsoft Entra and major Microsoft services audit logs (including Exchange Online, Microsoft Graph, SharePoint Online, and Teams). This improves investigation speed and containment for compromised sessions by tying downstream actions back to a specific authentication event, with coverage continuing to expand as additional Microsoft Entra services populate audit logs with these identifiers.

6. Fine-grained key partitioning

PR.AA.04 PR.AA.05

Datacenter key partitioning is now fully deployed across all cloud environments for baseline identity operations. New Azure cloud environments are also onboarding secure-by-default with split keys, strengthening isolation boundaries and reducing systemic cryptographic risk as the footprint expands.

7. Quantum-safe PKI systems

PR.AA.04

Planning for post quantum PKI has advanced from initial assessment to a defined transition approach. Microsoft has validated key drivers and scope for this work – including certificate services, identity platforms, and critical service trust chains – and is now developing an implementation plan for phased rollout.

2. Protect tenants and isolate production systems

We continue to reduce blast radius across the Microsoft cloud by eliminating legacy systems, securing tenant boundaries, enforcing least privilege, and hardening against lateral movement. Secure-by-default tenant creation with drift detection now operates at scale, establishing isolation at provisioning rather than restoring it through later cleanup. Together, these changes make tenant isolation more durable, consistent, and Secure by Default.

1. Remove legacy systems that risk security

ID.AM.02 ID.AM.08 PR.PS.01

Retirement of Azure Service Manager is complete across all remaining workload classes: classic compute, storage (~800,000 classic SQL storage accounts), networking, App Service Environment v1/v2, and API Management v1. Migration to Azure Resource Manager is at 100%, up from 98% in November 2025, eliminating classic administrator-role and management-certificate risks to further reduce the risk of credential replay and lateral movement. Satellite-tenant cleanup matured: 15,500 aged or unused tenants retired this cycle (down from 560,000, reflecting fewer discoveries, shifting the effort from broad cleanup to governance). Automated lifecycle checks scale permanent deletion of tenants with minimal manual effort.

2. Secure all tenants and their resources

ID.AM.02 PR.AA.05 PR.IR.01 PR.PS.01

Tenant resources are continuously hardened against baseline drift and compliance regression. Inventory and lifecycle coverage spans all tenant classes: production, productivity, auxiliary, and ephemeral – so every tenant is governed from creation through retirement. The ephemeral tenant fleet is aggressively lifecycle-managed with a default 90-day expiration.

This high-frequency rotation model significantly reduces residual risk by ensuring that test artifacts, credentials, and transient configurations are systematically purged, minimizing attack surface and preventing long-term exposure.

3. High security for Microsoft Entra ID apps

ID.AM.08 PR.AA.01 PR.AA.05 PR.IR.01

Application governance now operates as sustained, machine-scale enforcement. We decommissioned 1.4 million unused Microsoft Entra applications (an increase of 176,198 since November) and continued stripping unnecessary critical-privilege Graph directory permissions from production apps, shrinking the lateral-movement surface. Currently, 86.45% of multi-tenant applications declare an allowed-tenants list under enforced policy, and centrally applied configuration-time policy keeps every app on an authorized federated-credential issuer. Network restrictions on managed identities tripled from 5.9 million in November 2025 to 18.1 million today, and 99.99% of these managed identities are blocked from service-to-service authentication into production from non-production or non-Microsoft IP ranges.

4. Eliminate identity lateral movement

PR.AA.04 PR.IR.01 DE.CM.01

Eliminating lateral-movement pivots between tenants, environments, and clouds is one of the highest-leverage ways to reduce attack paths. Cross-boundary credential isolation advanced from 98% in November 2025 to 98.7%. In our productivity environment, guest invitation pathways are being tightened so external access is tied to clear ownership and lifecycle governance, and unmanaged guest access has been reduced by 228,000 with further cleanup in progress.

5. Continuous least-privilege enforcement

PR.AA.05

Emergency access is moving from persistent, privilege identities to on-demand, dual-authorization workflows. Dual-Key Break Glass (DKBG) controls now enable two-person authorized access to critical foundational services without standing privilege accounts. For JIT/PIM disruption scenarios, privilege access to Azure Resource Manager can be provisioned through RBAC administrator-mediated, dual-authorization role assignments for designated service owners – preserving operational resiliency without persistent exposure. All emergency access is fully audited and automatically generates security incident records – rollout began in April and is targeted for completion by end of July 2026 (eliminating 29,000 existing break-glass accounts). The high-privilege application reduction program is now 81% complete, with enforcement expanding across a broader application footprint to further shrink persistent, elevated access.

6. Secure devices used for access

ID.AM.01 ID.AM.02 PR.AA.01 PR.AA.06 PR.IR.01

All production access across Microsoft tenants now requires a locked-down endpoint. We have deployed 125,979 production-ready locked-down devices for production access (up from 112,000 in November 2025), as more workload classes move behind hardened-device gates. In productivity tenants, trusted-endpoint enforcement has expanded to additional high-risk scenarios, and 56,581 of our highest-risk users now operate exclusively on locked-down endpoints, up from 44,516 since November 2025.

3. Protect networks

Network security across Microsoft production environments is built on a mature, consistently enforced foundation. Asset inventory and lifecycle management now run at near-complete, steady-state coverage. Network isolation continues to scale, reducing direct back-end access, eliminating unnecessary internet-exposed endpoints, and tightening controls that limit lateral movement and shrink attack surface. Customer capabilities are also improving secure-by-default configurations and workload isolation. Ongoing work is focused on consistent isolation enforcement, broader adoption, and long-term resilience.

1. Inventory and security standards

ID.AM.01 ID.AM.03 PR.AA.05 PR.PS.01 PR.PS.04

Baseline network security capabilities for Microsoft production environments have reached a stable and repeatable operating state. Centralized telemetry and control systems now provide effectively complete visibility of production network assets through a complete device lifecycle management process. Core controls – including centralized authentication, authorization and accounting, IPv4/IPv6 access control lists, and per-device keys – are consistently enforced across production network devices, establishing a uniform baseline that limits lateral movement, contains the blast radius and strengthens accountability.

Closing objective: With asset inventory and lifecycle management operating at effectively complete coverage, these capabilities are now maintained through durable, secure-by-default provisioning and continuous validation processes. This objective has reached its intended end state and transitions to steady-state operations focused on durability and regression prevention.

2. Network isolation

PR.IR.01 PR.PS.01 PR.PS.04

The network segmentation strategy is advancing toward broad, consistent enforcement across production environments. Policy-driven isolation capabilities are now operating at scale, with adoption reaching 4.36 million resources in learning mode and approximately 1 million in enforced mode.

Direct exposure pathways are being systematically reduced, with public access removed from 732,000 resources and ongoing efforts to minimize internet-facing endpoints. We continue to eliminate unnecessary internet-facing endpoints, direct back-end access pathways are systematically reduced, and access policies are being refined by replacing permissive network rules with tightly scoped controls aligned to least-privilege principles.

Segmentation coverage is extending from critical and high-value services to all Microsoft services, strengthening containment boundaries, and reducing the potential for unauthorized lateral movement. Continued progress is focused on completing enforcement at scale and embedding these controls into standard provisioning and operational practices to ensure long-term consistency and resilience.

3. Secure customer cloud networks

PR.IR.01 PR.PS.01

Ongoing effort

Capabilities to better secure customer cloud networks continue to advance, with platform features enabling stronger, more consistent isolation of customer environments. Azure now provides integrated mechanisms for defining logical isolation boundaries around PaaS resources, supporting secure-by-default deployment patterns. This allows customers to benefit from resource isolation while ensuring the principle of least privilege is preserved by default through traffic pattern learning windows.

Ongoing efforts focus on expanding adoption of these capabilities and aligning them to secure deployment baselines, helping customers consistently reduce exposure risk while strengthening network isolation and perimeter defenses.

4. Protect engineering systems

Engineering security is now focused on enforcing controls at scale. Asset inventory, access governance, and Zero Trust source-code protections are operating in steady-state, while secure pipelines, signing-key protection, and supply chain controls continue to reduce risk. Remaining work is centered on durability and legacy cleanup, including classic pipelines and non-human identities.

1. Complete software asset inventory

ID.AM.02

Engineering asset visibility and access governance have stabilized into a mature, consistently operating capability across development environments. This model integrates asset inventory and ownership into a unified governance control plane, enabling consistent policy enforcement and precise incident scoping.

Closing objective: With coverage at a mature level and governance controls operating effectively at scale, this objective has achieved its intended outcome and transitions to steady-state operations. Ongoing efforts focus on sustaining coverage growth, maintaining policy consistency, and ensuring durability through continuous validation and regression prevention.

2. Zero Trust for secure code access

PR.AA.01-06 PR.DS.01-02 PR.DS.10

Zero Trust for source-code access is now broadly enforced across engineering environments. Code changes are governed through policy-based checks and approvals, improving auditability and reducing access risk. Most metrics are substantially complete; remaining gaps are concentrated in automated and non-human identity workflows, with current measurements reflecting improved visibility rather than weaker controls. A new permissions service, which manages least-privileged access, is currently adopted across 71% of high-value production repositories.

3. Secure code deployment

PR.AA.04 PR.PS.06

Secure code deployment controls remain strong. Durable hardening of Azure DevOps – including the use of hardware-backed protection for signing keys and expanded adoption of confidential computing – continues to strengthen software integrity and supply-chain security. Recent metric variation reflects expanded measurement scope, with final attribution under review.

4. Standardize secure development pipelines

PR.DS.01-02 PR.DS.10 PR.IR.01 PR.PS.06

Secure development pipelines continue to reduce risk at scale, with 93% of critical and high-value build pipelines and 90% of release pipelines using centrally governed templates. Across approximately 80,000 monitored pipelines, centrally governed templates enforce consistent security controls, including blocking known malicious endpoints. Policy maturity continues to improve, with broad isolation of network endpoints across pipelines. Remaining work focuses on broader template adoption and stricter enforcement.

5. Protect the software supply chain

GV.SC.01-04 GV.SC.07 GV.SC.09 ID.RA.10 PR.PS.06

Supply-chain security has shifted from reactive remediation to governed prevention. Microsoft has remediated more than 550,000 instances of critical and high-risk open-source vulnerabilities, and automated container patching now addresses about 3 million vulnerability instances each month. PR-time blocking of foreign checked-in binaries has also improved remediation effectiveness, increasing fix rates from 13% to 47% by preventing risky components earlier in the workflows. During the Shai Hulud v2 supply-chain attack, malicious packages were blocked before reaching internal systems. PR-time enforcement and dependency validation help stop risk before entering the development lifecycle.

5. Monitor and detect threats

Production monitoring and detection capabilities are established and operating at scale across production environments. Security logging from production nodes is supported with two-year retention. Standardized security logging continues to expand, and we are moving towards LLM-based approach for automatically collecting the service logs in a standardized format to raise the quality, consistency, and completeness. Centralized access to security data continues to mature, with most log categories meeting retention standards and governance processes driving gap remediation. Detection capabilities are also advancing beyond signature-based methods towards more behavior-driven approaches. AI-assisted detection authoring helps accelerate attack pattern identification, increase authoring throughput, and improve deployment velocity.

1. Complete production infrastructure inventory

ID.AM.01-02 ID.AM.05 PR.PS.04

Significant progress has been made in establishing centralized visibility across production infrastructure assets. Most assets including network devices, physical servers, and virtual machines are actively tracked through a unified inventory system. From these assets, security telemetry is centrally collected and retained for two years, providing a consistent foundation for long-term threat detection, audit readiness, incident response, and compliance. Our efforts are shifting to ensure freshness and accuracy of inventory systems and regression detection.

2. Security log retention standards

PR.PS.04 DE.CM.01-03 DE.CM.06 DE.CM.09

Standardization of security logging continues to expand across engineering systems, with more than 81% of services now emitting critical security logs in a standard format and retained for two years. Automated collection capabilities are reducing operational burden on service teams while improving coverage for investigations and threat detection. In parallel, LLM-based analysis is enhancing log quality and coverage by scanning code repositories, identifying telemetry gaps, and partnering with services teams to improve log comprehensiveness.

3. Centralized access to security logs

PR.PS.04 DE.AE.03

Centralized access to security telemetry continues to mature. Standardized retention policies and robust, scalable access controls enable security teams to efficiently retrieve and analyze logs for investigations and threat hunting. Major log categories now meet minimum retention standards, supported by structured processes to identify and address remaining telemetry gaps within defined SLAs. A centralized tracking framework further strengthens visibility and accountability, ensuring that identified gaps are prioritized and resolved. Ongoing efforts focus on improving consistency of access controls and strengthening governance processes to ensure durable, scalable access to security data.

4. Rapid anomaly detection and response

DE.AE.02-04 DE.AE.06-07 RS.CO.02 RS.MI.01-02

100+ new detections

Detection capabilities continue to evolve as we are moving beyond signature-based approaches towards behavior and baseline-driven threat detection. More than 100 new detections have been introduced, alongside enhancements to the existing 250+ detections, strengthening coverage of high-priority attack techniques across Microsoft infrastructure. Investments in detection engineering are improving both scale and speed, with AI-assisted methods increasing authoring throughput and accelerating deployment. Continued collaboration with threat research and Red Team operations is enhancing detection quality and relevance.

6. Accelerate response and remediation

Response and remediation capabilities continue to mature into a durable operating model that reduces time to mitigation and improves consistency at scale. Our focus is on strengthening vulnerability identification, triage, assessment, routing, and mitigation across cloud and AI environments, while expanding automation and tactical mitigation options to scale with increasing volume and shorten exposure windows. These capabilities improve speed, clarity, and coordination – enabling faster containment, more predictable remediation outcomes, and clearer customer guidance during security events.

1. Accelerate vulnerability mitigation

ID.RA.01 RS.AN.03 RS.MA.02-03 RC.RP.01-02

During this period, we have observed a meaningful industry-wide shift in AI capabilities, enabling the discovery of vulnerabilities at a scale and speed that was previously not possible through traditional human and tooling approaches alone. Microsoft has actively partnered with leading AI and security experts across the industry, including through initiatives such as Project Glasswing, to advance our ability to rapidly identify and mitigate vulnerabilities using these emerging techniques.

OpenAI and Microsoft have worked closely for years, and that partnership continues to deepen as both organizations apply advanced AI to cybersecurity. On April 22, 2026, Microsoft and Open AI announced an expanded collaboration between model development and real-world defense operations, through programs such as Trusted Access for Cyber, to strengthen protections for shared customers. This shift is already reflected in our operational outcomes, with a significant increase in security updates released in the June 2026 Patch Tuesday release. This includes many vulnerabilities discovered by codename MDASH. We expect this trend to continue as AI capabilities advance.

In addition to security updates, we continue to invest in automation and tactical mitigation capabilities that enable rapid response. This includes Microsoft Defender detections, information sharing with our Microsoft Active Protection Program partners, and novel mitigations. As an example, in response to a client-side vulnerability being actively exploited against customers, we deployed a mitigation capability without a customer update, protecting customers on the latest supported versions in under a day.

This area continues to evolve quickly, and we remain committed to publicly sharing our learning and recommendations so customers can better prepare for and defend against these changes. Read more here: Strengthening secure software at global scale: How MSRC is evolving with AI

2. Transparency of cloud vulnerabilities

ID.RA.01 RS.CO.02-03

1,989 CVEs + 93 No action CVEs

We continue to advance transparency in vulnerability management through transparent engagement with security researchers, consistent use of industry-standard frameworks, and expanded disclosure practices. The publication of CVEs enriched with CWE and CPE annotations is now an established practice, enabling customers and the broader ecosystem to better understand root causes and affected platforms. Support for machine-readable CSAF and VEX formats help customers automate consumption of Microsoft vulnerability and release information.

Transparency efforts emphasize completeness as well as volume. Publishing CVEs that do not require customer action provides a more accurate view of the threat landscape and reinforces trust through openness. Investments in enhancing researcher recognition (including bounty) and coordinated research initiatives like Zero Day Quest further strengthen early discovery, particularly in high-risk boundary areas between services.

Links:

Zero Day Quest 2026: $2.3 million awarded for vulnerability research

Congratulations to the top MSRC 2026 Q1 security researchers!

3. Enhance public messaging and engagement

GV.OV.01-03 GV.RR.01-04 RS.CO.03-04 RC.CO.03-04

Ongoing effort

Customer communication during security incidents is now mature and repeatable. The Customer Security Management Office (CSMO) coordinates customer facing communications across major incidents, using defined playbooks and cross-team alignment to deliver clear, actionable guidance faster.

Nation State Notifications (NSN) deliver actionable intelligence on breaches linked to nation-state actors, major ransomware campaigns, and fraud operations, including information about affected organizations and the threat actor involved. By surfacing these alerts directly in Microsoft Defender, alongside related alerts and incidents, NSN gives security teams timely visibility into high-impact threats and helps them respond faster and with greater confidence than siloed call-center notifications.

Link:

Delivered by MSTIC analysts

Appendix: Cybersecurity frameworks in action

SFI is a cybersecurity program tailored to Microsoft, focused on improving our security posture and revolutionizing the way we design, build, test, and operate our products and services, to achieve the highest security standards. This is done in part by stringently applying Zero Trust principles.

To help customers understand how to map our progress with their own cybersecurity programs, we have mapped SFI progress updates to the NIST CSF and provide Zero Trust implementation guidance.

The NIST CSF offers comprehensive coverage across the security landscape, enabling organizations to build cybersecurity programs that are both scalable and adaptable, the guidance is divided into the following six categories:

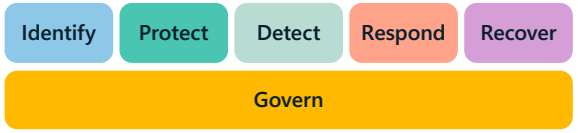

The following image provides a high-level mapping of the SFI progress report sections to the NIST CSF Categories, more detailed information will be provided within each section of this report.

NIST CSF Mapping Framework

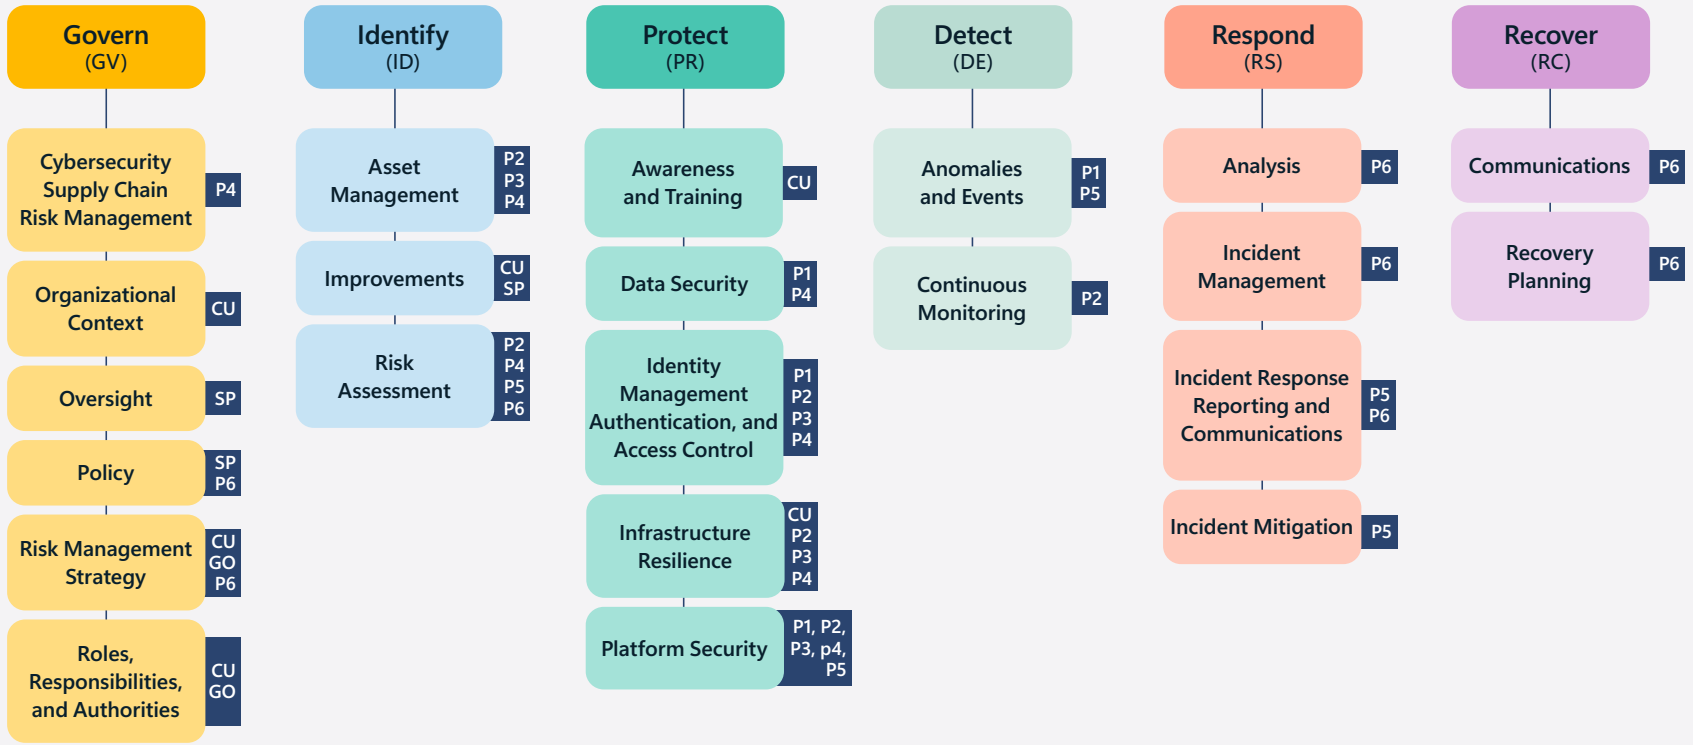

It is important to note that SFI represents just one part of our broader internal commitment to NIST CSF and other regulatory frameworks. Beyond SFI, we implement additional controls, processes, and technologies that reinforce our principle of Secure by Design. This layered approach helps ensure customers benefit not only from our secure infrastructure but also from the maturity and extensibility of our security practices.

CU Culture
GO Governance
SP Security principles
P1 Identity
P2 Tenants
P3 Networks
P4 Engineering
P5 Detect
P6 Respond

NIST CSF category descriptions

This section explains the NIST CSF categories used in mapping to SFI. Each NIST CSF function contains several categories, and each category contains at least one sub-category. To navigate these, we follow the standard naming format: function.category-sub-category

Example: The Governance (GV) function has a category of Organizational Context (OC), and the first sub-category (01) states "The organizational mission is understood and informs cybersecurity risk management". This would be written as GV.OC-01

Govern (GV)

Cybersecurity Supply Chain Risk Management (GV.SC)
Cyber supply chain risk management processes are identified, established, managed, monitored, and improved by organizational stakeholders.

Organizational Context (GV.OC)
The circumstances – mission, stakeholder expectations, dependencies, and legal, regulatory, and contractual requirements – surrounding the organization's cybersecurity risk management decisions are understood.

Oversight (GV.OV)
Results of organization-wide cybersecurity risk management activities and performance are used to inform, improve, and adjust the risk management strategy.

Policy (GV.PO)
Organizational cybersecurity policy is established, communicated, and enforced.

Risk Management Strategy (GV.RM)
The organization's priorities, constraints, risk tolerance and appetite statements, and assumptions are established, communicated, and used to support operational risk decisions.

Roles, Responsibilities, and Authorities (GV.RR)
Cybersecurity roles, responsibilities, and authorities to foster accountability, performance assessment, and continuous improvement are established and communicated.

Identify (ID)

Asset Management (ID.AM)
Assets (e.g., data, hardware, software, systems, facilities, services, people) that enable the organization to achieve business purposes are identified and managed consistent with their relative importance to organizational objectives and the organization's risk strategy.

Improvements (ID.IM)
Improvements to organizational cybersecurity risk management processes, procedures and activities are identified across all CSF Functions.

Risk Assessment (ID.RA)
The cybersecurity risk to the organization, assets, and individuals is understood by the organization.

Protect (PR)

Awareness and Training (PR.AT)
The organization's personnel are provided with cybersecurity awareness and training so that they can perform their cybersecurity-related tasks.

Data Security (PR.IR)
Data are managed consistent with the organization's risk strategy to protect the confidentiality, integrity, and availability of information.

Identity Management Authentication, and Access Control (PR.AA)
Access to physical and logical assets is limited to authorized users, services, and hardware and managed commensurate with the assessed risk of unauthorized access.

Infrastructure Resilience (PR.IR)
Security architectures are managed with the organization's risk strategy to protect asset confidentiality, integrity, and availability, and organizational resilience.

Platform Security (PR.PS)
The hardware, software (e.g., firmware, operating systems, applications), and services of physical and virtual platforms are managed consistent with the organization's risk strategy to protect their confidentiality, integrity, and availability.

Detect (DE)

Anomalies and Events (DE.AE)
Assets are monitored to find anomalies, indicators of compromise, and other potentially adverse events.

Continuous Monitoring (DE.CM)
Anomalies, indicators of compromise, and other potentially adverse events are analyzed to characterize the events and detect cybersecurity incidents.

Respond (RS)

Analysis (RS.AN)
Investigations are conducted to ensure effective response and support forensics and recovery activities.

Incident Management (RS.MA)
Responses to detected cybersecurity incidents are managed.

Incident Mitigation (RS.MI)
Activities are performed to prevent expansion of an event and mitigate its effects.

Incident Response Reporting and Communications (RS.CO)
Response activities are coordinated with internal and external stakeholders as required by laws, regulations, or policies.

Recover (RC)

Communications (RS.CO)
Restoration activities are coordinated with internal and external parties.

Recovery Planning (RC.RP)
Restoration activities are performed to ensure operational availability of systems and services affected by cybersecurity incidents.

Culture, Governance, and Security principles NIST CSF mapping

Culture

Focus area	CSF categories	CSF Sub-categories
Security is the top priority for Microsoft	Govern Policy Roles, Responsibilities, and Authorities	GV.PO-01 GV.PO-02 GV.RR-01 GV.RR-02 GV.RR-03 GV.RR-04
	Identify Improvement	ID.IM-01 ID.IM-02 ID.IM-03 ID.IM-04
Accountability	Govern Organizational Context	GV.OC-01 GV.OC-02 GV.OC-03 GV.OC-04
Education	Protect Awareness and Training	PR.AT-01 PR.AT-02

Governance

Focus area	CSF categories	CSF Sub-categories
Deputy CISOs	Govern Roles, Responsibilities, and Authorities	GV.RR-01 GV.RR-02
Cybersecurity Governance Council	Govern Roles, Responsibilities, and Authorities	GV.RR-01 GV.RR-02
Oversight of risk registers	Govern Risk Management Strategy	GV.RM-01 GV.RM-02 GV.RM-03 GV.RM-04 GV.RM-05

Security principles

Focus area	CSF categories	CSF Sub-categories
Security by Design	Protect Platform Security	PR.PS-01 PR.PS-01
Security by Default	Protect Platform Security	PR.PS-01 PR.PS-01
Security Operations	Identify Improvement	
	Protect Platform Security	PR.PS-04 ID.IM-03

Microsoft

Secure Future Initiative (SFI)

July 2026 progress report