



Secure Future Initiative (SFI)

Security above all else.

Contents

Summary

Introduction	03
Report background and methodology	04
Progress highlights	05

Culture and Governance

Culture	08
Governance	09

Security Principles

Secure by Design	11
Secure by Default	12
Secure Operations	13

Engineering Pillars

1. Protect identities and secrets	15
2. Protect tenants and isolate production systems	17
3. Protect networks	19
4. Protect engineering systems	21
5. Monitor and detect threats	23
6. Accelerate response and remediation	25
Appendix: CSRB Mapping	27

Introduction

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

It is the largest cybersecurity engineering project in history and the most extensive effort of its kind at Microsoft. To date, we have invested the equivalent of 34,000 engineers working full-time for 11 months to mitigate risk and improve security for Microsoft and our customers.

In September 2024, we published our first progress report, followed by the November 2024 update, which provided insights into our implementation and customer best practices. In this second progress report, we share detailed insights into the progress made since May 2024.

We have activated our culture to foster a security-first mindset in every employee at every level. We established new holistic governance structures to address cybersecurity risk and compliance enterprise wide. Teams across Microsoft are delivering innovations aligned with our security principles of Secure by Design, Secure by Default, and Secure Operations to better protect customers, including a new Secure by Design UX Toolkit, new capabilities like the Microsoft 365 Copilot Control System, and new controls that thwarted \$4 billion in fraud attempts.

We have made progress in every engineering pillar and objective. Each objective represents a significant body of work, most will take years to complete. Out of 28 objectives, 5 are nearing completion, 11 have made significant progress, and we continue to make

progress against the rest. We sought to prioritize the highest risks, most critical assets, and used platform engineering practices to scale the work and reduce toil. Key progress themes include:

- Stringently applying zero trust.
- Enforcing security standards at scale—phishing-resistant multi-factor authentication (MFA) and Microsoft Azure Managed Identities—to protect credentials and secrets.
- Ensuring a complete asset inventory.
- Implementing fine-grained permissions and segmentation.
- Testing controls and detections through continuous Red Team operations.
- Streamlining operations so we can respond more rapidly without sacrificing security.

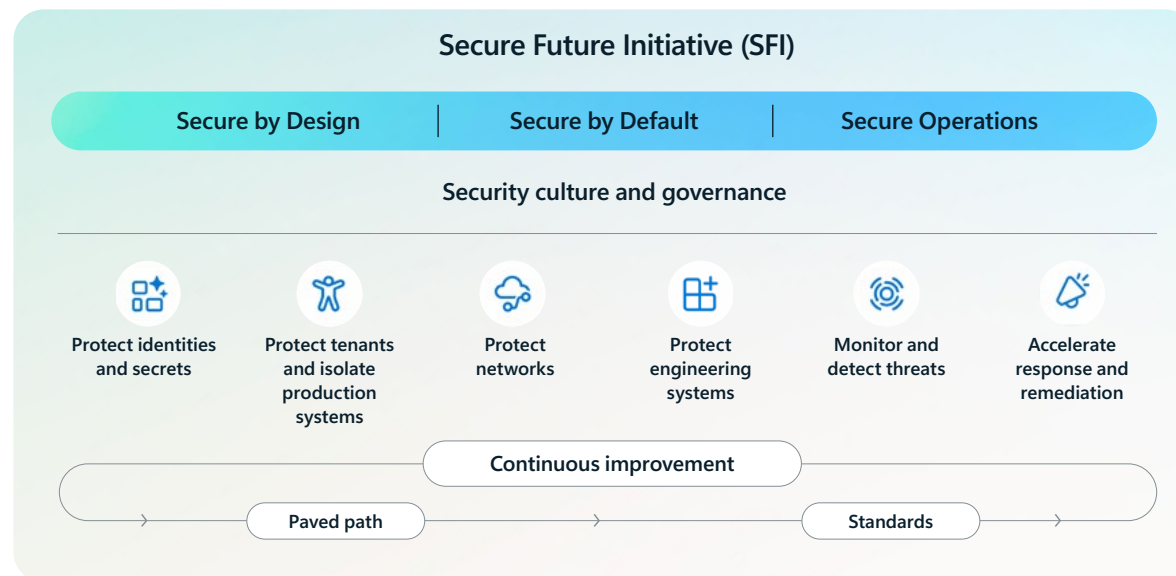
Insights and learnings from this progress inform ongoing innovations in our Microsoft Security portfolio – Microsoft Entra, Microsoft Defender, and Microsoft Purview – that helps better protect customers and Microsoft.

Our progress will not be linear. The threat landscape will continue to evolve, resulting in new vulnerabilities and security incidents. Technology will advance, creating new ways to improve security and new issues to address. Each of these is an opportunity to work with our customers and the industry to strengthen our collective defenses.

We are grateful for the partnership and look forward to innovating together for a safer future.

Security above all else.

“The threat landscape will continue to evolve, resulting in new vulnerabilities and security incidents. Technology will advance, creating new ways to improve security and new issues to address. Each of these is an opportunity to work with our customers and the industry to strengthen our collective defenses. We are grateful for the partnership and look forward to innovating together for a safer future.”



Report background and methodology

Culture, governance, and security principles:

For each of these areas we share examples of progress made.

Engineering pillars:

In May 2024, we [announced](#) the expansion of SFI to include six prioritized engineering pillars and 28 aligned objectives. Engineering owners were assigned to each engineering pillar. They established an initial body of work and scope to advance the objective, according to Microsoft priorities. The work was then translated into standards and key results.

Since that time, pillar owners have expanded on the initial body of work, adding new scope, standards or key results based on changes in the threat landscape, risk prioritization from the Cybersecurity Governance Council, and learnings from ongoing implementation.

In September 2024, we disclosed progress made in each engineering pillar against specific standards. In this April 2025 progress report, we are disclosing progress made towards completing an objective.

We calculated progress by recording the completed percentage of standards and key results, those defined initially and any additions. Some objectives focus on a limited scope of products. We do not disclose the full body of work we are undertaking to advance an objective (such as standards and key results) but do share examples of progress made against specific standards and key results. In a few cases, progress is not reflected as percent complete.

As noted in the introduction, each objective represents a significant body of work to improve security in a specific area and reduce risk for Microsoft and our customers. Some objectives will take several years to complete. Others, like our work in post-quantum cryptography and the orderly sunseting of cryptographic techniques as they age, will take much longer.




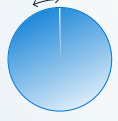
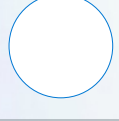
Progress made will not be linear. Standards and key results may change over time due to emerging risks or technology. Progress may regress due to an increase in the body of work prioritized to complete the objective or advance rapidly due to innovation.

Measuring investments in cybersecurity risk reduction and communicating progress in a durable manner across the magnitude of engineering activities SFI covers is an imperfect science. While this report cannot address every nuance of Microsoft efforts, it is intended to communicate SFI progress in a way that demonstrates our Microsoft security investments and underscores important advancements relevant for the awareness and activities of our customers and the industry.

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

Out of 28 objectives, 5 are nearing completion, 11 have made significant progress, and we continue to make progress against the rest. We prioritized the highest risks, most critical assets, and used platform engineering practices to scale the work, drive clear prioritization, and reduce toil.

Progress legend

Icon	Description	Total objectives
	0-32% Initial progress	5
	33-65% Progress	3
	66-94% Significant progress	11
	95-99% Nearing completion	5
	No percentage	4

Progress highlights

This section includes a summary of progress for culture, governance, and our security principles:

- [Secure by Design](#)
- [Secure by Default](#)
- [Secure Operations](#).

Culture

To foster a security-first mindset in every employee, we introduced the Security Core Priority and continue to improve our training around security.

- As of December 2024, every employee had a Security Core Priority and discussed individual impact with their manager during performance check-ins.
- 50,000 employees have participated in the Microsoft Security Academy to improve their security skills.
- More than 99% of employees completed the Security Foundations and Trust Code courses, increasing awareness of cybersecurity best practices.

[See page 8 →](#)

99%

of employees completed 2024 Security Foundations and Trust Code courses.

Governance

To improve overall cybersecurity risk and compliance we have expanded and refined our governance oversight and risk management processes.

- We added a Deputy CISO for Business Applications and consolidated responsibility for Microsoft 365 and Experiences and Devices (E+D) into one role.
- All 14 Deputy CISOs have completed a risk inventory and prioritization for their product or function.

[See page 9 →](#)

New

Deputy CISO for Business Applications appointed.

Principles

Teams across Microsoft are delivering rapid innovation aligned to our security principles to better protect our customers and Microsoft. These include:

- A new design review process for AI development at Microsoft, led by the Artificial Generative Intelligence Safety and Security Organization.
- A new Secure by Design UX Toolkit, which has been tested with 20 product teams and rolled out to 22,000 employees.
- Eleven capabilities from Microsoft Azure, Microsoft 365, Microsoft Windows, and Microsoft Security that help improve security by default.
- Consistent application of secure operations across all aspects of AI systems, as described in our Responsible AI Annual Transparency Report.
- New policies, behavioral-based detection models, and investigation methods that thwarted \$4 billion in fraud attempts.

[See pages 11-13 →](#)

Progress Highlights continued

Engineering

We continue to make progress in every pillar and objective. Out of 28 objectives, 5 are nearing completion, 11 have made significant progress, and we continue to make progress against the rest. As a result of SFI our platforms and services are more secure and we have improved our ability to detect and respond to threats.

Protect identities and secrets:

We have improved identity security for Microsoft services and customers.

- New defense-in-depth protections for Microsoft Entra ID and Microsoft Account (MSA) token signing keys already stored in hardware-based security modules. The Microsoft Account (MSA) signing service has been migrated to Azure confidential VMs.
- 90% of identity tokens from Microsoft Entra ID for Microsoft apps are validated by one consistent and hardened identity Software Development Kit (SDK).
- To mitigate risk from advanced attacks, 92% of employee productivity accounts now use phishing-resistant multifactor authentication (MFA).

Protect tenants and isolate production systems:

We continue to remove legacy/unused resources, and increase isolation, to reduce the risk of lateral movement.

- We transitioned more than 88% of resources to Azure Resource Manager, removed a total of 6.3 million tenants (an additional 550,000 since September), and all new tenants are now automatically registered in our security emergency response system.
- We use an automated lifecycle management solution for all Microsoft Entra ID applications in the production environment.
- Authentication to 4.4 million production environment managed identities is now restricted to specific network locations, further protecting these critical assets.

Protect networks:

Progress made against all objectives has improved the security of our network and delivered new innovations to help customers protect their networks.

- More than 99% of network assets have been inventoried and use enhanced security standards.
- We continue to add additional layers of defense in depth by applying network isolation and segmentation to our network.
- We introduced four new security capabilities to help customers secure their networks: Network Security Perimeter (NSP), DNS Security Extensions (DNSSEC), Azure Bastion Premium, and a private subnet feature.

Protect engineering systems:

We have improved the security of systems we use to build, test, and deploy code.

- 99.2% of pipelines have a complete inventory, which is enforced at creation and validated within 24 hours.
- MFA protects 81% of production code branches through proof-of-presence checks.
- Broad adoption of Central Feed Services, which helps to provide developers with a governed open-source feed.

Monitor and detect threats:

To improve our ability to investigate and respond to threats.

- We track 97% of our production infrastructure assets centrally.
- Engineering teams continue to adopt our security logging standard, including the two-year minimum retention policy.
- We added more than 200 additional detections against top tactics, techniques, and procedures (TTPs). Applicable detections will be integrated into Microsoft Defender.

Accelerate response and remediation:

We are addressing more vulnerabilities, more quickly, and continue to improve security-related customer communications.

- 73% success rate addressing cloud vulnerabilities within our reduced time to mitigate, with significantly expanded program scope.
- As part of Zero Day Quest, researchers identified 180 new vulnerabilities in the high impact areas of cloud and AI, enabling us to address them proactively.
- We introduced new processes and playbooks to improve security incident communications to customers.

Culture and Governance

Culture	8
Governance	9

Culture

We've made significant progress fostering a security-first mindset in every employee by activating our culture and people practices.

- Every employee now has a Security Core Priority to discuss with their manager as part of performance check-ins, which we call Connects.
- The Microsoft Security Academy provided tailored learning to more than 50,000 employees and our internal hub for security content has contributed to a 25% reduction in incidents.

Performance and development

The Security Core Priority formalized security accountability in all roles. 100% of performance check-ins include a Security Core Priority. Almost a third (30%) of employees have customized their priority to reflect security work specific to their roles, demonstrating deeper engagement in security ownership.

Training and development

We've expanded the Microsoft Security Academy, making it available to all employees through our internal learning portal.

Our annual security training course, Security Foundations training, has been particularly effective. Early data from completed responses indicates that employees who completed the training were 50% less susceptible to phishing and 20% more likely to report phishing attempts. This training also received a high favorability score of 83 points, indicating strong engagement and relevant content.

The 2024 Trust Code Course (Standards of Business Conduct) provided employees with information about SFI to ensure they put security first in all that they do and reminded them of the importance of speaking up when they see something that does not seem right. The 2024 Trust Code Course was completed by 99% of employees.

Finally, we are launching a Global Cybersecurity Ambassador Program to provide functions across the company with a go-to resource for cybersecurity awareness needs and issues. Ambassadors will serve as the voices of cybersecurity awareness, sharing innovative practices and actions needed to protect Microsoft and our customers and partners.



“The 2024 Security Foundations and Trust Code courses were completed by 99% of employees.”

Key learnings

Performance evaluations can drive meaningful engagement when done right

Embedding security topics into performance evaluations isn't enough on its own, it must be implemented in way that resonates with employees and encourages action. That is why we developed Security Core Priority resources to help managers and employees create meaningful goals and measures of success with a security-first mindset.

Since the launch in August 2024, employees and managers have visited the resources over 200,000 times. The introduction of the Security Core Priority, plus supporting resources, ensured that security became a measurable part of every employee's impact and resulted in strong employee sentiment.

Reinforcing that, when done thoughtfully, security objectives can be embedded into performance evaluation in a way that aligns with employee motivation rather than feeling like a checkbox.

Governance

To improve overall cybersecurity risk and compliance we have expanded and refined our governance oversight and risk management processes.

Cybersecurity Governance Council

The Council is responsible for overall cybersecurity risk and compliance. Since our initial overview of the Council's work in [September 2024](#), we have expanded governance oversight and refined risk management processes.

Key advancements include:

- **Enhanced Deputy Chief Information Security Officer (CISO) coverage:** We added a Deputy CISO for Business Applications and consolidated Microsoft 365 and Experiences and Devices (E+D) under a single Deputy CISO.
- **Enterprise-wide risk tracking:** All Deputy CISOs have completed a comprehensive risk inventory for their platform and function, aligning risks to current threat intelligence and product domains. We centrally track the highest priority risks, combining insights from our cybersecurity threat operations teams (threat-based risks) and Deputy CISOs (product and function risks).

We mitigate prioritized risks at scale through SFI Engineering.

The Microsoft Senior Leadership team continues to review progress bi-weekly. Updates are provided to the Microsoft Board of Directors quarterly.

Key learnings

Embedding security governance early

Integrating Deputy CISOs from key product and functional areas into the Council has advanced security as a core part of development. That makes it more than just a checkpoint, enabling earlier risk mitigation and improved resilience at scale.

Areas covered by Deputy CISO

- Artificial Intelligence
- Azure
- Business Applications
- Commerce
- Consumer
- Core Systems and Mergers and Acquisitions
- Customer Security Management Office
- Experiences and Devices
- Gaming
- Government
- Identity
- Microsoft Corporate
- Microsoft Security
- Regulated Industries

Cybersecurity policy initiatives and governments

Amy Hogan-Burney, Vice President and Deputy General Counsel of Customer Security and Trust, serves as Secretary of the Council. The Council's agenda includes cybersecurity regulatory updates from around the globe, including impacts of emerging regulations on Microsoft and its customers. For example, as a member of the European Commission's Expert Group on Cybersecurity of Products with Digital Elements, formed to implement the 2024 European Union Cyber Resilience Act, Microsoft is actively shaping best practices in vulnerability reporting, security by design, and open-source security, working alongside regulators, industry leaders, and customers to advance security readiness and best practices. Microsoft is also pleased to have been a signatory of the Cybersecurity and Infrastructure Agency, Secure by Design Pledge in 2024. You can learn more about our progress [here](#).

[CISA Secure by Design paper](#) published with agencies of other governments in October 2023



In conjunction with the Microsoft security improvements, SFI includes policy initiatives that encourage governments to conduct cyber activities responsibly, making the Internet a safe place for all. Cyber mercenaries are private sector firms that develop and sell offensive cyber capabilities, primarily to government customers.

For the past two years, Microsoft has been privileged to [participate](#) in the Pall Mall Process, an intergovernmental and multistakeholder initiative led by the UK and French governments to identify responsible practices to limit harms from commercial cyber intrusion capabilities. In April, following several rounds of consultations, the initiative released their "[Code of Practice](#)" for States to reduce proliferation and irresponsible use. At the Munich Security Conference in February, Microsoft contributed to a Foreign Policy Digital Front Lines [report](#) outlining the breadth of the pernicious market for cyber mercenary services. For our own part, Microsoft has also recently adopted an internal policy governing how we evaluate vendors and customers that may be affiliated with cyber mercenary firms.

Security Principles

Secure by Design	11
Secure by Default	12
Secure Operations	13

Security Principles

SFI builds on three core principles to help ensure that our products are secure from inception through deployment and ongoing use. We define them as:

- **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.

The Microsoft Security Development Lifecycle (SDL) helps teams adopt these principles by providing a practical security approach that is risk-driven and agnostic to development methodology or technology.

Read on for examples of how teams across Microsoft have applied these principles to help our customers – and Microsoft – be more secure.

Secure by Design

To promote the safety and security by design of Microsoft generative AI solutions, we formed the Artificial Generative Intelligence Safety & Security organization, or AeGIS for short.

This team provides engineering product teams with the expertise required across the AI Red Team, incident response, design review processes, security training, and developing core software architecture and technical mitigations. You can read more about the AI Red Team's work and lessons learnt, which have been applied to Microsoft AI systems, in this [study](#).

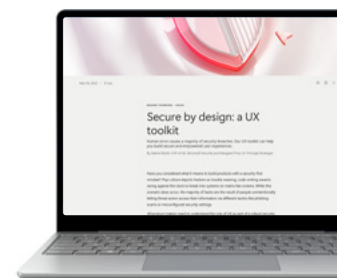
AI Red Teams study lessons



To improve user experience (UX) and security integration in all products, we launched a Secure by Design UX Toolkit and [published a customer-facing version](#). Poor UX can expose customers to phishing, misconfiguring settings, sharing sensitive information unintentionally, and vulnerabilities. By prioritizing UX, we aim to reduce user-targeted attacks, build customer trust through intuitive security features, and minimize human errors with clear indications of secure actions.

The toolkit has been deployed to 22,000 employees, embedding security best practices in product development and ensuring our interfaces are designed to be intuitive, non-intrusive, and help protect customer data.

Secure by Design UX toolkit



Early results are promising:

- **Windows** utilized the toolkit to develop a secure-by-design feature roadmap, including enhanced sign-in security for the Recall feature, and implementing secure defaults across settings, reducing the attack surface.
- **Azure** launched a fraud prevention feature incorporating MFA before logging into the Azure Portal, preventing unauthorized party abuse.
- **Defender Experts** shifted culture to prioritize threat actor prevention, identifying potential access control and impersonation issues, leading to significant product changes.

Building on the principle of Secure by Design, Microsoft is taking a proactive approach to ensuring our products and services are "Fraud-Resistant" by design. In January 2025, we implemented a new fraud prevention policy and crypto-mining detection models to protect customers from fraudsters who take over accounts for crypto-mining activities, which increase customer costs. The policy includes fraud prevention assessments and fraud controls that are integrated into the product lifecycle. For instance, on March 1, 2025, the updated crypto-mining policy, under the Acceptable Use Policy removed exception path language due to significant drawbacks and the absence of legitimate requests.

Security Principles continued

Secure by Default

Teams across Microsoft are delivering new innovations to better protect customers, by default.



Microsoft 365

- **Microsoft 365 CCS:** The Copilot Control System (CCS) is a system of integrated controls and capabilities for Copilot and agents, designed to enable IT administrators and security professionals to effectively secure, manage, and analyze the use of agents across their organization.
- **Microsoft 365 Admin Center:** We are rolling out MFA enforcement for all Microsoft 365 admin center users. Additionally, we introduced a new AI administrator role for efficient administration of Microsoft 365 Copilot and enterprise AI services, without the extensive permissions required for the global admin role.

Windows 11

- **Windows Resiliency Initiative:** WRI is the Microsoft effort to enhance Windows resilience through security and reliability investments such as secure-by-default practices, reducing admin privileges, strengthening app and driver controls, and improving identity protection.
- **Windows Hotpatch:** Announced at Ignite 2024, this provides a 60% reduction in time to adopt security updates, assisted by applying updates seamlessly without system restarts.
- **Windows Smart App Control:** Now AI-enhanced and on by default, this predicts app safety based on 84 trillion daily security signals, blocking malware-connected apps and ensuring verified apps run smoothly.

Microsoft Azure

- **Azure AI Foundry:** Microsoft built this platform to provide customers with the tooling required to design, customize, and manage their own AI applications and agents with built-in content moderation, prompt injection detection and red-team tooling.
- **Azure Integrated HSM:** The Azure Integrated HSM, an in-house security chip inside every new Azure server, strengthens key protection by default using encryption and signing keys within a dedicated Hardware Security Module (HSM), meeting demanding FIPS 140-3 Level 3 security requirements.
- **Azure Portal:** Multifactor authentication (MFA) can block over 99.2% of account compromise attacks, making it one of the most effective security measures available. As of March 31, 2025, MFA is required for 100% of sign-ins to Azure portal, further showcasing our education to the security of all customers.

Microsoft Security

- **Microsoft Entra ID:** To protect customers by default, we are enforcing MFA for Microsoft Entra ID administrator accounts for Microsoft Sovereign Clouds, using Microsoft-managed policies and secure defaults. We have completed enforcement for all tenants in one Sovereign Cloud.
- **Microsoft-managed Conditional Access policies:** In February we started rolling out two new policies that limit device code flow and legacy authentication. These policies help to reduce the risk of phishing and improve overall security posture, by default.
- **Microsoft Defender for Cloud Apps:** We introduced new capabilities to help protect OAuth applications by default across Microsoft 365, Google, and Salesforce, and 18 new detections for changes in activity.

Security Principles continued

Secure Operations

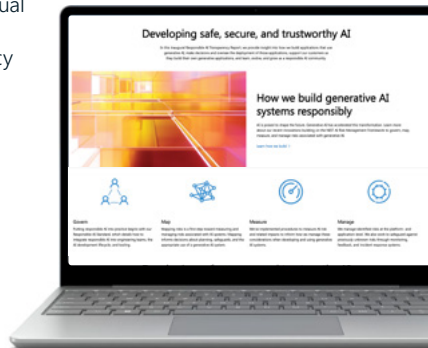
Security controls and monitoring will continuously be improved to meet current and future threats.

Our annual [Responsible AI Transparency Reports](#) outline how we apply secure operational principles across all aspects of our AI system. The next report will provide updates on our responsible AI development practices, customer support initiatives, and proactive efforts to collaborate across sectors.

We also created and scaled the Central Fraud and Abuse Risk team (CFAR) in our effort to holistically protect our digital ecosystem. By improving security policies, developing behavioral-based detection models, and utilizing both automated and manual investigation methods, our detection teams fortified the protection of customer accounts, improved platform security, and thwarted \$4 billion in fraud attempts.

“Our detection teams fortified the protection of customer accounts, improved platform security, and thwarted \$4 billion in fraud attempts”

Microsoft Annual Responsible AI Transparency Report



14B

Analyzed approximately 14 billion sales transactions

49,000

Rejected 49,000 fraudulent partnership enrollments

Engineering Pillars

We have made progress in every engineering pillar and objective. Each objective represents a significant body of work, most will take years to complete. Out of 28 objectives, 5 are nearing completion, 11 have made significant progress, and we continue to make progress against the rest. We sought to prioritize the highest risks, most critical assets, and used platform engineering practices to scale the work and reduce toil.

Protect identities and secrets	15
Protect tenants and isolate production systems	17
Protect networks	19
Protect engineering systems	21
Monitor and detect threats	23
Accelerate response and remediation	25

Engineering Pillars continued



Protect identities and secrets

Over the past 11 months, we have made significant progress in Objectives 2-4.

90% of identity tokens from Microsoft Entra ID for Microsoft apps are validated using one standard identity SDK, which provides a consistent and hardened implementation, improving security. To mitigate advanced MFA attacks, phishing-resistant MFA protects 100% of production system accounts and 92% of employee productivity accounts. Additionally, more than 19 million resources in Microsoft Azure now adhere to our safe secrets standard.

We have also made progress with Objectives 1, and 5-7.

Microsoft Entra ID and Microsoft Account (MSA) access token signing keys are now protected in hardware security modules (HSM) and virtualization-based security in Windows, with automatic rotation. Since September 2024, we have added new defense-in-depth protections, migrated the MSA signing service to Azure confidential VMs, and we are migrating the Entra ID signing service to Azure confidential VMs. These improvements help mitigate the attack vectors that we suspected the actor used in the 2023 Storm-0558 attack on Microsoft. We have improved our ability to detect forged tokens and are expanding asynchronous validation to additional Microsoft workloads. We are partitioning critical infrastructure identity keys to reduce impact if they are compromised. Our work to prepare for the post-quantum cryptographic world is in early development.

Key insight

Validating the effectiveness of new controls is critical, which we do in part through Red Team exercises.

These exercises simulate sophisticated adversary objectives to rigorously test designs, mitigations, and detection mechanisms. By red-teaming their own systems, customers can validate control implementations, identify gaps, and address them rapidly as part of ongoing operations or future improvements.

Objectives

Progress

- 1 Protect cryptographic signing keys**

Protect identity infrastructure signing and platform keys with rapid and automatic rotation with hardware storage and protection.


- 2 Adopt standard SDKs for identity**

Strengthen identity standards and drive their adoption through use of standard SDKs across 100% of applications.


- 3 Phishing-resistant MFA**

Ensure 100% of user accounts are protected with securely managed, phishing-resistant multifactor authentication.


- 4 Safe secrets standard**


Ensure 100% of applications are protected with system-managed credentials.


- 5 Stateful validation for identity tokens**

Ensure 100% of identity tokens are protected with stateful and durable validation.


- 6 Fine-grained key partitioning**

Adopt more fine-grained partitioning of identity signing keys and platform keys.


- 7 Quantum-safe PKI systems**

Ensure identity and public key infrastructure (PKI) systems are ready for a post-quantum cryptography world.



Objective 1: Protect cryptographic signing keys

In September, we reported on our implementation for Entra ID and Microsoft Account (MSA) to use hardware security module (HSM) based storage for access token signing keys, utilizing virtualization-based security (VBS) in Windows. Since then, we have applied new defense-in-depth protections, migrated the Microsoft Account (MSA) signing service to run on Azure confidential VMs which provides additional hardware-based isolation between token signing processes and the underlying hosts, and we are migrating the Entra ID signing service to Azure confidential VMs.

To assess the new defense-in-depth improvements we performed Red Team research and response drills to assess these improvements. This validated that our improved auditing telemetry and reduced key validity periods significantly improved our ability to investigate attacks. This research has also informed our detection strategies and provided insights into how we can defend against more sophisticated attacks.

Objective 2: Adopt standard SDKs for identity

We continue to standardize security token acquisition and validation, and today 90% of tokens issued by Microsoft Entra ID for Microsoft apps are validated using one standardized implementation. Alongside the effort to implement the standard SDKs is the work required to remove and block legacy protocols that no longer meet our stringent security requirements. This approach of adopting standards and retiring legacy protocols will continuously improve our identity systems to defend against emerging threats.

Objective 3: Phishing-resistant MFA

We remain 100% deployed in our production environment and we have made notable progress increasing adoption in our productivity environment (92%). To achieve 100% enforcement in our productivity environment, we are working to improve passkey support in non-Windows platforms and addressing logistical challenges related to the procurement and delivery of physical security keys. We continue to reduce the risk of compromise during new employee setup by enforcing video-based verification, now at 99%.

Objective 4: Safe secrets standard

Since September 2024, we have significantly increased the use of system-managed credentials in Entra ID applications and data storage resources like Azure Storage, Azure SQL Database, and Cosmos DB. We've expanded this standard to include additional Azure resources such as Azure Event Hub, Azure Service Bus, and Azure Cognitive Services. More than 19 million resources are now adhering to our safe secrets standard.

Objective 5: Stateful validation for identity tokens

In September, we noted progress made in our ability to detect forged security tokens using expanded telemetry from the identity SDKs we are adopting as part of Objective 2. We continue to expand asynchronous validation of identity tokens to cover more Microsoft applications and workloads.

Objective 6: Fine-grained key partitioning

As part of our broader investments in signing key protection for Entra ID and Microsoft Accounts (MSA), we are partitioning foundational key systems in the Entra infrastructure to isolate datacenter management layers and Azure regional services, which will prevent a compromised key from allowing traversal to higher-security privilege layers.

Objective 7: Quantum-safe PKI systems

The first phase of our work to ensure identity and public key infrastructure (PKI) systems are ready for a post-quantum cryptography world is to add support for required algorithms into the hardware and operating systems we use. Since our last update, we announced that we have added support for quantum-resistant algorithms in the Windows core cryptographic function library ([SymCrypt](#)).

Engineering Pillars continued



Protect tenants and isolate production systems

Over the past 11 months, we have made significant progress against Objectives 1, 3, and 4.

To mitigate the risk of legacy and unused resources, we have transitioned more than 88% of resources from Azure Service Manager to Azure Resource Manager. Additionally, since September, we have removed an additional 550,000 tenants, bringing the total to 6.3 million. Microsoft Entra ID applications in our production environment have been moved to a fully automated application lifecycle management solution. All new Microsoft Entra ID applications are now registered with our security emergency response system, and we have restricted access to 4.4 million production environment managed identities to protect these critical assets. To reduce the risk of lateral movement, we are piloting a project to move customer support workflows and scenarios into a dedicated tenant.

We have made progress against Objectives 2, 5, and 6.

Security baselines are enforced across all types of Microsoft tenants, and a new tenant provisioning system automatically registers new tenants in our security emergency response system. We have prioritized stringent enforcement of least-privilege for critical assets and high-impact scenarios, migrating ~1,000 high-privilege applications to new standard authentication protocols. Access to production resources is restricted to production-ready, locked down, physical devices. As noted in November 2024, we have moved 28,000 high-risk users, working on sensitive workflows, to a locked-down Azure Virtual Desktop (AVD) infrastructure. We continue to refine our user experience for these hardened endpoints.

Key insight

If an attacker compromises an asset, they will often attempt to move laterally to other assets. Modeling this as a graph can be highly beneficial; for example, a Microsoft Entra ID application, with a role-based access control (RBAC) permissions on an Azure subscription, can be visualized as two connected nodes (one for the application and one for the subscription). Modeling assets as a graph will reveal unknown vulnerabilities and classes of known issues that need to be mitigated to reduce lateral movement vectors.

Objectives

Progress

- Remove legacy systems that risk security**

1 Maintain the security posture and commercial relationships of tenants by removing all unused, aged, or legacy systems.


- Secure all tenants and their resources**

2 Protect 100% of Microsoft, acquired, and employee-created tenants, commerce accounts, and tenant resources to the security best practice baselines.


- Higher security for Entra ID apps**

3 Manage 100% of Microsoft Entra ID applications to a high, consistent security bar.


- Eliminate identity lateral movement**

4 Eliminate 100% of identity lateral movement pivots between tenants, environments, and clouds.


- Continuous least-privilege enforcement**

5 100% of applications and users have continuous least-privilege access enforcement.


- Secure devices used for access**

6 Ensure only secure, managed healthy devices will be granted access to Microsoft tenants.



Objective 1: Remove legacy systems that risk security

Legacy systems can pose significant security risks, as they are often targets for threat actors. Since September we have eliminated 550,000 unused and aged tenants, which continues to reduce the potential attack surface. We continue to make significant progress in our transition away from Azure Service Manager, successfully removing or migrating more than 88% of resources previously managed by Azure Service Manager to Azure Resource Manager. Customers that are currently using this service should also consider migrating resources to [Azure Resource Manager](#).

Objective 2: Secure all tenants and their resources

To further improve security, we have expanded security baselines to all types of Microsoft tenants, including production, dev/test, and lower-risk tenants like employee-created demo/learning tenants. For lower-risk tenants, we implemented strict lifecycle policies that delete tenants when they expire.

The new tenant provisioning process now automatically registers all new tenants into our security emergency response system. This automatic registration eliminates human error during provisioning and improves our ability to detect and address potential security incidents. Based on learnings from tracking down employee created tenants, we developed a new tenant classification machine learning model that connects signals from various data sources to improve the precision of our internal tenant inventory. After training and validating the model across pilot scopes, the model was successfully deployed and adopted in our commercial cloud.

Objective 3: High security for Microsoft Entra ID apps

In September, we reported a complete iteration of lifecycle management for Microsoft Entra ID apps. After extensive testing we have switched our production environment to fully automatic application lifetime management to ensure continuous security posture management, reducing security vulnerabilities and operational inefficiencies. More than 2,000 multi-tenant applications have been converted to single-tenant to reduce lateral movement. All new Microsoft Entra ID applications are now integrated into our security emergency response system. Existing applications unable to integrate are being disabled. We have also added network restrictions to limit authentication to 4.4 million production environment managed identities to mitigate future attacks.

Objective 4: Eliminate identity lateral movement

To further isolate our systems and reduce the risk of lateral movement we created a tenant layering standard. This standard allowed us to categorize tenants into layers and define the valid direction for service principal creation. We are also piloting a project to separate our customer support workflows and scenarios into a dedicated tenant, isolated from other environments.

Objective 5: Continuous least-privilege enforcement

Least-privilege is a continuous effort to ensure users and applications have only the necessary access rights. We have prioritized stringent enforcement of new standards for critical assets and high-impact scenarios. Given most services depend on other services to deliver value, it crucial to ensure least-privilege adherence in two common scenarios – services acting on behalf of a user and services that are not acting on behalf of a user. To address these scenarios, we developed new standard authentication protocols that significantly limit token privileges for applications requiring service-to-service protocols, and have migrated ~1,000 high-privilege applications to these new protocols.

Objective 6: Secure devices used for access

To better protect our production environment, we have prioritized locking down the devices used to access these resources. As reported in November, we have deployed 98,000 production-ready, locked-down devices. Any user accessing production systems must use these locked-down devices. To further improve security for productivity environments, we are moving high-risk users, working on sensitive workflows, to locked down Azure Virtual Desktop (AVD) endpoints. After onboarding 28,000 users, as reported in the November update, we have been focused on ensuring optimal user experience for those users. This includes deploying to 6 new Azure regions to reduce network latency and fine-tuning virtual machine parameters to improve responsiveness.

Engineering Pillars continued



Protect networks

We are nearing completion of Objective 1.

Currently, more than 99% of network devices are logged in a central repository for full lifecycle management. These devices use centralized authentication and audit trails, are configured with ACLs for IPv4/v6 to restrict lateral movement, and have safeguards in place to prevent key compromise or abuse.

We have made initial progress against Objective 2, which will create additional layers of defense in depth against attackers by applying network isolation and micro-segmentation to the entire Microsoft network.

We are establishing a plan to reimplement ACLs to a durable standard and enforce least-privilege ACLs on every endpoint. Currently, 20% of first-party IPs are tagged and 93% of first-party services have established plans for allocating IPs from tagged ranges and provisioning IP capacity. Finally, we have introduced four new capabilities to help customers isolate and secure their network resources: Network Security Perimeter (NSP), DNS Security Extensions (DNSSEC), Azure Bastion Premium, and a private subnet feature.


Key insight

By focusing on comprehensive inventory, proper telemetry coverage, and implementing secure configurations, customers can significantly reduce security vulnerabilities. Furthermore, network isolation and micro-segmentation provide critical layers of defense, preventing potential threats from spreading to critical assets.

Objectives

- 1 Inventory and security standards**
Secure 100% of Microsoft production networks and systems connected to the networks by improving isolation, monitoring, inventory, and secure operations.
- 2 Network isolation**
Apply network isolation and microsegmentation to 100% of the Microsoft production environments, creating additional layers of defense against attackers.
- 3 Secure customer cloud networks**
Enable customers to easily secure their networks and network isolate resources in the cloud.

Progress



4 new capabilities

Objective 1: Inventory and security standards

Multiple investments have been made to harden the network. These include, but are not limited to, strong and centralized authentication, access controls (ACLs), virtual private networks (VPN), secure access workstations (SAW), removal of dual homing, security monitoring across the network fleet, ensuring network device keys are unique per device and continuously rotated to reduce blast radius, and configuring service accounts use one-time passwords (OTP) generated to prevent key/secret abuse.

Currently, 99.97% of network devices are logged in a central repository with metadata for lifecycle management. Additionally, 99% of devices use centralized authentication and audit trails, 99.8% are configured with ACLs for IPv4 and IPv6 to restrict lateral movement and provide traffic visibility.

Objective 2: Network isolation

Network isolation provides an additional layer of defense against intrusion threats and lateral movement to higher-value targets. Our plan is to systematically re-implement ACLs to a durable standard, centrally enforcing least-privilege access control lists (ACLs) on every endpoint in the fleet. Given the scope of this effort, progress is expected to continue for several quarters. This effort depends on properly identifying all resources on the network. Currently, 20% of existing first party IPs are tagged and 93% of first-party services have established plans for allocating IPs from tagged ranges and provisioned IP capacity.

We are adopting Network Security Perimeter (NSP) for Microsoft services, starting with high-value targets. NSP enables us to standardize network isolation mechanism across Microsoft services, so it is easier to manage. More than 21 million resources are targeted for NSP, currently 354,000 are in NSP learning mode (with plans to adopt enforcement mode soon), 7,000 are in enforce mode, and 317,000 have public access disabled. We are using Service Level Network Managers (SLNM) to centrally enforce least-privilege access and have reached 50% adoption, with more than 7,000 services deploying centralized ACLs.

Objective 3: Secure customer cloud networks

To empower customers to deploy and maintain secure and properly configured networks, we announced, at Microsoft Ignite 2024, the public preview of the Network Security Perimeter (NSP), enabling customers to achieve PaaS isolation with unified firewall, granular access control, centralized logging, and egress controls to prevent data exfiltration. We introduced DNS Security Extensions (DNSSEC) that enhances DNS security with the validation of responses, ensuring origin authority and data integrity. We launched Azure Bastion Premium, featuring Private Only Mode and Graphical Session Recording for heightened security and compliance, and the Developer version for cost-effective, secure remote access (RDP and SSH sessions) in development environments. We introduced a private subnet feature to remove default outbound [access](#) so that customers can implement private-only subnets, eliminating a common impediment to secure-by-default configuration.

Engineering Pillars continued



Protect engineering systems

We are nearing completion of Objectives 1-3.

99.2% of Azure DevOps release pipelines now have a complete inventory. 79% of our repositories and pipelines are fully inventoried at creation as part of our StartRight* process, with the remaining either inventoried or disabled within 24 hours using our StayRight** process. To better protect source code, we have significantly reduced the number of members in admin roles in our engineering system and enforced multifactor authentication through proof-of-presence for production code branches, now at 81%. We have reduced access to code signing services and our live secret detection and remediation capabilities have improved to 99.5%.

We have made significant progress against Objectives 4-5.

92% of commercial cloud production build pipelines are now centrally governed, making builds more consistent, efficient, and trustworthy. Additionally, we have expanded efforts to address high-severity open-source software vulnerabilities and achieved broad adoption of Central Feed Services (CFS). This lowers risk of ingesting compromised package versions by scanning for viruses, checking for typo-squatting, blocking malware, and quarantining for new package versions.

* **StartRight**: a solution designed to streamline and support compliance to SFI requirements in the resource creation process within Azure DevOps.

** **StayRight**: manages drift for inventory by enforcing SFI requirements through constant monitoring for events like new assets creation, ownership changes, and classification errors.

Key insight

Enhancing ownership at a granular level has been helpful in addressing code accountability and timely responses to security bugs and other issues. Additionally, tools and processes like StartRight* and StayRight** are critical to ensure that new assets are created in conformance with security standards and any drift is corrected automatically.

Objectives

Progress

- 1 Complete software asset inventory**

Build and maintain inventory for 100% of the software assets used to deploy and operate Microsoft products and services.
- 2 Zero trust for source code access**

100% of access to source code and engineering systems infrastructure is secured through Zero Trust and least-privilege access policies.
- 3 Secure code deployment**

100% of source code that deploys to Microsoft production environments is protected through security best practices.
- 4 Standardize secure development pipelines**

Secure development, build, test, and release environments with 100% standardized, governed pipelines and infrastructure isolation.
- 5 Protect the software supply chain**

Secure the software supply chain to protect Microsoft production environments.

Objective 1: Complete software asset inventory

The progress made in the software asset inventory objective has significantly reduced risks associated with untracked and unmanaged software assets. By completing inventory for build and release pipelines, and enhancing visibility and control through granular ownership for repositories, we have greatly improved the management of software assets. Contributing inventory data into the Azure Resource Graph (ARG) provides visibility into relationships between software assets, their deployment status, and ownership information, enabling remediation efforts.

As a result, 99.2% of Azure DevOps release pipelines now have complete inventory and 79% of repositories and pipelines have been fully inventoried at creation through our StartRight process. The remaining are created through programmatic APIs that our system cannot catch – instead our StayRight process catches these within 24 hours to ensure we get the inventory and visibility.

Objective 2: Zero trust for source code access

To better protect source code and engineering system access, we are implementing several measures: 1) reducing unused and overly permissive access to admin roles, 2) enforcing safe default configurations for all Microsoft Azure DevOps organizations, 3) migrating all Azure DevOps OAuth Apps to Microsoft Entra Apps, and 4) enforcing proof-of-presence (PoP) for code changes. We have significantly decreased the number of members in admin roles, thereby lowering the risk of code compromise, and reduced the number of users who can publish to build feeds, which helps prevent the introduction of malicious artifacts into the source code. We have also disabled default automatic builds for forked repositories in ADO projects. Proof-of-presence enforcement for production branches is now at 81%.

Objective 3: Secure code deployment

To better protect source code that deploys to Microsoft production environments, we are reducing access to code signing services and continuing our work to detect and remediate any secrets in code. We have significantly reduced risks by locking down 56% of code signing access to identities in the production tenant. Live secret detection and remediation capabilities are now operating at 99.5%, further improving the security and reliability of both development and production systems.

Objective 4: Standardize secure development pipelines

We are moving commercial cloud production pipelines to new centrally governed pipeline templates. These templates help reduce risk by centralizing the enforcement of certain Security Development Lifecycle (SDL) requirements. 92% of commercial cloud production build pipelines are now centrally managed, significantly improving security at this key development stage. To further improve security we have expanded this to include production release pipelines, 54% of which are now centrally managed.

Objective 5: Protect the software supply chain

We are driving four standards to help ensure open-source software (OSS) used in Microsoft production environments is sourced from governed internal feeds and free of known critical and high-severity public vulnerabilities. Component Governance (CG), a Software Composition Analysis (SCA) tool that tracks OSS usage and vulnerabilities in OSS, has achieved broad adoption and is enabled by default. The Centralized Feed Service (CFS), which provides governed feeds for consuming open-source software, has reached broad adoption across the production repositories for the supported package types.

Engineering Pillars continued



Monitor and detect threats

Objective 1 is nearing completion.

We have inventoried and now centrally track more than 97% of our production infrastructure assets. In addition, 99% of network devices, and more than 95% of nodes/machines, have central security log collection with our two-year retention policy enforced.

We have also made significant progress against Objectives 2-3.

To improve our ability to investigate and respond to incidents, we developed a standard schema for service level audit logs, adopted by most Microsoft service teams. Five out of seven major security log categories now follow the two-year minimum retention standard. To streamline access to telemetry in centralized security logs and improve response time to incidents, we embedded data analysts with the security investigation team.

We are continuously improving our ability to automatically detect and respond to threats. Since September, we have added more than 200 detections against top tactics, techniques, and procedures (TTPs) across the Microsoft infrastructure, enhancing our ability to detect threat actor activities. Applicable detections will be integrated into Microsoft Defender.

Key insight

Security threat detection validation is crucial for ensuring that security detections are effective and comprehensive. It involves verifying that detections cover all relevant platforms and environments, which helps identify any gaps in the current detection process. Accurate validation, using realistic attack signal and emulation, prevents a false sense of security and helps ensure that detections are deployed correctly and are actionable across different environments.

Objectives

Progress

- | | |
|--|-------------------------------|
| <p>1 Complete production infrastructure inventory</p> <p>Maintain a current inventory across 100% of Microsoft production infrastructure and services.</p> | |
| <p>2 Security log retention standards</p> <p>Retain 100% of security logs for at least two years and make six months of appropriate logs available to customers.</p> | |
| <p>3 Centralize access to security logs</p> <p>100% of security logs are accessible from a central data lake to enable efficient and effective security investigation and threat hunting.</p> | |
| <p>4 Rapid anomaly detection and response</p> <p>Automatically detect and respond rapidly to anomalous access, behaviors, and configurations across 100% of Microsoft production infrastructure and services.</p> | <p>200+ detections</p> |

Objective 1: Complete production infrastructure inventory

Effective security monitoring and threat detection require a centralized accurate asset inventory. We are centrally tracking more than 97% of production infrastructure inventory (e.g., network devices, physical server machines, and virtual machine nodes). Additionally, 99% of network devices and more than 95% of nodes/machines now have central security log collection with two-year retention.

Objective 2: Security log retention standards

Over the past six months, significant advancements have been made in enhancing threat investigation capabilities. Specifically, we have established a standard for service-layer security audit logs, now adopted by most Microsoft services, which includes a minimum two-year retention period. This retention period has been enforced centrally for most security logs, including five out of the seven major security log categories for Microsoft production infrastructure and services. We will expand enforcement of this standard for all logs emitted and retained locally.

Objective 3: Centralize access to security logs

The security logs that are being retained with a central two-year retention policy can be accessed through central logging systems. To improve the efficiency of security investigation efforts, we have embedded data scientists in the security investigation team to swiftly identify and resolve any telemetry gaps. Additionally, specialized accounts for the investigation team have been created, enabling secure access to telemetry and enhancing investigation effectiveness. This streamlined process has enabled service teams to more rapidly respond to security incidents.

Objective 4: Rapid anomaly detection and response

To ensure effective detection of threat actor activities, we have collaborated with various security teams to identify and deploy more than 200+ additional detections against top tactics, techniques, and procedures (TTPs) across the Microsoft infrastructure. We work closely with service architects to tailor detections to unique system architectures, scaling our capabilities with automated models to distinguish normal behavior from potential threats. We then validate these detections using tools and Red Team operations to identify and address gaps. These custom detections and Microsoft Defender for Cloud are being deployed across the Microsoft environment. Applicable detections will be integrated into Microsoft Defender.

Engineering Pillars continued



Accelerate response and remediation

Over the past 12 months, we have made progress against all objectives.

In November, we reported a 90% success rate in addressing high-severity cloud vulnerabilities within our target time. We expanded the program’s scope, achieving a 73% success rate while mitigating more vulnerabilities. To accelerate coordinated vulnerability mitigation, we work to uncover exploits used in the wild and then work with the appropriate vendors to ensure they are properly mitigated. We also announced Zero Day Quest, the largest event focused on research into high-impact areas like cloud and AI. Researchers discovered 180 new critical and important vulnerabilities, which we are working to proactively address. We continue to publish no-action critical cloud vulnerabilities and exposures to provide transparency for our customers. Finally, we established new processes and playbooks to improve security incident communications for customers.


Key insight

Automating vulnerability patching is crucial for efficiency, especially at cloud-scale. We have deployed automated operating system upgrades to 86% of our first party Virtual Machine Scale Sets (VMSS-based services), resulting in more than 91 million upgrades in 2024. Customers should strongly consider automated vulnerability patching to remediate issues across their enterprise.

Objectives

- 1 Accelerate vulnerability mitigation**
Reduce the Time to Mitigate for high-severity cloud security vulnerabilities with accelerated response.
- 2 Transparency of cloud vulnerabilities**
Increase transparency of mitigated cloud vulnerabilities through the adoption and release of Common Weakness Enumeration™ (CWE™), and Common Platform Enumeration™ (CPE™) industry standards for released high severity Common Vulnerabilities and Exposures (CVE) affecting the cloud.
- 3 Enhance public messaging and engagement**
Improve the accuracy, effectiveness, transparency, and velocity of public messaging and customer engagement.

Progress



34 no action CVEs

Ongoing effort

Objective 1: Accelerate vulnerability mitigation

In November, we reported a 90% success rate in addressing high-severity cloud vulnerabilities within our target timeframe. Since then, we have expanded the scope of this program to include more environments, products, and lower severity vulnerabilities. Considering this expanded scope our success rate is now 73%.

To improve responsible vulnerability mitigation, in the high impact areas of cloud and AI, we launched the [Microsoft Zero Day Quest](#) in November 2024. To date, nearly 100 researchers have discovered 180 new critical and important vulnerabilities, which we are working to address proactively. Learnings from the Zero Day Quest will be shared across Microsoft to improve cloud and AI security.

Microsoft teams work with the industry to discover and fix zero day exploits in the wild. For instance, we identified a North Korean threat actor using an unknown vulnerability in Google's v8 engine and collaborated with Google to fix it in Chrome and Edge within one week. Learn more [here](#).

Objective 2: Transparency of cloud vulnerabilities

We continue to use industry standards like CVE, CWE, and CPE annotations to provide customers with additional information on cloud vulnerabilities and increased transparency.

In 2024, we published 1025 CVEs with CWE and CPE annotations. We also started publishing no-action critical cloud CVEs. We are pleased to see other cloud providers expanding their CVE programs to do the same. Of the 34 no-action-required cloud CVEs published, the majority were from Microsoft employees, showcasing our ongoing commitment to transparency.


















Objective 3: Enhance public messaging and engagement

In 2024, Microsoft invested in additional resources to support customers in major security incidents, including the creation of a customer-facing executive engagement team within the Office of the CISO (OCISO) and tripling the number of global security advisors. These security advisors proactively inform customers about security risks and incidents, provide one-to-one briefings, and oversee security-related escalations.

OCISO manages cross-product/service, customer impacting security incidents and has invested in its Customer Security Management Office (CSMO) to drive urgency, rigor, and consistency in these incident communications to customers. In September 2024, CSMO began implementing new processes and playbooks to help expedite and improve our communication with customers about security incidents. They work with security, engineering, and product teams to understand the incident, assess customer impact, and partner with other teams to provide actionable information to customers. Since the start of 2025, all OCISO-managed incidents have adhered to these new processes.

Appendix: CSRB Mapping

This table maps progress in this report to the Cyber Safety Review Board (CSRB) recommendations made in [March 2024](#) that apply to Microsoft and/or all cloud service providers (CSPs). Like SFI pillars and standards, our work to address these recommendations and maintain company practices will be ongoing, given the breadth or complexity of certain recommendations.

Recommendation category and CSRB number	Status
Microsoft corporate security culture	1 Complete 
	2 Complete 
	3 Complete 
CSP cybersecurity practices	6 Significant progress 
	9 In progress 
Audit logging norms	4 Complete 
	5 Significant progress 
	10 Significant progress 
Digital identity standards and guidance	11 In progress 
	13 In progress 
CSP transparency	14 In progress 
	15 In progress 
	16 In progress 
	17 In progress 
Victim notification process	18 Open to collaborate 
	19 In progress 
	20 Open to collaborate 

Note: CSRB recommendations 7, 8, 12, and 21-25 are not applicable to Microsoft or CSPs.



Secure Future Initiative (SFI)

April 2025 progress report

©2025 Microsoft Corporation. All rights reserved.