



Zero Trust: A new era of security

Introduction

The digital security landscape is transforming in ways most couldn't predict as little as five years ago, pushing classic security strategies and tools to their breaking point. Attackers' ability to bypass conventional access controls is ending any illusions that traditional security strategies and tools can keep corporate resources secure in the digital age.

The root of the issue is that attacks readily exploit the assumption that assets are safe on a "trusted" corporate network.

Attackers are getting in the network

An attacker can compromise a single endpoint within the trusted boundary by using phishing attacks and then quickly expand their foothold across the entire network by using reconnaissance, credential theft, and lateral movement. While network controls can block and detect some classic attacks, they simply cannot ensure an entire network is trustworthy.

Assets are leaving the network

Some assets and devices no longer reside on the corporate networks. Today's employees are bringing their own devices and working remotely, data is being accessed outside the corporate network and shared with external collaborators, corporate apps are being hosted in a multi-cloud environment, and sensitive data is being stored in SaaS apps. As many of these interactions don't traverse networks controlled by the organization, large gaps in coverage occur with traditional network access controls.

Organizations that fail to evolve from traditional defenses are much more vulnerable to breaches. As such, organizations can no longer draw traditional perimeters around their networks and expect to stay secure. Fortunately, there is a way forward: Zero Trust.

In this e-book, we will discuss the core principles of the Zero Trust model for security and walk through how Microsoft can help with your Zero Trust security strategy.



Microsoft's evolution on networks

In Microsoft's own security operations, we're also finding that network controls are less and less effective at detecting threats.

As of the writing of this document, it has been over two years since the last primary detection of an attack on our corporate environment came from a network detection tool. Networking tools are still present, but their use is mostly for supporting investigations and advanced threat hunting.

Security for a world without boundaries

Zero Trust is an access model that assumes everything is on the open internet (an untrusted network), even resources behind the firewalls of the corporate network. We can no longer trust the integrity of the network, so operating under the principle of “trust no one, verify everything” helps us strengthen security. In practical terms this means before a user can access any resource, all requests are fully authenticated, authorized and encrypted in real time.



Rather than assuming a password is enough to identify a user, validate additional context (such as multi-factor authentication, normal sign-in location and time).

Instead of blindly trusting the user's device because it's on the corporate network, verify that the device is healthy and compliant.

Instead of allowing or denying full access to devices on the network, and only grant access to the specific services, applications, or data required, while continuously monitoring for suspicious activity.

Zero Trust shifts the focus of governing access from the network to intelligent access controls that take advantage of dynamic user and device risk signals and other telemetry to make more informed access decisions for an organization's data and resources on a case-by-case basis.

By gating access to individual resources using dynamic trust decisions, organizations can take advantage of fine-grain control to further refine user experience and security assurances.

Dual perimeter strategy

Organizations will need to operate two simultaneous access control regimes to manage the full set of assets in an enterprise estate:

Identity perimeter—a modern perimeter based on gating access to resources with identity controls. This approach is fundamental for securing new asset types, cloud-native apps, modernized LOB apps, third-party services, and other assets that are out of reach of the network perimeter.

Network perimeter—network controls to protect legacy assets that weren't designed for a hostile internet. Some organizations are using microsegmentation approaches to strengthen network access controls (which frequently adhere to the same Zero Trust principles described in this document).

Our experience securing global clouds has found that device, identity, and application specific controls provide higher confidence detections, better coverage, and richer context than network controls. As such, we have focused our security investments heavily on identity-based access controls. We will refer to those controls throughout this document.



At Microsoft, we've been on a Zero Trust journey for many years. The underlying idea and approach of Zero Trust was introduced in 2004 by a security consortium known as the Jericho Forum. They formalized and promoted the Jericho Commandments¹, laying out a framework for "de-perimeterisation" to support good security in this world of new platforms and work styles. The Zero Trust approach maps well to these principles and concepts.

Microsoft continues to evolve how we implement our security operations and design of our products using a Zero Trust model. This has resulted in the development of intelligent access controls integrated across our technologies and assets, and the ongoing use of attack simulation to improve security controls for today's cyber threats.


Microsoft's IT organization has been shifting our security strategy towards a full Zero Trust implementation for several years and has recently reached a major milestone of our internet-first initiative. We are in a production pilot to put client devices on the internet full time, which will shrink our network controls-based perimeter to only datacenters and supporting systems. This shift required multiple projects that each individually improved our security posture along the way—including implementing two-factor authentication and conditional access for iOS and Android devices, increasing device security assurances, modernizing applications to cloud (or publish to internet access), adjusting network bandwidth shape, and more.

1. Defining [Jericho Commandments PDF](#)



Implementing a Zero Trust security model

Migrating to a Zero Trust security model allows you to simultaneously improve security over conventional network-based approaches and better enable users where and when they need access. A Zero Trust model requires signal to inform decisions, policies to make access decision, and enforcement capabilities to implement those decisions effectively. Let's explore:



Signal to make an informed decision.

Zero Trust considers many signal sources—from identity systems to device management and device security tools—to create context-rich insights that help make informed decisions.

Decision based on organizational policy.

The access request and signal are analyzed to deliver a decision based on finely-tuned access policies, delivering granular, organization-centric access control.


Enforcement of the policy across resources.

Decisions are then enforced across the entire digital estate—such as read-only access to a SaaS app or remediating compromised passwords with a self-service password reset.

How Microsoft technologies enable a Zero Trust security model

At Microsoft, we believe a Zero Trust strategy can only be accomplished when there is end-to-end integrated coverage throughout the entire digital estate. We deliver that coverage by enriching each of these components of access control to deliver our end-to-end Zero Trust vision for our customers. The next three sections explore how we make this happen and how you can apply this strategy to your own organization.





Signal

When an access request comes in, we need to gather all the relevant context of the session to determine its overall risk and make an informed decision. This context is comprised of information from the access request itself as well as any other relevant information related to it—including the identity of the user plus the state of their device, the apps they're using, the sensitivity of the data they're trying to access, and more. This context is then used to deliver a dynamic trust decision to determine whether that particular session can be trusted at that time.

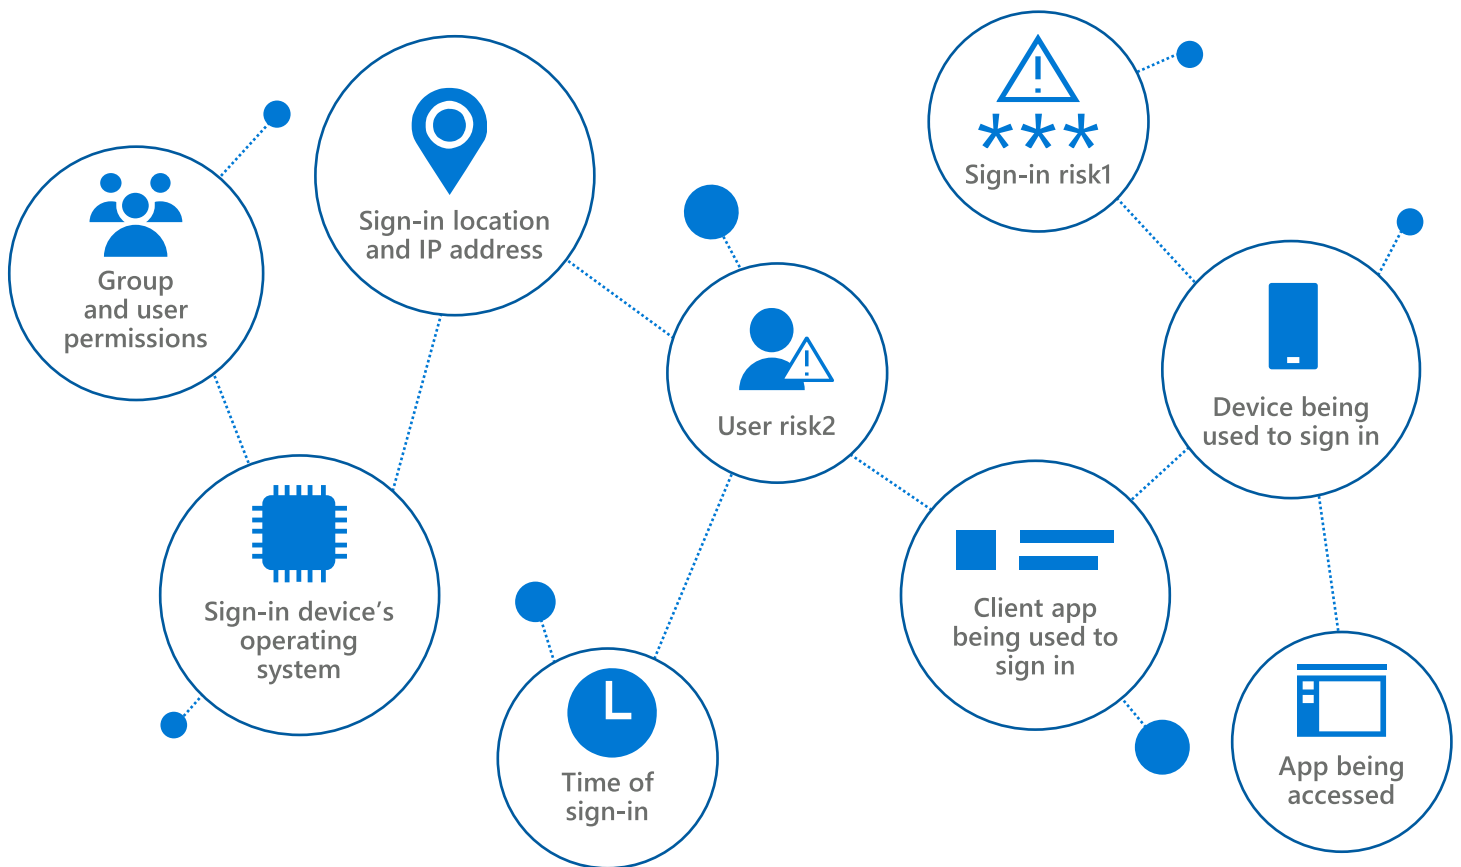
Using multiple rich signal sources, Microsoft can help you make more confident and informed access decisions that better protect your organization and lessen the impact on end-user productivity. Each signal source we gather reduces your exposure to risk. For example:

If a hacker is using **stolen credentials** to try and access your resources, behavior analytics can detect the abnormal IP address and time of login to prompt a multifactor authentication request.

If an attacker then steps up the attack and **compromises a legitimate device**, we can include device integrity into the session analysis and identify the attempt as high risk and deny the request.

At Microsoft, we rely on **Azure Active Directory (Azure AD)** to provide strong, adaptive, standards-based identity verification. In addition to the basic validation of usernames, passwords, and multi-factor or biometric authentication, Azure AD analyzes a variety of data to provide important context on user, device, location, and session risk for every access request.

This contextual data is used to make well-informed access decisions for organizational resources. Those data sources include:




Identity-related information is only part of the story; we also need to be sure that the device being used isn't compromised. **Microsoft Defender Advanced Threat Protection (MDATP)** and **Microsoft Intune** work together to provide critical device context used to determine trustworthiness.

Microsoft Intune provides a signal on whether the iOS, Android, Windows, or macOS device is compliant with the organization's policies (such as requiring personal devices to be enrolled in corporate device management before allowing access).

This assurance that Intune provides can be further strengthened by MDATP, which continuously monitors the state of devices, detecting compromises and flagging the devices' risk state to Intune.

1. Probability that the sign-in isn't authorized by the identity owner
2. Probability that a bad actor has compromised a given user

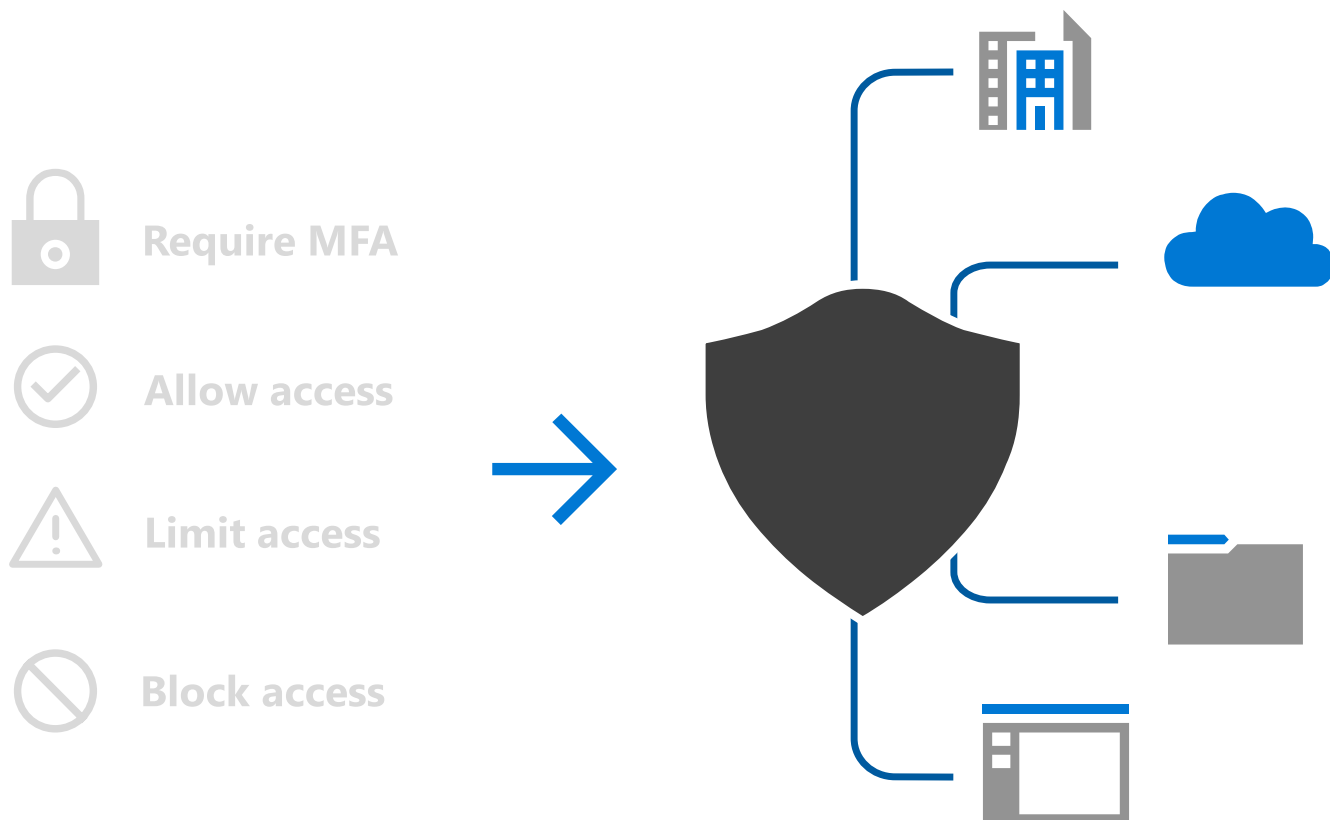


Decision

Now that we have all the information from the signals, we need to apply a policy to the access decision based on the organization's security posture and risk appetite. As mentioned earlier, trust decisions in the real world need more flexibility than a simple access/block decision in order to support end-user productivity.

Azure AD Conditional Access is the policy engine in our holistic Zero Trust approach. Conditional Access provides a flexible set of policies that can be configured by organizations to granularly control the circumstances in which users can access corporate resources. Conditional Access enables organizations to set policies based on the signal sources listed previously and then decide whether to (1) allow access, (2) deny access, (3) limit access, (4) require additional authentication challenges (such as multi-factor authentication), or (5) require additional remediations.

Conditional Access works robustly with any app or service configured for access with Azure AD—such as SaaS apps like Salesforce.com, custom apps running in the cloud, and of course Microsoft products and services. Azure AD can also be used to govern access to on-premises web apps via the application gateway. This helps ensure protection across the entire digital estate.



Enforcement

Because Conditional Access works with any application configured for access with Azure AD, it can deliver decisions across the environment with fine-grained control and seamless end-user experiences. Here are some ways it can deliver decisions across devices, web apps, and documents:

Extend policies to on-premises apps with [Azure AD Application Proxy](#), which unifies security in a hybrid environment.

Use [Microsoft Cloud App Security](#) to protect data in SaaS apps and enforce policies—such as blocking downloads, monitoring low-trust user sessions, applying read-only mode, and restricting user sessions from non-corporate networks.

Block or revoke access on compromised devices and remediate threats to bring devices back into compliance with Microsoft Defender ATP and Microsoft Intune.

Extend access control beyond data stores and apps with [Azure Information Protection \(AIP\)](#). AIP applies protections—such as encryptions, identity, and authorization policies—that stay with the document and emails, independently of their location.

Segment physical on-premises, hybrid, and cloud networks and gate individual access using user, device, and application-aware rules enforced using the least privilege access principle.



Closing thoughts

Zero Trust turns the complexity of the modern workplace into an asset. Every new point of connection provides rich signal for a smart policy engine to make informed access decisions. And paired with robust enforcement, you can ensure that the right people are getting the right level of access across the digital estate, improving both security posture and end-user productivity.

As you consider your organization's journey to Zero Trust, know that it typically has multiple steps. While identity is always going to be the foundational component, it's important to be thoughtful about choosing your investments and aligning them with your current business needs. Consider how to get quick wins and incremental value for each investment. Ultimately, every step forward will make a difference in reducing risk and returning trust in the integrity of your digital estate.

Resources

To learn more about Zero Trust, download our [infographic](#) and watch our [video](#)

For more on Microsoft IT's Zero Trust journey, read [Implementing a Zero Trust Security Model at Microsoft](#)

Visit our [Microsoft Security website](#) for more information on our approach to security