

Secure Future Initiative (SFI)

Microsoft Ignite, November 2024 update

Secure Future Initiative

Secure by design

Secure by default

Secure operations

Security culture and governance



Protect identities and secrets



Protect tenants and isolate production systems



Protect networks



Protect engineering systems



Monitor and detect threats



Accelerate response and remediation

Continuous improvement



Paved path



Standards



Executive summary

In May 2024, Microsoft CEO Satya Nadella made security the company's top priority. Since that time, we have dedicated the equivalent of 34,000 engineers to advance the objectives laid out in SFI, making it the largest cybersecurity engineering project in history.

In September 2024, we published our first progress report, including progress made in every area and each engineering pillar. Customers asked for insight into our implementation and guidance.

In this update—released at Microsoft Ignite—we share insight into our efforts to foster a culture that prioritizes security above all else, establish governance and engineering frameworks to help manage cybersecurity risk at scale, and deter escalating nation-state threat actor activity.

Each engineering pillar outlines the work we are doing to reduce risk in specific areas and steps we have taken to improve security by default, by design, and in our operations. For example, we enforce multifactor authentication (MFA) by default for all new tenants, are making progress enforcing Microsoft-managed Conditional Access policies, and are in the process of enforcing MFA for the Microsoft Azure Portal and Microsoft 365, Microsoft Entra, and Intune admin centers. We also block secrets from being exposed at push for customers using GitHub Advanced Security—which is free for public GitHub repositories. Additionally, at Ignite, we announced [Zero Day Quest](#), our inaugural hacking and research challenge focused on advancing AI and cloud security through collaboration with the community.

Each engineering pillar also includes detailed customer guidance. Key themes include reducing risk related to credentials, improving productivity and production device security, minimizing excessive permissions, maintaining comprehensive asset inventories, adhering to security baselines, and implementing centralized and standard audit logging. In addition to the guidance provided, customers can learn more about Zero Trust implementation [here](#).

Table of contents

- Executive summary2
- Table of contents3
- Culture4
- Governance6
- Cybersecurity and governments8
- Engineering pillars10
- 1. Protect identities and secrets11
- 2. Protect tenants and isolate production systems15
- 3. Protect networks17
- 4. Protect engineering systems20
- 5. Monitor and detect threats24
- 6. Accelerate response and remediation27

Culture

[Growth mindset and continuous learning are foundational to our culture.](#) As such, we are focused on embedding a security-first mindset into our people practices, beginning with the following:

- **Performance and development:** Security is now a factor in performance evaluation at every level. All employees are required to develop their Security Core Priority and discuss initial impact with their manager by December 31, 2024.
- **Learning and development:** We have introduced additional security training resources, such as the Microsoft Security Academy and 100% of employees are required to take our annual Security Foundations and Standards of Business Conduct trainings.

In this update we share insights into our implementation and best practices for human resources and corporate functions.

Security Core Priority

The Security Core Priority is how we hold employees accountable for building more secure products and services, ensuring customer safety, and protecting data at Microsoft. It reinforces that every employee, regardless of function, has a responsibility to directly contribute to the security of Microsoft and our customers. The Security Core Priority is divided into two parts:

- **Core and Common:** This includes completing required security training, prioritizing security in daily work, implementing more secure approaches in alignment with SFI, and taking advantage of MSProtect—an internal hub for all security-related content—to learn more.
- **Optional:** Each employee customizes this with security work specific to their role.

Employees discuss the impact delivered against their Security Core Priority during performance check-ins called Connects. Managers consider this input when determining impact and recommending annual rewards. This approach ensures security is integrated into the performance evaluation of all employees, regardless of their role.

Our work to activate the Security Core Priority is ongoing and reinforced through continuous training, resources, rewards processes, and communications. We remain on track for every employee to have drafted their Security Core Priority and discussed impact with their manager by December 31, 2024.*

**Implementation will vary in certain countries based on local regulations.*

Microsoft Security Academy

Continuous learning is at the core of our growth mindset culture. As such, our learning strategy is to transform skilling into a core component of our employee experience and emphasize the development of differentiated skills that are critical to business success.

To address security skills specifically, we created the Microsoft Security Academy. Content from this academy will be used by academies for each division and customized to address function-specific scenarios.

This approach allows us to provide every employee with core content that can help them quickly acquire the security skills they need and customized content for their role and function.

Security Foundations

Security Foundations is required annual training for all employees. The objectives for this training are to foster cybersecurity awareness and protective behaviors among all employees and to help reduce human error and prevent cyberattacks. This is crucial because [74% of all security breaches involve human error](#), highlighting the need for comprehensive security training.

Training modules focus on phishing prevention, access management, and AI, and they always include an executive call to action. In October, Microsoft launched a new Security Foundations module, which focused on increasing awareness of the Microsoft SFI. We are on track for 100% of employees to complete this training by November 30, 2024.

Learnings and best practices:

- Use existing systems (for example, performance management) to embed new attributes into your culture more quickly.
- Identify the most strategic skilling investments to enable the highest priority security-related outcomes.
- Basic security hygiene best practices can help prevent many cyberattacks. Reinforce these best practices through engaging, short, mandatory training.

Governance

Microsoft has invested in several governance and engineering structures that help us identify, prioritize, and address cybersecurity risk across the company. These include:

- **Cybersecurity Governance Council:** Responsible for overall cyber risk and compliance. Reviews emerging threats, prioritizes risks, and ensures compliance with cybersecurity regulations.
- **Execution framework:** Addresses risk at scale through engineering actions with repeatable and durable patterns.
- **Senior leadership engagement:** Ensures continuous oversight and accountability.

Below are additional insights into these structures and best practices from our learnings to date.

Reviewing and prioritizing cybersecurity risk

Microsoft Chief Information Security Officer (CISO) Igor Tsyganskiy and 14 Deputy CISOs form the recently established Cybersecurity Governance Council.

Each Deputy CISO represents and is accountable for a security domain—an engineering division into which they report or a foundational security function reporting to the CISO. Deputy CISOs reporting into an engineering division report into division leaders, ensuring local responsibility for and awareness of risk.

Deputy CISOs are accountable to identify, prioritize, and track the status of cybersecurity risks for their domain. They meet twice a week to review status and prioritize new risks. This helps ensure risks benefit from the collective review of 14 Deputy CISOs with decades of experience across industries. Prioritized risks are reviewed with SFI engineering leaders for action and monitored to completion.

Cybersecurity risk and government cybersecurity policy initiatives

Governments worldwide increasingly focus on the security of information technology, using cybersecurity policy initiatives to encourage better protection against cybersecurity threats by manufacturers and customers. Looking to the future, cybersecurity regulatory requirements are forecast to increase.

To facilitate compliance with cybersecurity policy initiatives, the Cybersecurity Governance Council—with support from cybersecurity policy and regulatory subject matter experts—reviews and evaluates requirements. Notable examples of how we have worked to address these requirements include the commitment Microsoft has made to the US Cybersecurity and

Infrastructure Security Agency's voluntary [Secure by Design Pledge](#) and the European Union's [NIS2 Directive](#). As Microsoft fulfills its responsibilities with these initiatives, the emerging best practices identified can be translated into repeatable SFI standards and implemented.

Actioning and driving scale through engineering

Leaders responsible for each SFI pillar mitigate prioritized risks through engineering actions. They meet three times a week to align on strategy and drive execution across the company. This cross-organization SFI engineering effort is led by Microsoft President of Identity and Network Access Joy Chik, Executive Sponsor for SFI.

As we address each risk, we create standards and paved paths. These platform engineering concepts enable us to scale work using repeatable and durable patterns—strengthening ongoing developer productivity and accelerating progress. We prioritize the sequencing of standards based on risk assessment.

The goal for every standard is 100% adoption of, or compliance with, that standard. Once a standard has been established to support enduring progress, we track teams as they move toward completing that standard and then work to stay in compliance. Standards help promote security by design and by default.

Reporting and accountability

The Cybersecurity Governance Council and SFI Engineering leadership provide regular updates on SFI progress to the Microsoft Senior Leadership Team. As we have made progress on SFI goals, the cadence has shifted from weekly to biweekly.

The Microsoft CISO, Vice President for Customer Security and Trust, and the SFI Executive Sponsor report on cybersecurity risk, regulatory compliance, and SFI progress to the Microsoft Board of Directors quarterly.

Learnings and best practices:

1. Increasingly, technical organizations with diverse cybersecurity risk domains:
 - a. May need a distributed information security officer structure to provide domain knowledge and enable rapid decision making.
 - b. Should encourage close coordination with, and support from, regulatory subject matter experts.
2. Security should be baked into existing engineering frameworks, from product design through development and operations, that embrace continuous improvement.
3. Cybersecurity updates to senior leadership must be frequent and go beyond status reporting, with in-depth attention to complex topics which span multiple domains.

Cybersecurity and governments

As [noted at the launch](#) of our SFI work in November 2023, the evolving threat landscape and the critical role of information technology in society require rapid alignment between public and private sectors on international cyber norms. These include:

- States should recognize cloud services as critical infrastructure, if they haven't done so already, and off limits to targeting in cyber operations.
- States should not conduct cyber operations that threaten critical infrastructure, nor should they allow others to do so via cyberattacks originating within their territory.
- States should not indiscriminately compromise the security of cloud services for the purposes of espionage.
- States should construct cyber operations to avoid imposing costs on those who are not the target of operations.

Microsoft has worked in recent months to raise the visibility of escalating nation-state activity that may transgress these expectations and to support efforts to strengthen and uphold cyber norms going forward.

Advancing recognition of international cyber norms

Launched on October 15, the [2024 Microsoft Digital Defense Report](#) (MDDR) provides a comprehensive overview of the threat landscape visible to Microsoft security teams in the past year. The report specifically highlights the breadth of malicious activity conducted by nation-state threat actors tracked by our threat intelligence teams, breaking down this activity based on the sectors these actors target. The 2024 report suggests that over 50% of nation-state cyber operations observed targeted critical infrastructure sectors like IT, government services, transportation, manufacturing, and communications—sectors that should not be threatened under the [UN norms for responsible state behavior online](#).

Beyond the MDDR, Microsoft also [reported publicly](#) in blog posts in October on the activity of a threat actor, Star Blizzard, associated with the Russian government, and the subsequent actions taken against that threat actor. After tracking a persistent spear-phishing campaign by Star Blizzard beginning in 2023 and targeting civil society organizations attempting to exfiltrate sensitive data, the Microsoft Digital Crimes Unit successfully secured a court order to disrupt the threat actor's technical infrastructure. The resulting public reporting from Microsoft also highlighted how this action, by a nation-state threat actor, appears to violate the UN norms for responsible state behavior online.

Action required: Deter escalation of attacks

In the 2024 MDDR, Microsoft outlines for the first time the necessary steps that governments can and should take—together with other stakeholders—to deter escalating nation-state attacks that target critical infrastructure. These include:

- Clarifying existing norms and creating more inclusive diplomatic processes for establishing rules for responsible state behavior online.
- Strengthening public attributions of nation-state cyberattacks to make them more frequent, robust, and indicative of what international norms/laws were violated by a nation-state cyberattack.
- Encouraging governments to signal and impose more significant consequences for cyber operations that violate norms and put civilians at risk in order to discourage adversaries.

Looking forward: RAISE, the Roundtable for AI, Security and Ethics

As important as it is to clarify and uphold existing international cyber norms, it is also important to consider what new expectations are necessary for emerging technologies like AI. This is why, in recent months, Microsoft partnered with the UN Institute for Disarmament Research to launch the [Roundtable for AI, Security, and Ethics](#) (RAISE), an initiative dedicated to AI for national security, grounded in international legal and normative frameworks.

RAISE has brought together experts from industry, academia, civil society, and government—including representatives from China, Ecuador, India, Israel, Japan, Namibia, Russia, Switzerland, the UK, and the US—who have led initial work to identify shared interests, enhance cooperation, and generate actionable recommendations on this important topic. The goals of RAISE as an ongoing initiative are to reduce the risks of AI in national security, support multilateral AI governance, and promote AI to enhance security globally.

Security above all else



Protect identities
and secrets



Protect engineering
systems



Protect tenants and isolate
production systems



Monitor and
detect threats



Protect networks



Accelerate response
and remediation

Engineering pillars

1. Protect identities and secrets

In September, we shared insights into the progress made in Pillar 1. In this update, we focus on work done to reduce credential-related risk and include customer guidance.

- Threat actors frequently use tactics that exploit **compromised user accounts or stolen credentials and tokens** for applications and storage systems.
- To help secure customers, multifactor authentication (MFA) [is now on by default for new tenants](#) and will be enforced for the Microsoft Azure Portal, Microsoft Entra admin center, Intune admin center, and Microsoft 365 admin center.
- To mitigate against more sophisticated attacks, we are [enforcing phishing-resistant MFA](#) across our productivity environments. We are also **adopting managed identities** to protect against and limit the impact of compromised credentials.
- Customers can mitigate related risks by adopting **phishing-resistant MFA** and modifying their applications to use **Azure Managed Identity**.

Phishing-resistant authentication methods

Problem description

Attackers are increasingly trying to evolve their tactics to bypass MFA protection by intercepting security codes, using stolen phone numbers, overwhelming users with MFA push notifications, and using fake replicas of legitimate websites to capture first- and second-factor credentials. Microsoft Entra ID helps enable enforcement of phishing-resistant authentication to address these emerging threats. As we implemented this in our productivity tenants, we faced two main challenges:

- **To provide phishing-resistant authentication** methods across a highly mixed environment of user personas (administrators, developers, marketing teams, and so on) and device/OS platforms including Windows, macOS, iOS, Android, Linux, and Virtual Desktops (VDI).
- **To secure account setup and recovery processes** to protect against attacks in these initial account stages and support users in self-onboarding and account recovery wherever they work in a geo-distributed environment.

Microsoft implementation

First, we segmented our users and devices based on risk category, team, and geography. Analyzing this data led us to develop new Microsoft Entra product capabilities, including passkeys in Microsoft Authenticator and improved Microsoft Entra sign-in experiences. Closed-loop feedback from these internal tests and deployments helped product teams improve feature releases and

1. Protect identities and secrets

customer guidance with the broader goal of simplifying phishing-resistant authentication deployments for Microsoft customers.

Second, we implemented a phased, multichannel rollout (via email, Microsoft Teams, and Authenticator app notifications) using a communication campaign that lasted several weeks. This campaign informed users of the actions to register phishing-resistant authentication methods and provided for adequate help desk infrastructure support. At the same time, we implemented video-based user verification based on the National Institute of Standards and Technology (NIST) standards for mid-assurance identity proofing (NIST SP 800-63-4). This solution helps to verify a user's identity, including remote employees, to enable users to self-service the setup of phishing-resistant credentials.

Third, we enforced users' access to resources per user and device groupings using Microsoft Entra ID's Conditional Access capability. Throughout this process, we partnered closely with the Microsoft global help desk team to perform constant temperature checks on incoming support-ticket levels and to unblock users as issues were identified.

Customer guidance

Strengthen your organization's security posture by implementing these best practices for phishing-resistant authentication.

Define security standards

1. **Phishing-resistant methods:** Employ phishing-resistant authentication methods such as certificate-based authentication, Windows Hello for Business, macOS platform Single Sign On (SSO), FIDO2 security keys, and Microsoft Authenticator passkeys. Use Microsoft Entra Conditional Access and Authentication Strengths controls to enforce phishing-resistant methods for access.
2. **Identity proofing:** Implement strong identity-proofing solutions in user onboarding and recovery processes, using capabilities like Microsoft Entra's Temporary Access Pass.

Adopt and measure

1. **User personas:** Determine user personas and prioritize them based on risk and complexity.
2. **Platform readiness:** Verify that platform versions support phishing-resistant authentication.
3. **Track registration campaigns:** Educate users about phishing-resistant authentication, register credentials on each device they use (including portable devices), and monitor credential usage.
4. **Support monitoring:** Track support-ticket volumes, identify and address issues that completely block users.

Microsoft Entra customers can get started on their phishing-resistant authentication journey using the detailed deployment guide at <https://aka.ms/PasswordlessGuide>.

Azure Managed Identities for service-to-service authentication

Problem description

Secrets used for authentication of clients to Azure resources (such as Microsoft Entra app passwords, storage access keys, and storage SAS tokens) have risks of leaks and exfiltration when mishandled in operational processes (such as during resource setup, testing, or secret rotation). Attackers can use these secrets to move laterally, disrupt systems, or exfiltrate data. Azure Managed Identities for Azure resources is designed to replace these manual processes of managing client secrets in favor of a platform-managed identity for applications and resources. As we adopted managed identities broadly across Microsoft, we needed to address the following challenges:

- **Assignment of the resources to be updated and training of the engineering owners** that are responsible for adopting managed identities across the services at Microsoft.
- **Gaps in scenario support across the ecosystem** that limit adoption of managed identities in complex or unsupported platforms.
- **Implementation of policies and governance** to help confirm new resources will be set up in alignment with standards, eliminating regressions using traditional secrets.

Microsoft implementation

First, we identified all the resource types in our environments and offerings that provide and require secrets for workload authentication. These include Microsoft Entra ID apps, Azure Storage, SQL, and Cosmos DB, among others. We established a clear standard across Microsoft that client authentication for these resources must move to Azure Managed Identity.

Second, we generated action items for all teams, pairing them with detailed guides for each resource type to help developers update their services. We tracked the adoption progress across the company and regularly reviewed status and obstacles with organization leaders. Product owners responsible for addressing obstacles, such as adding support in SDKs or product offerings to support Microsoft Entra ID authentication, provided regular status updates on their plans to support broad adoption.

Third, we built support for policies to prevent resources being created with client secrets and to block regressions for resources already using managed identities. These policies enabled us to audit and track all resources that allow traditional client secrets and enforce the standard for blocking their usage.

Customer guidance

Strengthen your organization's security posture by implementing these best practices for client authentication using Azure Managed Identities instead of client secrets.

Define security standards

- **Azure Managed Identity:** Use Azure Managed Identity for service-to-service (S2S) authentication to Microsoft Entra ID apps and Azure resources. Use resource settings and Azure Policy to disable use of local authentication methods/client secrets to help ensure resources remain aligned with the standard.

Adopt and measure

1. **Identify secret usage:** Use Azure Resource Graph and Microsoft Graph to identify secret usage across resource types.
2. **Adopt Managed Identities:** Implement Azure Managed Identity–based solutions and safely deploy those updates, maintaining failback mechanisms to help avoid service disruptions.
3. **Verify migration:** After migration, verify applications are no longer using unmanaged secrets and that all authentication paths have been updated.
4. **Disable local authentication:** Once verified, set “disable local authentication” to true on migrated resources.
5. **Policy enforcement:** Create Azure policies to verify that new resources are created in a compliant state and existing resources maintain compliance.

We recommend customers start adopting Azure Managed Identity for all S2S authentication flows. This feature is available for most Azure resources today. Azure Managed Identity as a federated identity credential for Microsoft Entra ID applications will soon be available in public preview. For more information, visit <https://aka.ms/AlwaysUseManagedIdentity>.

2. Protect tenants and isolate production systems

In September, we shared insights into the progress made in Pillar 2. In this update, we focus on work done to reduce risk from production and productivity devices and include customer guidance.

- Attacks using **device-based compromise to pivot to other environments** and gaps in access control policies represent frequently used tactics by threat actors we track.
- To help secure customers, we have [introduced a Microsoft Entra Conditional Access template](#), currently in public preview, which requires device compliance.
- To mitigate these risks, we have deployed 98,000 production-ready, locked-down devices. We've also moved 28,000 high-risk users to a customized and locked-down VDI solution. These measures help us **enforce strict device compliance** to protect against and limit the impact of device compromise on user identity.
- Customers can mitigate these risks by applying the new Microsoft Entra Conditional Access template, **implementing strict device compliance policies**, and increasing the **segmentation and isolation** of their systems.

Trusted devices for developers and employees

Problem description

Using a single device for both production and non-production access can increase security risks. If a non-production identity is compromised, attackers may exploit it to gain control of the device, leading to lateral movement into production environments. This poses significant threats as all identities on the device become potential targets. Attackers can manipulate the device remotely, steal tokens, and exfiltrate data, potentially pivoting into production environments even with low-privilege operations.

Microsoft implementation

First, we reviewed the device policies for previously established segmented user-identity directories for production access. These production identity systems are distinct from the employee corporate directory. By restricting production account use to a known scope of trusted production-ready locked-down machines, we reduce the attack surface of production identities.

2. Protect tenants and isolate production systems

Second, we accelerated the issuance of production-ready locked-down devices to new populations of users. We distributed an additional 8,000 new devices since the September update, bringing the total fleet of production access devices to 98,000.

Third, we used Microsoft Entra ID's Conditional Access capability to enforce user authentication for production identities to originate from trusted locked-down physical devices or combination of physical and VDI endpoints.

Customer guidance

To strengthen your organization's security posture, consider these best practices for safeguarding production environments accessed by employees.

Define security standards

1. **Device security:** [Secure devices as part of your privileged access controls](#) to protect user identities on the device.
2. **Application control:** Devices should restrict software to a minimal scope. For production access devices, exclude non-essential applications and utilities potentially vulnerable to malware and phishing attacks.
3. **Governance:** Applications and software versions must undergo careful security assurance scrutiny before addition.
4. **Least privilege:** In addition to user authorization, device compliance adds defense-in-depth to the implementation of least-privilege access.

Adopt and measure

1. **Identity separation:** Recognize that device security is part of the identity perimeter for any environment. Create separate Microsoft Entra tenants for environments intended to be segmented. Establish separate production identities for production environments, segmented from non-production environments.
2. **Device compliance policies:** Test and implement the new Microsoft Entra Conditional Access template in public preview, [review device compliance policies in Microsoft Intune](#), and set rules for device access for each environment.
3. **Policy enforcement:** Apply Conditional Access policies and [cross-tenant access settings](#) to enforce segmentation and device compliance for access.

By implementing these measures, organizations can help reduce the attack surface of production identities and strengthen their overall security posture.

3. Protect networks

In September, we shared insights into the progress made in Pillar 3. In this update, we focus on the work done to improve network security to reduce risk and include customer guidance.

- Network security vulnerabilities can be exploited by threat actors **to gain unauthorized access to protected resources** and move laterally across networks.
- To help customers, **Azure Virtual Network Encryption** is generally available in all regions and Domain Name System Security Extensions support is available in public preview. Developers can use the Bastion Developer SKU to establish secure connections without exposing public IPs.
- To mitigate these risks, we have increased **isolation and segmentation** of both management (control-plane) and services (data-plane) and improved network device health and safety. We also inventoried **over 99.3% of physical assets** and implemented mandatory access control lists (ACLs) to isolate the management of these assets.
- Customers can also mitigate network security risks by **hardening their physical infrastructure, restricting network access to permitted locations, and deploying network defenses** for their cloud and on-premises workloads.

Isolating and segmenting networks

Problem description

Threat actors exploit network security vulnerabilities to access protected resources, often moving laterally across them. Physical infrastructure risks stem from weak security practices, including poor password management and misconfigured firewalls. Virtual resources such as virtual machines (VMs) and containers are vulnerable when provided insufficient network isolation and overly permissive connectivity, potentially exposing them to untrusted networks and risking data exfiltration.

Microsoft implementation

First, we identified, inventoried, and attributed accountability for more than 99.3% of physical assets, establishing a foundation for network security. We also implemented service tagging of new IP address allocations, enabling traffic identification across the Microsoft network up to the service, subscription, and VM levels. These capabilities help detect malicious activity and simplify ACL management for both infrastructure and services.

Second, we hardened network devices and improved network device lifecycle management policies. We developed a scalable and automated method for secret rotation, making secrets unique per device. Network devices were onboarded onto our automation platform for continuous rotation. We implemented unique per-device authentication and one-time passwords for service accounts, to reduce blast radius in case of secret leakage or exposure, and implemented micro-segmentation ACLs to further secure the management of the network.

Third, we removed exception cases for mandatory controls, such as limiting access to a known scope of trusted production-ready locked-down machines plus VPN to the backend network, and we continue to accelerate the deployment of per-device access controls on backend datacenter devices, alongside developing short-lived credentials for devices, to significantly reduce the impact of exposed secrets.

Customer guidance

As customers grow their workloads in Azure, they may face similar challenges with IP address management and network isolation. We recommend the following best practices:

Define security standards

1. **Identify:** Have an inventory of all devices running in your physical network and resources in the cloud. Confirm proper IP allocation and attribution of IPs in your network so you can identify the source and destination of any traffic flowing in your network.
2. **Isolate:** Apply micro-segmentation in your network and least-privilege access. This extends to the physical network in your on-premises environment and your network resources in the cloud, including infrastructure-as-a-service and platform-as-a-service resources. Use least-privilege access policies to prevent lateral movement within your network.
3. **Layered defense:** Build a layered defense approach using monitoring, ACLs/network security groups, network address translation gateways, and bastions to help secure your network terrain.
4. **Automation and safe deployment:** Take advantage of automation and safe deployment practices to implement changes gradually and safely.
5. **On-premises network security:** Apply similar security measures to on-premises networks, including secrets management, secrets rotation, and least-privilege access to network devices.

Adopt and measure

1. **Inventory and track resources in your network:** Create an inventory of your network assets and use tooling like IP address management to manage and track your IP allocations.
2. **Follow Microsoft security benchmarks:** The [Azure Security Benchmark v3—Network Security](#) provides a comprehensive overview of the key features and benefits, detailed instructions on how to set it up and use it effectively, troubleshooting tips for common issues, and best practices to maximize its potential.
3. **Protect your network perimeter:** Apply ingress and egress control using Azure Bastion and Azure Network Address Transition Gateway, implement intrusion detection/prevention systems to monitor for malicious activity complemented by endpoint solutions like Microsoft Defender, enable Azure DDoS Protection to defend against large-scale attacks, and protect applications and APIs against attacks by using a web application firewall (WAF) like Azure Application Gateway or Azure Front Door.
4. **Implement lifecycle management practices:** To help to ensure your network devices are compliant and secure, periodically rotate all secrets via automation, making sure they are unique and have a short lifespan. Regularly update network device firmware to fix security vulnerabilities. Design secure ACLs to limit lateral movement within your network and ensure visibility of communications. Invest in telemetry, monitoring, and alerting for relevant security events to proactively prepare for incident response.

By implementing these network security measures, customers can work towards building a more resilient and defensible network infrastructure. While the journey to enhanced network security may be complex, the tools and practices outlined above provide a solid foundation for addressing current and emerging network threats in the cloud and on-premises.

4. Protect engineering systems

In September, we shared insights into the progress made in Pillar 4. In this update, we focus on the work done to reduce the risk of secrets and credentials in code and include customer guidance.

- Secrets and credentials play a crucial role in **granting privileged access to applications and data and facilitating communication between systems**. Exploiting exposed credentials and gaining lateral movement are behaviors frequently used by threat actors we track.
- To help customers, GitHub Advanced Security scans GitHub and Azure DevOps Git repositories to detect and report exposed secrets, and also **blocks new secrets from being exposed at push**. GitHub Advanced Security is free for public GitHub repositories.
- To address these risks, we are working to **remove secrets from** code and other unsecured storage and transmission methods and have **implemented standards of strong authentication protocols** that do not rely on weak mechanisms such as plaintext credentials and that actively detect, block, and remove exposed secrets and credentials.
- Customers can mitigate these risks by blocking secrets at push and implementing strong authentication protocols to **reduce the potential of leaked credentials and secrets** leading to successful attacks.

Removing secrets from code

Problem description

Secrets or credentials are commonly used in the process of development. API keys, database credentials, identity and access management permissions, secure shell keys, certificates, and more are necessary for applications and systems to talk to each other in a secure manner and have become increasingly more common with the rise of DevOps and cloud computing. Across the industry, developers often use poorly secured methods in passing these credentials to the applications—such as choosing a weak, shareable credential like a password, storing the applications certificate on their personal development machine, sharing them via email, or checking the credential into the code—where anyone with access to the repository can extract.

Attackers always look for systems of access. For engineering systems, this translates into attackers looking for vulnerabilities like plaintext credentials that may have been inadvertently checked in or improperly stored. They use these credentials to access applications, sometimes gaining elevated access to sensitive resources. These could be resources in production, such as a critical storage account containing customer data or a seemingly non-production resource. But even in the latter

4. Protect engineering systems

case, we've seen threat actors pivot from non-production resources to production resources due to unexpected links to other identities or other resources with production access.

To reduce these risks, we accelerated our secret detection and remediation program through improvements in our ability to prevent, detect, and revoke credentials in source code.

Microsoft implementation

First, we implemented managed identities where possible. The most significant solution was to switch to passwordless authentication when supported—such as Azure Managed Identities—as they eliminate the risk of leaking credentials altogether.

Second, we focused on live secrets. The number of findings in full scans of git history were in the tens of thousands. We intentionally focused on secrets that were proven to be live to drive effective remediation. This minimized false positives and resulted in faster reaction time from developers, because, when a live secret is detected, they trust that it must be remediated.

Third, we drove a conversion for service providers at Microsoft that produce plaintext credentials to generate [highly identifiable secrets](#), rather than custom strings. This allowed us to have deterministic detections with zero false positives and enabled us to block any operation that introduces such a credential.

Fourth, we blocked all highly identifiable secrets at git push in all of our Azure DevOps repositories. Git push is an operation from developers when, after a coding session on their machine, they will push their work to the cloud-based repository. Blocking this operation, with no option to suppress or bypass the finding, prevents the credential from being committed to the cloud, hence ensuring a clean source-code control history. We have blocked close to 3,000 credentials from entering our source code this way. This provides a much easier remediation as a development environment secret can be simply deleted at that stage, instead of requiring a rotation once it is committed to code. In addition, we frequently scan the repositories, including git history, for any credential that might have been missed, and especially for legacy credentials that are not highly identifiable.

Fifth, for any plaintext credential that is determined to be live, we find attribution—who owns the resource protected by this secret—and require a fast remediation, which our system systematically verifies through a regular liveness check.

We will continue to work with Microsoft services teams and third-party providers to ensure full support for passwordless authentication protocols which will enable us and our customers to remove all secrets from code.

Customer guidance

Define security standards

As customers create more software in DevOps platforms which build applications and services that take advantage of cloud resources, the persistence of credentials will unintentionally proliferate. To reduce the risk of leaked credentials, we recommend the following:

1. **Wherever possible, adopt secretless authentication mechanisms:** To enhance security, customers should use alternative authentication mechanisms such as Microsoft Entra authentication, OAuth, or Azure Managed Identities instead of relying on plaintext credentials like personal access tokens (PATs).
2. **Adopt highly identifiable secrets:** For any systems that still require tokens or other plaintext credentials to function, those credentials should be generated by service providers to a unique and highly identifiable standard. User-proposed passwords or passphrases should be eliminated or minimized, as these are not as secure as automatically generated formats due to the potential for reuse. Customers should also immediately revoke any secrets that are no longer needed or are suspected to be compromised.
3. **Use secure vaults to store secrets:** Customers should store all credentials in secure vaults like Azure Key Vault or GitHub Secrets and never hard-code them in source code or configuration files. As part of a Zero Trust standard, the critical initial step is to establish the policy that "source code must not contain credentials allowing access to resources." Regular monitoring and auditing of secret usage are essential to ensure tokens are used as intended and to set up alerts for unusual activity.
4. **Apply least privilege:** Grant only the necessary permissions for the task and set short expiration times for any credentials to enforce rapid renewal and reduce exposure. Use clear and descriptive names when creating credentials to help identify and manage tokens during auditing. These best practices help mitigate the security risks associated with plaintext credentials like PATs.

4. Protect engineering systems

Adopt and measure

Continually scanning and preventing the persistence of credentials reduces risk.

1. **Continually scan for credentials in source:** All repositories should be covered by a credential scanning methodology, including test, personal, and experimental repositories. While other security vulnerabilities may be less critical in applications that have been retired, or in test code, exposed credentials can be taken advantage of from production and non-production resources alike. Just because the source files may not be critical, it does not mean that they might not contain critically important secrets to detect.
2. **Turn on push-protection:** By turning on push-protection in products such as GitHub Advanced Security (which can be used for hosted repositories in either GitHub or Azure DevOps), customers can prevent secrets from entering source code control repositories and git history, eliminating the possibility of leaking credentials.
3. **Prioritize live secrets:** Detected credentials should take advantage of validity checking to prioritize remediation. Credentials that are valid may be used by anyone who has obtained access to the source to compromise other systems. Validity checking is also useful in ensuring that developers have truly revoked a credential. Only once the credential is no longer valid may you consider the finding closed, regardless of user attestation.
4. **Build a robust remediation and monitoring process:** Any detected credential persisted outside an approved store should be considered compromised and revoked immediately. Revoking a credential is the only way that you can be sure it is unable to be used by a malicious actor. If a replacement credential is needed, make sure that it is stored and accessed securely moving forward. Post-revocation, inspect audit logs for potentially suspicious use of the credential and investigate.

By implementing these practices, organizations can work towards improving their use of strong authentication practices in code, while maintaining a balance between security needs and developer productivity. GitHub Advanced Security [code scanning](#) is the preferred tool as it can detect secrets at push. Additional tools like GitHub [CodeQL](#) and Microsoft Defender for Cloud [DevOps](#) provide functionality later in the engineering cycle. Effective monitoring and detection of exposed secrets and credentials requires ongoing vigilance.

5. Monitor and detect threats

In September, we shared insights into the progress made in Pillar 5. In this update, we focus on the work done to improve security logging and include customer guidance.

- Security audit logs play a crucial role in detecting the behaviors frequently used by threat actors we track. They enable **real-time monitoring and detection of unauthorized access** or suspicious activities, helping to identify potential security breaches early. They also provide detailed records of who did what and when, expediting security investigations.
- To help customers, Microsoft [expanded cloud logging](#) to provide deeper security visibility, including detailed logs of more than 30 types of data and **standard log retention for 180 days**. These logs are available to Microsoft 365 customers by default and at no additional cost.
- As part of our improvements for threat detection and response, we have implemented standards which require **consistent audit logging across all services** and **log retention for two years**. We have already established central management and a two-year retention period for identity infrastructure security audit logs.
- Customers can implement logging standards to **improve threat detection and response**.

Standardized audit logging, collection, and retention

Problem description

Microsoft identified two key audit-logging challenges that created obstacles during security investigations and threat hunting: inconsistent logging formats and disparate log storage and retention.

Although many services captured audit logs, these logs were stored in different locations and lacked a standardized format. Incident responders had to spend significant time searching for logs, deciphering formats, and correlating events, which increased cognitive load and extended investigation times. Inconsistent logging formats also resulted in logging gaps, as some services did not log all the data needed for effective security investigations. This inconsistency prevented us from developing standard detections across services.

We also found inconsistent log retention across services, which can result in missing critical information needed for complete forensic analysis, such as initial access points, compromised systems and accounts, privileged actions, and lateral movement paths. This incomplete picture can hinder the determination of the attacker's methods and objectives, leading to inadequate

remediation efforts and potentially leaving vulnerabilities unaddressed. These issues can allow attackers to remain undetected longer, slowing down investigations and increasing the risk of data breaches.

Microsoft implementation

First, we standardized audit-logging capabilities, which is critical for comprehensive threat detection. We developed a standard audit library to unify logging practices across services, and we identified and fixed gaps in telemetry data as part of adopting the standard library.

Second, we implemented centralized log collection to enhance visibility and analysis capabilities. We created specialized accounts for security investigators, overcoming previous identity and access challenges and enabling deeper investigations. This centralization improves threat detection and investigation efficiency and enhances event correlation across systems.

Third, we extended log retention, a crucial requirement for thorough investigations. Extended retention allows for better detection of long-term attack patterns. It also aids in forensic analysis of historical security incidents.

Fourth, we enhanced threat detection with improved machine learning and AI. By integrating advanced analytics and machine-learning models into our threat detection systems, these enhancements help identify complex and evolving attack patterns. Continuous refinement of detection algorithms improves accuracy over time.

Customer guidance

To enhance threat monitoring and detection capabilities, we recommend the following practices:

Define security standards

1. **Use a standard log library:** All services and applications in your environment should utilize a common audit library to maintain a consistent logging scheme. This includes standardizing the format for key information such as user identities, actions taken, event timestamps, and methods used. Check that key security fields are defined as part of this standard log library.
2. **Centralized log storage:** Make certain all logs from services are directed to a central location or can be accessed through a central repository.
3. **Maintain log integrity and security:** Ensure the audit logs are immutable. To achieve this, make sure the data is encrypted both in transit and at rest. Additionally, storing the logs in an immutable storage solution helps maintain data integrity.

Adopt and measure

1. **Inventory and coverage of audit logs:** Maintain an up-to-date inventory of all applications and services within your organization. Confirm that all APIs across your services log audit data. Monitor API logging coverage (that is, the percentage of API actions being logged) to detect any gaps and regressions.
2. **Review and analyze logs:** Set up real-time monitoring and alerts for coverage gaps and unexpected or suspicious log-access activities. Include appropriate logs in threat-modeling requirements. Conduct regular reviews by security response or investigator teams to verify logs capture the necessary investigation information.
3. **Enable audit logs for cloud infrastructure and applications:** This provides complete end-to-end visibility during security investigations. Use [Microsoft Purview auditing solution](#) for Microsoft Entra and Microsoft 365 applications logs (for example, Exchange Online, SharePoint, OneDrive). Enable Azure infrastructure audit logs centrally for your tenant by taking advantage of [Azure Policy](#). This helps your organization respond effectively to security events, forensic investigations, internal investigations, and compliance obligations.
4. **Adopt a centralized log management system:** Adopt a centralized log management system, such as [Azure Monitor](#), to store and analyze logs from across your environment. Use security information and event management solutions such as [Microsoft Sentinel](#) which enhance security investigations by correlating data from various sources, allowing investigators to easily query data and create quick detections. This unified view of activities will improve the efficiency of threat detection and investigation. Additionally, utilize automated tools to manage log retention and to help ensure compliance with retention policies.
5. **Develop threat detections:** Integrate the audit logs with threat intelligence feeds to develop security threat detections and utilize machine learning and analytics to identify anomalies. You can also correlate events from various services and applications to uncover complex attack patterns, and regularly review and update your detection mechanisms to adapt to evolving threats.

By implementing these practices, organizations can work towards improving their threat detection capabilities while maintaining a balance between security needs and privacy considerations.

6. Accelerate response and remediation

In September, we shared insights into progress made in Pillar 6. In this update, we focus on the work done to improve security incident response and communications and include customer guidance.

- We addressed 90% of vulnerabilities **within our reduced time to mitigate** for high severity cloud vulnerabilities accelerated response time. To further enhance our mitigation efficiency, we are investing heavily in automation by taking advantage of AI technologies and implementing process improvements across the company.
- Since the beginning of the year, we have published close to **800 Common Vulnerabilities and Exposures (CVEs)** to transparently communicate vulnerability mitigations and accelerate broad adoption. 12 of these were for cloud services vulnerabilities which require no customer action. These CVEs include Common Weakness Enumeration (CWE) and Common Platform Enumeration (CPE) annotations, which we adopted earlier this year. The result is more actionable and useful CVE reports for customers, regulators, and security vendors.
- We continue to improve how we **apply consistent standards for effective and accurate communication on security incidents**. This includes CVEs, nation state notifications (NSNs), and our goal of providing holistic security- and privacy-related notifications in product portals to our customer's security and privacy contacts.

Apply consistent standards for effective and accurate communication

Problem description

Security events require increased transparency and faster customer communication. Clear and actionable guidance should be distributed using well-understood delivery mechanisms to ensure that the relevant personnel and resources are aware of issues that may impact their organization and can react quickly.

Microsoft implementation

As we continue to learn from security incidents and incorporate customer feedback, we have launched a cross-company initiative to provide clear, timely, and actionable communications and customer engagement for security and privacy events. Our goal is to reimagine how we support customers by providing holistic security- and privacy-related notifications in product portals that are readily accessible to the appropriate organizational points of contact, such as the customer's CISO organization and the security operations center.

Earlier this year, we started enriching our CVE information with CWE and CPE annotations. We now publish CVEs for critical severity cloud vulnerabilities, even when customers do not need to manually implement a mitigation as Microsoft has already taken action to resolve the vulnerability.

If a CVE impacts an Azure service and requires customer action to mitigate, customers with affected resources will be notified in the Azure Service Health portal. The alerts will specify impacted resources at tenant or subscription level and provide mitigation guidance, actions required, and support avenues. These additions to our CVE program help provide customers and partners with transparent and actionable information on prevailing security concerns, even in cases where immediate action may not be required.

We increased the level of nation-state actor intelligence shared and improved our NSN process. NSNs are delivered to Microsoft customers when their Microsoft 365, Azure, or consumer email account is targeted or compromised by a nation-state-sponsored or affiliated cyberthreat actor (nation-state actor). Over the past few years, Microsoft has delivered over 20,000 NSNs. We synthesize the insights and threat intelligence from nation-state incidents in our annual [Microsoft Digital Defense Report](#) for our customers' and partners' benefits. Looking to the future, our goal is to customize solutions to help ensure we can deliver NSNs to customers through consistent and well-known product portals (for example, Microsoft 365 admin center, Azure Service Health) that will be prominently displayed and routed to security-focused customer contacts.

Customer guidance

Use these guidelines and best practices to prepare for future security incident communications.

Create Azure, Microsoft 365, and Dynamics 365 communication pathways

1. Stay informed about Azure security issues by [setting up alerts](#) (email, SMS, webhook) in [Azure Service Health](#).
2. Stay informed about Microsoft 365 and Dynamics 365 security issues by [setting up alerts](#) in the [Microsoft 365 admin center](#).
3. Create alert assignments specific to incident responders, vulnerability management, engineering, etc., for all tenants and subscriptions. Typically, bad actors target less-used resources, so it's important your security alerting covers all workloads.
4. Make sure you have contactable email addresses that are reviewed regularly for accuracy.
5. Build internal organizational rules and notifications to help ensure the appropriate privacy and security contacts are notified by your global admins of issues that have a security or privacy element.
6. For additional incident readiness guidance and best practices, visit <https://aka.ms/IncidentReadiness>.

Understand and prepare for NSNs

1. If you receive an NSN, it means that the Microsoft Threat Intelligence Center has observed activity indicating that a nation-state actor has targeted your Microsoft 365, Azure, or consumer email account, or that your account has been compromised by a nation-state actor. "Targeted" means the threat actor's attempt to compromise your account was unsuccessful; "compromised" means the attempt was successful, and you should take immediate steps to resecure your account and investigate the potential intrusion.
2. NSNs do not indicate that Microsoft systems, products, or services have been compromised.
3. NSNs provided to organizations include a brief description of the nation-state actor's observed behavior or activity and actionable information, such as the observed IP address used by the nation-state actor, to begin your investigation and address affected accounts or infrastructure.
4. Microsoft delivers NSNs to the targeted or compromised Microsoft 365 enterprise (organization) account with a direct phone call to the Tenant Global Administrator or another appropriate security contact at the customer organization. Typically, if we are unsuccessful in multiple attempts to reach the organization by phone, we will send the NSN via email. Microsoft delivers NSNs via email to consumer accounts.
5. Infrequently when a high-volume of NSNs are needed, such as during a broad-scale phishing attack, NSNs are delivered via the Microsoft 365 admin center. These message center notifications are visible to Tenant Global Administrators and users who have been assigned the message center privacy reader role for the affected tenant.

By understanding these notification practices, organizations can improve their readiness to respond to security notifications and incidents effectively.