THE 6 PILLARS OF ZERO TRUST

fedscoop

Microsoft Azure

WHY AGENCIES NEED TO TAKE A MULTIDIMENSIONAL VIEW TO ZERO TRUST

The importance of authenticating trust not just within networks, but as part of a holistic Zero Trust architecture.

By *Susie Adams* and *TJ Banasik*

Susie Adams is Chief Technology Officer for Microsoft Federal.

TJ Banasik, CISSP-ISSEP, ISSAP, ISSMP, is Sr. Program Manager, Microsoft Azure Government End-to-End Customer Engineering.

Cybersecurity has always been about protecting our digital resources, our IT systems and of course, the people who depend on them. Within cybersecurity circles, that's traditionally meant focusing on three things: the need to "protect, detect and respond."

In practice, however, those of us working in enterprise IT departments have traditionally concentrated our time and resources on protecting networks and devices. We didn't spend a lot of time on detection, either because of other priorities or because there were only so many resources. The result for agency and program managers, as well as IT departments all too often, has meant having to respond to security incidents reactively instead of proactively.

IT departments in the commercial sector, and increasingly in the government sector, are coming to terms with the need to take a fundamentally different approach to security. It's not just that the bad guys are getting smarter or learning to leverage the cloud as fast or faster than we are.

It's also the fact that the modern enterprise perimeter is increasingly virtual. It's no longer just limited to the physical assets in your data center. Data and applications now live on premises and in multiple clouds — some you control, and some you count on third parties to manage.

In this new world, the security paradigm has changed: Identities are the new firewall, devices are the new perimeter and "assume breach" is the new security model.

> *In this new security world, identities are the new firewall, devices are the new perimeter and 'assume breach' is the new security model.*

It's against that backdrop that the federal IT community, like the commercial sector before it, has begun to embrace the principles of Zero Trust security. It's a world where we start with the premise that data systems are open to the internet and therefore, trust must be established and explicitly verified where and whenever a resource is required.

However, there's a tendency to think about Zero Trust too narrowly — that so long as we apply Zero Trust to our networks and follow the security guidance from NIST and OMB, our systems will be secure.

We would respectfully argue that Zero Trust is much bigger than that.

We believe that Zero Trust has to be established around six foundational pillars, not just your networks. These pillars represent the primary dimensions of an enterprise's entire digital estate, regardless of where that estate actually exists, and involve attaining trusted assurance around identities, devices, apps and APIs, data, and infrastructure as well as your networks.

We offer a more detailed look at those six pillars and discuss what Zero Trust maturity looks like for each of them in the next article of this report.

The idea is that in order to truly establish a Zero Trust environment, we have to take steps to assure the trustworthiness of users and their activities within and across all six of these dimensions, not just within networks, as part of a holistic Zero Trust architecture.

For those who are familiar with the Cybersecurity and Infrastructure Security Agency's latest

"Trusted Internet Connections 3.0 Reference Architecture," and CISA's *"TIC 3.0 Interim Telework Guidance,"* as well as NIST's *Zero Trust Architecture draft* 800-207, these Zero Trust principles will go a long way in improving your security posture, whether you're using Microsoft Azure or other cloud services, because the principles are universal.

The big question, of course, that every federal IT leader will ask is, 'How do I actually go about doing this in a world where I barely have enough security folks let alone those who really understand cloud?"

The short answer is, start now by focusing on what matters most to your agency and what's already high on your priority list. Then develop a plan to move up the Zero Trust maturity model for each pillar, featured on the pages that follow in this report.

Many agencies, for instance, are already devoting a good deal of resources to meet federal mandates around identity and access controls. The more granular you can get with conducting security access reviews and determining conditional access, using cloud-based tools, such as Azure Active Directory, and Zero Trust principles, the faster you can narrow your attack surface.

Another opportunity for agencies that are modernizing is to give greater thought to network-based segmentation and control. The more you can develop your microperimeters and segment your assets, the more policy enforcement points you can put into place to segment, validate and control traffic. Cloud-based security is much different than on-premises security, so consider using cloud workload protection platforms and security management tools, such as Azure Security Center, to harden your cloud workloads.

As you consider these and other measures, give proper attention to employing a modern security information event management platform that can operate in a multicloud environment and orchestrate and automate policies quickly. On average, compromised credentials will be used within 24 hours. A skilled attacker or advanced persistent threat group will be able to compromise domain administration credentials within about four days.

We've put together an entire series of blogs on a variety of Zero Trust use cases, with practical steps and recommendations, to help agencies get started. You can see a preview of those blogs featured in the pages of this report.

We've also put together a *blueprint for automating* Zero Trust cybersecurity tasks that allows you to deploy policy objects in your environment and takes much of the guesswork out of adopting Zero Trust tools.

The bottom line is the traditional enterprise perimeter is fast dissolving. While Zero Trust is undoubtedly the way forward, agencies should also be mindful to see all the dimensions to which Zero Trust applies, beyond Zero Trust networks, while at the same time, seeing the possibilities for making more immediate gains with the right tools.

For the cloud and tools to enable your agency's Zero Trust framework, learn more about Azure Government at azure.com/gov

THE 6 PILLARS OF ZERO TRUST AND WHERE TO START

How to assess your Zero Trust readiness and maturity.

Cloud applications and the mobile workforce have redefined the security perimeter. Employees are bringing their own devices and working remotely. Data is being accessed outside agency and enterprise networks and shared with external collaborators such as partners and vendors. Enterprise applications and data are moving from on-premises to hybrid and cloud environments.

The new perimeter isn't defined by the physical location(s) of the organization—it now extends to every access point that hosts, stores, or accesses corporate resources and services. Interactions with agency resources and services now often bypass on-premises perimeter-based security models that rely on network firewalls and VPNs.

Agencies which rely solely on on-premises firewalls and VPNs lack the visibility, solution integration and agility to deliver timely, end-to-end security coverage.

Today, organizations need a new security model that more effectively adapts to the complexity of the modern environment, embraces the mobile workforce, and protects people, devices, applications, and data wherever they are located. This is the core of Zero Trust.

Understanding Zero Trust

Instead of believing everything behind the enterprise firewall is safe, the Zero Trust model assumes breaches will occur and verifies each resource request as though it originates from an uncontrolled network. Regardless of where the request originates or what resource it accesses, Zero Trust teaches us to "never trust, always verify."

In a Zero Trust model, every access request is strongly authenticated, authorized within policy constraints and inspected for anomalies before granting access. Everything from the user's identity to the application's hosting environment is used to prevent breach. We apply micro-segmentation and least privileged access principles to minimize lateral movement.

Finally, rich intelligence and analytics powered by artificial intelligence (AI) help us identify what happened, what was compromised, and how to prevent it from happening again. To successfully leverage threat intelligence, you must have a large, diverse set of data to apply to your processes and tools. Microsoft's capabilities are powered by the data, machine learnings, and human insights to enable richer intelligence.

Guiding principles of Zero Trust

1. Verify explicitly.
Always authenticate and authorize based on all available data points, including user identity, location, device health, service or workload, data classification, and anomalies.

2. Use least privileged access.
Limit user access with Just-In-Time and Just-Enough Access (JIT/JEA), risk-based adaptive polices, and data protection to protect both data and productivity.

3. Assume breach.
Minimize blast radius for breaches and prevent lateral movement by segmenting access by network, user, devices, and application awareness. Verify all sessions are encrypted end to end. Use analytics to get visibility, drive threat detection, and improve defenses.

Controlling access with policy

Today, organizations need to be able to provide secure access to their resources regardless of user or application environment. Before we allow access, we want to assess a user's location, their role in the organization, the health of their device, the type of service and classification of the data they're requesting access to, and more. To do this effectively, we need to use signal and automated policy enforcement to deliver the right balance between security and optimal user experience.

A Zero Trust security model relies on automated enforcement of security policy to ensure compliant access decisions throughout the digital estate. The framework of controls built into your security solutions and tools enables your organization to fine-tune access policies with contextual user, device, application, location, and session risk information to better control how users access corporate resources and backend resources communicate. These policies are used to decide whether to allow access, deny access, or control access with additional authentication challenges (such as multi-factor authentication), terms of use, or access restrictions.

Building Zero Trust into your agency

A Zero Trust approach should extend throughout the entire digital estate and serve as an integrated security philosophy and end-to-end strategy. This is done by implementing Zero Trust controls and technologies across six foundational elements. Each of these six foundational elements is a source of signal, a control plane for enforcement and a critical resource to be defended and a focus for investment:

Identities

Whether they represent people, services, or IOT devices – define the Zero Trust control plane. When an identity attempts to access a resource, we need to verify that identity with strong authentication, ensure access is compliant and typical for that identity, and follows least privilege access principles.

Devices

Once an identity has been granted access to a resource, data can flow to a variety of different devices—from IoT devices to smartphones, BYOD to partner managed devices, and on-premises workloads to cloud hosted servers. This diversity creates a massive attack surface area, requiring we monitor and enforce device health and compliance for secure access.

Applications and APIs

provide the interface by which data is consumed. They may be legacy on-premises, lift-and-shifted to cloud workloads, or modern SaaS applications. Controls and technologies should be applied to discover Shadow IT, ensure appropriate in-app permissions, gate access based on real-time analytics, monitor for abnormal behavior, control of user actions, and validate secure configuration options.

Data

Ultimately, security teams are focused on protecting data. Where possible, data should remain safe even if it leaves the devices, apps, infrastructure, and networks the organization controls. Data should be classified, labeled, and encrypted, and access restricted based on those attributes.

Infrastructure

Whether on-premises servers, cloud-based VMs, containers,or micro-services, infrastructure represents a critical threat vector. Assess for version, configuration, and JIT access to harden defense, use telemetry to detect attacks and anomalies, and automatically block and flag risky behavior and take protective actions.

Networks

All data is ultimately accessed over network infrastructure. Networking controls can provide critical “in pipe” controls to enhance visibility and help prevent attackers from moving laterally across the network. Networks should be segmented (including deeper in-network micro segmentation) and real-time threat protection, end-to-end encryption, monitoring, and analytics should be employed.

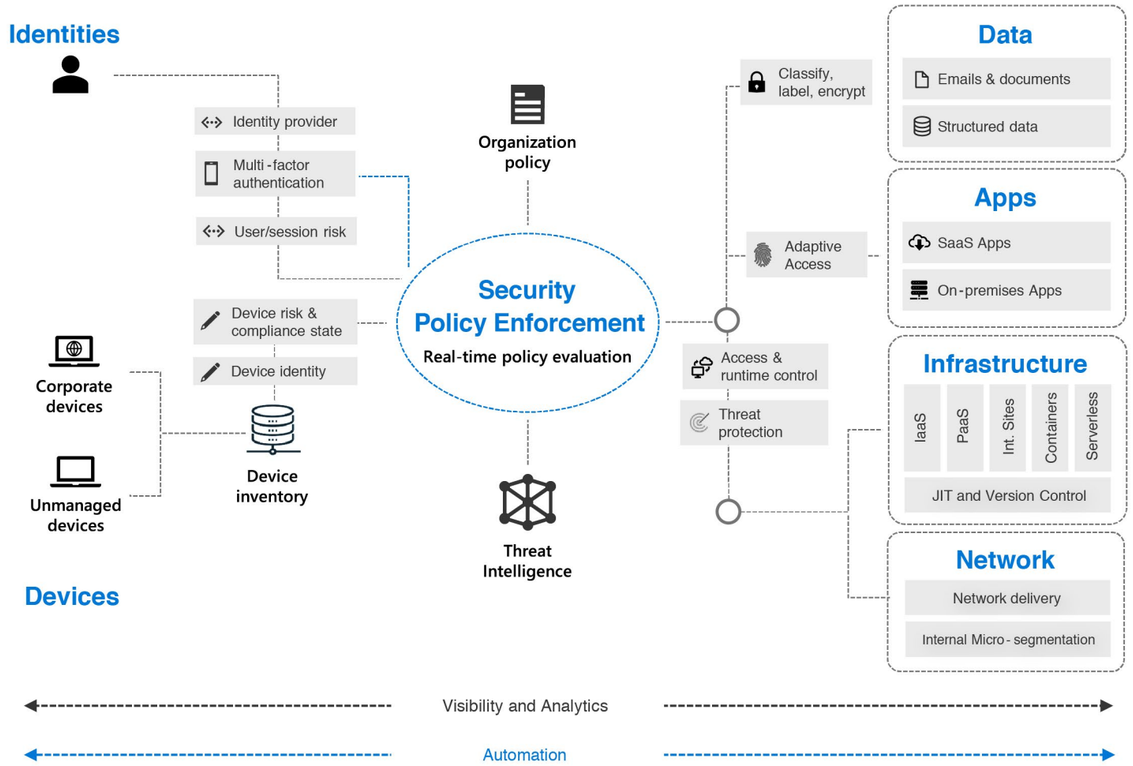

Zero Trust across the digital estate

In an optimal Zero Trust implementation, your digital estate is connected and able to provide the signal needed to deliver end-to-end coverage and make informed access decisions using automated policy enforcement.

Improving visibility and embracing security automation

Because Zero Trust relies heavily on signal and solution integration to be successful, this is a great time to work towards providing greater visibility into your threat landscape and embracing security automation. The Security Operations Center (SOC) should have a multi-tier incident response team in place that uses advanced threat detection and AI-driven alert management capabilities to cut through the noise and deliver prioritized security alerts. Response to common incidents, such as denying access to infected devices, should be automated to improve response times and reduce risk exposure.

MATURITY MODEL

TRADITIONAL
This is where most organizations generally sit today if they haven't started their Zero Trust journey:

- On-premises identity with static rules and some SSO.
- Limited visibility is available into device compliance, cloud environments, and logins.
- Flat network infrastructure results in broad risk exposure.

ADVANCED
In this stage, organizations have begun their Zero Trust journey and are making progress in a few key areas:

- Hybrid identity and finely-tuned access policies are gating access to data, apps, and networks.
- Devices are registered and compliant to IT security policies.
- Networks are being segmented and cloud threat protection is in place.
- Analytics are starting to be used to assess user behavior and proactively identify threats.

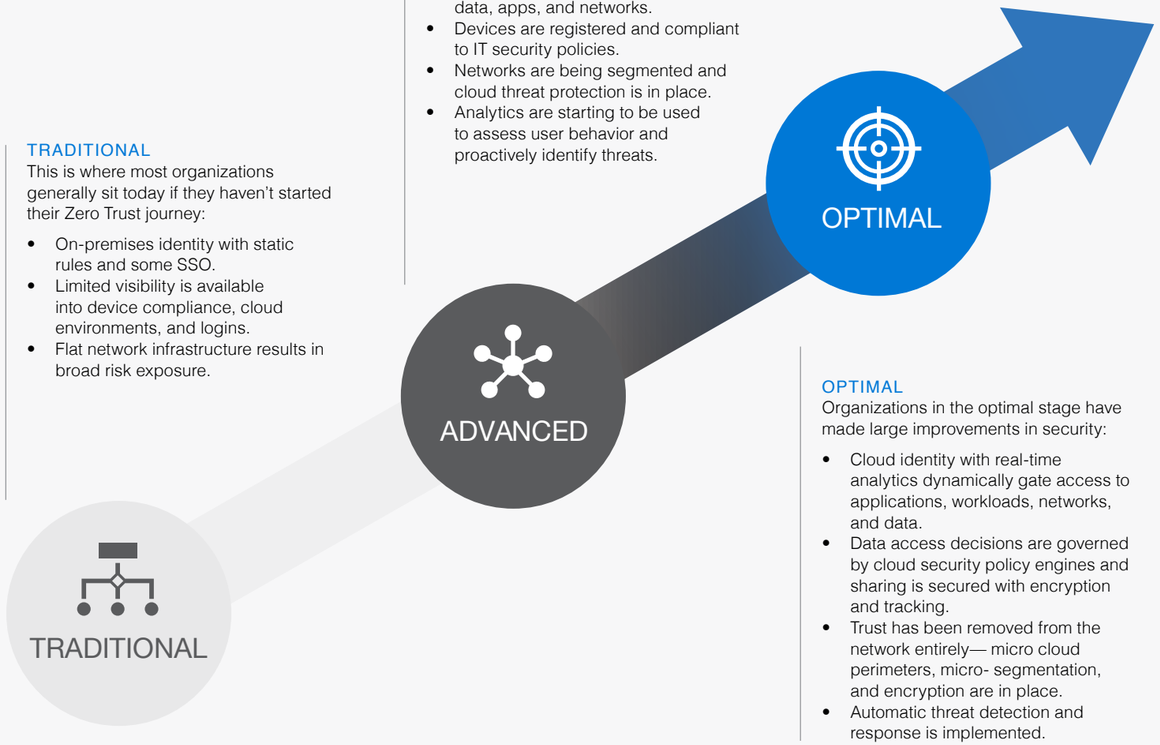

OPTIMAL
Organizations in the optimal stage have made large improvements in security:

- Cloud identity with real-time analytics dynamically gate access to applications, workloads, networks, and data.
- Data access decisions are governed by cloud security policy engines and sharing is secured with encryption and tracking.
- Trust has been removed from the network entirely— micro cloud perimeters, micro- segmentation, and encryption are in place.
- Automatic threat detection and response is implemented.

Not every Zero Trust model implementation is the same

Different agency requirements, existing technology implementations, and security stages all affect how a Zero Trust security model implementation is planned. Using our experience in helping government customers to secure their IT operations as well as implementing our own Zero Trust model, we've developed the following maturity model to help you assess your Zero Trust readiness and build a plan to get to Zero Trust.

We've also developed an expanded maturity model to help you assess your own Zero Trust readiness across your user identities, devices, application, data, infrastructure, and networks.

	TRADITIONAL	ADVANCED	OPTIMAL
Identities	On-premises identity provider is in use No SSO is present between cloud and on-premises apps Visibility into identity risk is very limited	Cloud identity federates with on-premises system Conditional access policies gate access and provide remediation actions Analytics improve visibility	Passwordless authentication is enabled User, device, location, and behavior is analyzed in real time to determine risk and deliver ongoing protection
Devices	Devices are domain joined and managed with solutions like Group Policy Object or Config Manager Devices are required to be on network to access data	Devices are registered with cloud identity provider Access only granted to cloud managed & compliant devices DLP policies are enforced for BYO and corporate devices	Endpoint threat detection is used to monitor device risk Access control is gated on device risk for both corporate and BYO devices
Apps	On-premises apps are accessed through physical networks or VPN Some critical cloud apps are accessible to users	On-premises apps are internet-facing and cloud apps are configured with SSO Cloud Shadow IT risk is assessed; critical apps are monitored and controlled	All apps are available using least privilege access with continuous verification Dynamic control is in place for all apps with in-session monitoring and response
Infrastructure	Permissions are managed manually across environments Configuration management of VMs and servers on which workloads are running	Workloads are monitored and alerted for abnormal behavior Every workload is assigned app identity Human access to resources requires Just-In-Time	Unauthorized deployments are blocked and alert is triggered Granular visibility and access control are available across all workloads User and resource access is segmented for each workload
Networks	Few network security perimeters and flat open network Minimal threat protection and static traffic filtering Internal traffic is not encrypted	Many ingress/egress cloud micro-perimeters with some micro-segmentation Cloud native filtering and protection for known threats User to app internal traffic is encrypted	Fully distributed ingress/egress cloud micro-perimeters and deeper micro-segmentation ML-based threat protection and filtering with context-based signals All traffic is encrypted
Data	Access is governed by perimeter control, not data sensitivity Sensitivity labels are applied manually, with inconsistent data classification	Data is classified and labeled via regex/ keyword methods Access decisions are governed by encryption	Classification is augmented by smart machine learning models Access decisions are governed by a cloud security policy engine DLP policies secure sharing with encryption and tracking

Tools to drive your Zero Trust implementation

As you begin to assess your Zero Trust readiness and begin to plan on the changes to improve protection across **identities, devices, applications, data, infrastructure,** and **networks**, consider these key investments to help drive your Zero Trust implementation more effectively. Through our own experience, we've found each of the following to be critical to closing important capability and resources gaps:

1. **Strong authentication.**
   Ensure strong multi-factor authentication and session risk detection as the backbone of your access strategy to minimize the risk of identity compromise.
2. **Policy-based adaptive access.**
   Define acceptable access policies for your resources and enforce them with a consistent security policy engine that provides both governance and insight into variances.
3. **Micro-segmentation.**
   Move beyond simple centralized network-based perimeter to comprehensive and distributed segmentation using software-defined micro-perimeters.
4. **Automation.**
   Invest in automated alerting and remediation to reduce your mean time to respond (MTTR) to attacks.
5. **Intelligence and AI.**
   Utilize cloud intelligence and all available signals to detect and respond to access anomalies in real time.
6. **Data classification and protection.**
   Discover, classify, protect, and monitor sensitive data to minimize exposure from malicious or accidental exfiltration.

Next Steps

While a Zero Trust security model is most effective when integrated across the entire digital estate, most organizations will need to take a phased approach that targets specific areas for change based on their Zero Trust maturity, available resources, and priorities. It will be important to consider each investment carefully and align them with current business needs. The first step of your journey does not have to be a large lift and shift to cloud-based security tools. Many organizations will benefit greatly from utilizing hybrid infrastructure that helps you use your existing investments and begin to realize the value of Zero Trust initiatives more quickly.

Fortunately, each step forward will make a difference in reducing risk and returning trust in the entirety of your digital estate.

Microsoft is currently on its own Zero Trust journey. Head over to our *IT Showcase* to learn more about how we've approached our Zero Trust journey, our current progress, and upcoming milestones.

For a practical approach on ways to improve your Zero Trust posture, the following pages feature a preview of our blog series discussing practical steps to implementing Zero Trust security using Azure solutions.

To see where you are in your journey, take the *Zero Trust assessment* and find out more on ways *Microsoft Azure Government* can help advance your agency's Zero Trust maturity.

A PRACTICAL APPROACH TO IMPLEMENTING ZERO TRUST IN AZURE

Microsoft experts offer six approaches to building Zero Trust using Azure solutions.

Experts from Microsoft Azure have put together a series of blogs discussing practical steps to implementing Zero Trust security using Azure solutions. Here's a preview of how to get started improving identity and access management; protecting cloud workloads; investigating insider attacks; and monitoring cloud security, policy enforcement and supply chain risk management.

Read the entire series from *TJ Banasik*, CISSP-ISSEP, ISSAP, ISSMP, Sr. Program Manager, Microsoft Azure Government End-to-End Customer Engineering; *Mark McIntyre*, Senior Director, Cybersecurity Solutions Group; and *Adam Dimopoulus*, Senior Program Manager, Azure Global Customer Engineering.

IMPLEMENTING ZERO TRUST WITH STRENGTHENED IDENTITY AND CONDITIONAL ACCESS MANAGEMENT

In cloud-based architecture, identity provides the basis of a large percentage of security assurances. While legacy IT infrastructure has traditionally relied on firewalls and network security solutions for protection against outside threats, these controls are less effective in cloud architectures with shared services being accessed across cloud provider networks or the internet.

It's challenging or impossible to write concise firewall rules when you don't control the networks where these services are hosted, and when different cloud resources spin up and down dynamically.

Those challenges make managing access, based on identity authentication and authorization controls in the cloud, essential to ensuring that data and resources are protected by limiting access only to those who need it.

Microsoft has several offerings, including *Azure Active Directory*, to manage identity and access to any cloud and on-premises application, with advanced protection that automates, detects and remediates identity-based risks. Azure AD's Conditional Access capabilities are the policy decision point for access to resources based on user identity, environment, device health, and risk — verified explicitly at the point of reference.

Microsoft Azure Government has developed a 12-step process for securing identity and access management in federal information systems that is aligned with the identity management principles within NIST, OMB and CISA Zero Trust frameworks, including:

12 STEPS TO IMPLEMENTING ZERO TRUST IDENTITY MANAGEMENT PRINCIPLES IN AZURE

12. **Perform Continuous Training**
11. **Leverage Adaptive Access Control**
10. **Secure Administrative Rights**
9. **Enforce Least Privilege**
8. **Make Trust Determinations**
7. **Determine User Access to Resources**
6. **Evaluate Credentials and Authentication**
5. **Implement Secure Authentication**
4. **Control User Authentication**
3. **Conduct User Account Provisioning**
2. **Manage Identity & Access**
1. **Employ an Identity Management System**

Read the full blog.

ZERO TRUST SECURITY AND PROTECTING CLOUD WORKLOADS

The Microsoft *Zero Trust vision paper* outlines three principles of Zero Trust implementation in cloud workloads: Verify explicitly, provide least-privilege access and assume breach. These principles also assume the need to continuously measure trust and risk.

While security architecture design is a key starting point for Zero Trust, this model must be continually monitored throughout the enterprise security lifecycle. The Azure Security Center serves as both a cloud protection platform and Cloud Security Posture Management solution, facilitating continuous monitoring of security controls in dynamic environments.

Microsoft Azure Government has developed a nine-step process for helping protect cloud workloads in federal information systems using *Azure Security Center*. The framework for the security solution is aligned with the security protection principles within the NIST, OMB and CISA Zero Trust frameworks:

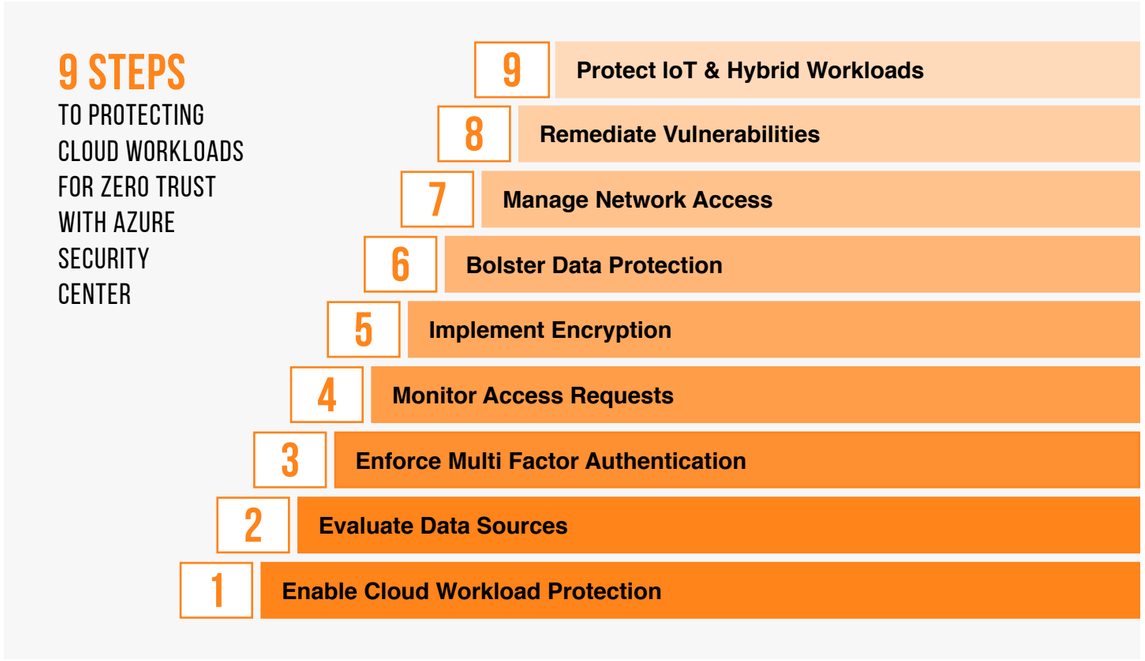

Read the full blog.

MONITORING CLOUD SECURITY AND THE ZERO TRUST SECURITY FRAMEWORK

Microsoft Azure Government has developed an 11-step process for monitoring cloud security in federal information systems that aligns with NIST, OMB and CISA Zero Trust frameworks. And we've created a range of security monitoring solutions including:

- *Azure Sentinel* provides a scalable, cloud-native, security information event management (SIEM) and security orchestration automated response (SOAR) solution.
- *Azure DDoS Protection* protects against distributed denial of service (DDoS) attacks.
- *Azure Firewall* is a cloud-based security service that protects Azure Virtual Network resources.
- *Azure Web Application Firewall* protects your applications from common web vulnerabilities such as SQL injection and cross-site scripting, and lets you customize rules to reduce false positives.
- *Network Security Group* contains security rules that allow or deny inbound network traffic to, or outbound network traffic from, several types of Azure resources.

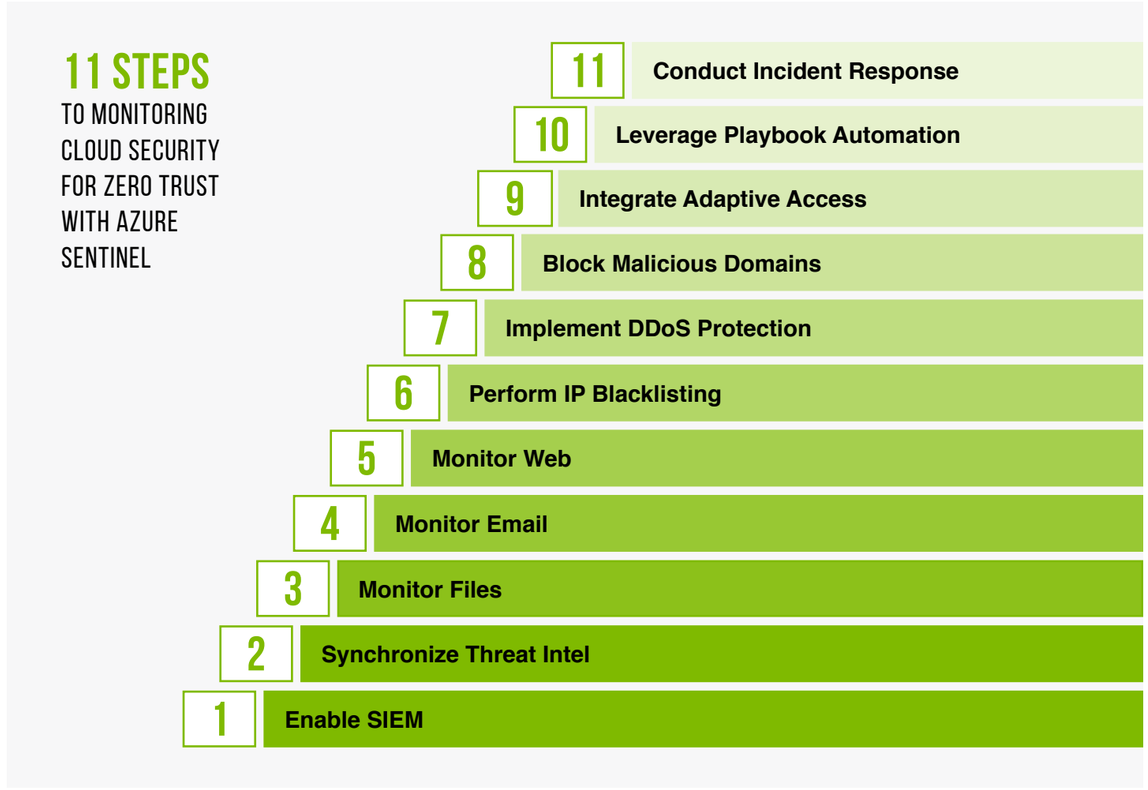

Read the full blog.

ZERO TRUST SECURITY AND POLICY ENFORCEMENT

The Zero Trust model requires strict access control to ensure every access request is authenticated, inspected and authorized within policy constraints. Therefore, it's important to evaluate the 5 W's of each access request to determine: who is requesting access, what resource is requested, when the environment is being accessed, where is the user located, why is access requested — and how is the user authenticating.

A central element to evaluating that access is the ability to enforce policies. A policy enforcement point solution evaluates each access request and applies respective control requirements such as multi-factor authentication for an access request from an unexpected location.

In Azure Government, we cross-referenced TIC 3.0, NIST Cybersecurity Framework and NIST SP 800-53 to align requirements for implementing Zero Trust models to create a 10-step process for enforcing access control policy in federal information systems. The process can take advantage of several offerings for policy enforcement:

- *Azure Policy* helps you manage and prevent IT issues with policy definitions that enforce rules and effects for your resources.
- *Azure Security Center* provides unified security management and advanced threat protection across hybrid cloud workloads along with a compliance dashboard.
- *Azure Blueprints* simplifies deployments by packaging artifacts, such as Resource Manager templates, role-based access controls and policies into a single blueprint definition.

Read the full blog.

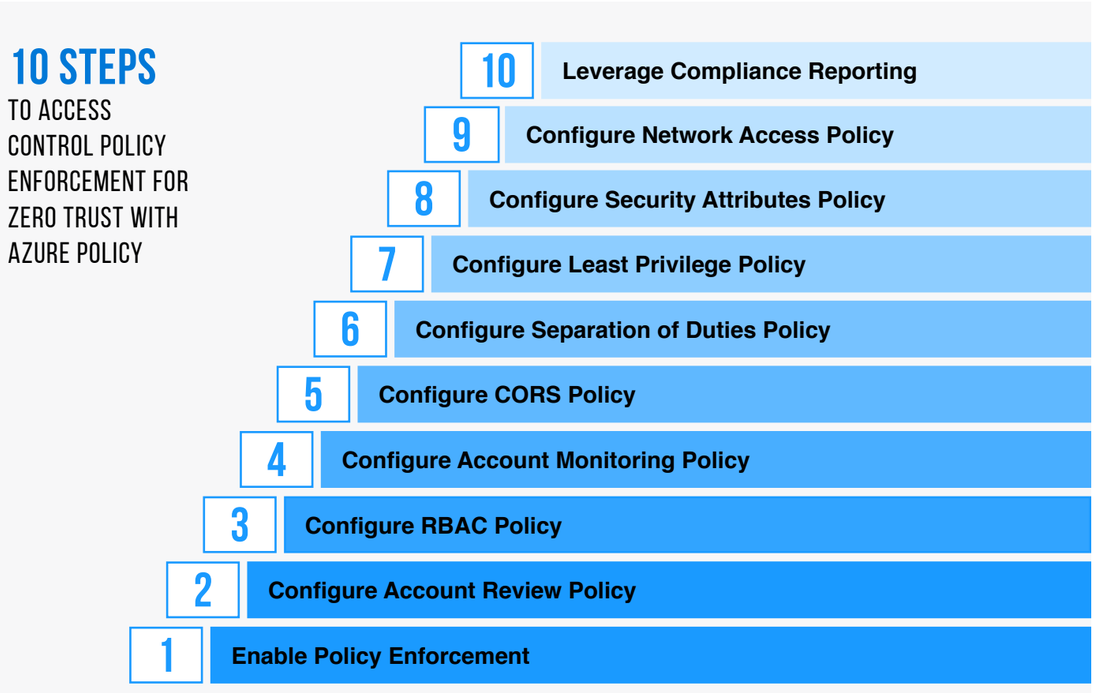

Azure blueprint for Zero Trust lets you easily create, deploy, and update compliant environments.

Simplify large scale Azure deployments by packaging key environment artifacts — such as Azure Resource Manager templates and role-based access controls and policies — in a single blueprint definition. Easily apply the blueprint to new subscriptions and environments and fine-tune control and management through versioning. Azure blueprint for Zero Trust lets you:

- **Streamline environment creation** - Easily create and manage your cloud governance templates, access controls, and policies as a single compliant package so environments are ready to be configured.
- **Enable compliant development** - Speed deployment of compliant applications to production through a self-service model, and easily deploy compliant environments matched to production standards.
- **Lock foundational resources** - Avoid unwanted changes and misconfigurations—even by subscription owners—that could affect multiple applications.

*Find out more about **Azure blueprint for Zero Trust.***

ZERO TRUST SECURITY TO IMPROVE INSIDER THREAT MONITORING

Because Zero Trust assumes breach, tools will verify each request as though it originates from an uncontrolled network and will look for anomalies prior to granting access — and then limit access to the least privilege required to perform respective job functions.

Solutions that use threat detection algorithms can detect access patterns that are out of normal behavior and deny the compromised account (or insider threat) access to resources. In order for these solutions to work properly, defining governance — that incorporate TIC 3.0 Zero Trust principles and respective NIST 800-207 and NIST standards — is critical. Microsoft Azure Government has developed an 8-step process and several solutions to facilitate insider threat monitoring for federal information systems:

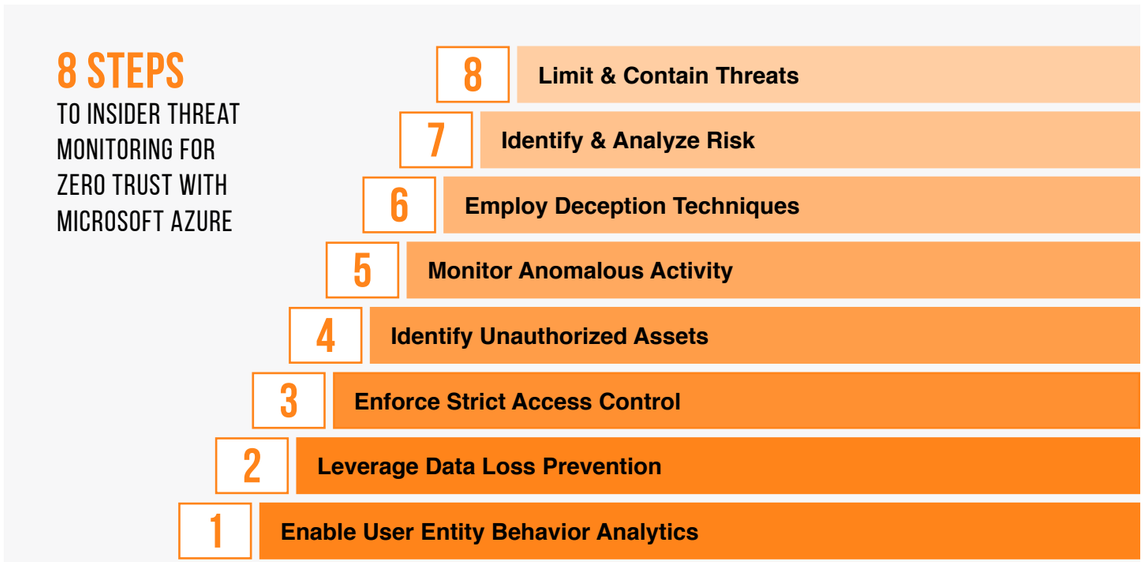

Read the full blog.

SUPPLY CHAIN RISK MANAGEMENT AND ZERO TRUST

Supply chain risk management requires evaluating service providers on a holistic basis by taking into consideration factors such as vendor security controls, enterprise switching costs and supply chain risk management — and assuming breaches will occur.

Threat actors target supply chain partners to gain a foothold inside an enterprise network and initiate various attacks, such as:

- Compromising software building tools to imprint malware into all software generated from the building tools.
- Replacing software update repositories with malicious replicas that distribute malware across entire software ecosystems.
- Stealing code-signing certificates to make malicious software appear as legitimate code.
- Intercepting hardware shipments to inject malicious code into hardware, firmware and field programmable gate arrays (FPGAs).
- Pre-installing malware onto IoT devices before they arrive to target organizations.

This nine-step process is a starting point to facilitate supply chain risk management, which also requires alignment of people, processes, policy and technology:

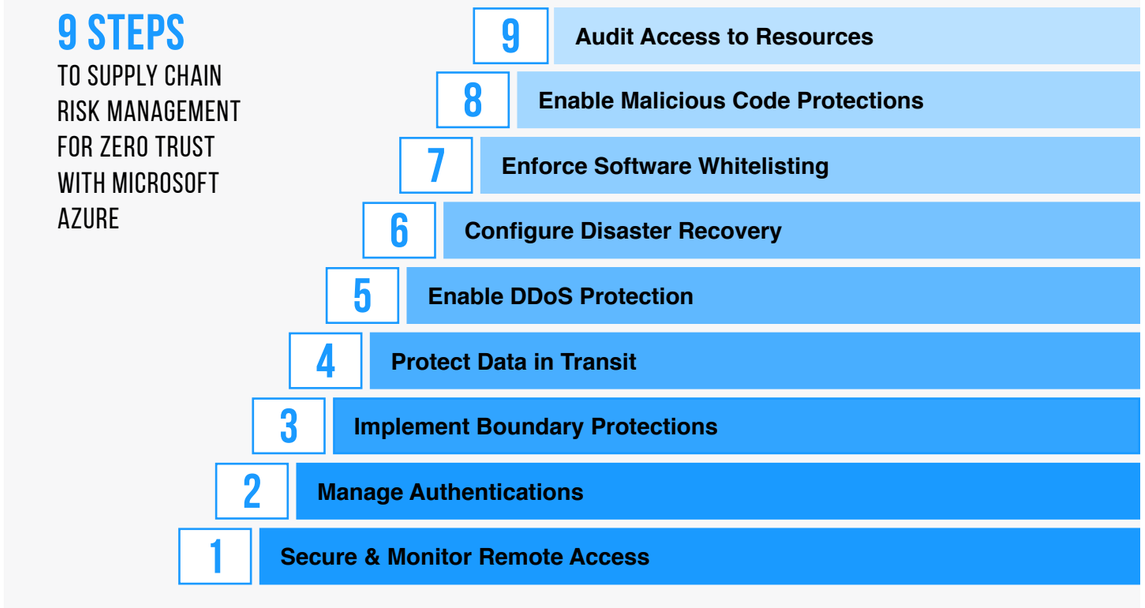

Azure offers a variety of solutions to facilitate supply chain risk management.

Read the full blog.

fedscoop

Microsoft Azure