

From MFA to Zero Trust:

A Five-Phase Journey to
Securing the Federal Workforce



Table of Contents

Introduction	03
The Zero Trust Approach	04
Introducing the Three Pillars of Zero Trust	05
Using this Guide to Implement Zero Trust for the Federal Workforce	05
Phase 1: Establish User Trust	06
Phase 2: Device and Activity Visibility	09
Phase 3: Device Trust	12
Phase 4: Adaptive Policies	15
Phase 5: Zero Trust for The Federal Workforce	18
Summary	20

Introduction

A zero trust security architecture helps solve the many complex and unique challenges that federal agencies and their workforces face. The principles of zero trust can be applied to secure both cloud and legacy applications and data, employees and contractors as well as personal or agency-managed devices. According to NIST SP 800-207, “Zero trust assumes there is no implicit trust granted to assets or user accounts based solely on their physical or network location (i.e., local area networks versus the internet) or based on asset ownership (enterprise or personally owned). Authentication and authorization (both subject and device) are discrete functions performed before a session to an enterprise resource is established.”

Designing a zero trust architecture for your organization and putting it in place can feel daunting.

This guide lays out a practical approach in five phases for implementing Zero Trust for the Federal Workforce, which comprises an agency’s users and their devices, and how they access applications. The approach is iterative. Begin with a specific set of people, expand coverage for their applications and then for their devices. Once we are always verifying trust within this well-defined scope, apply a set of reasonable policies to enforce trust and protect the organization. Finally, integrate this scope with the broader organization’s IT and security functions and shift to continuous improvement.

Following these steps, an organization can incrementally achieve a zero trust transformation.

This guide lays out a practical approach in five phases for implementing *Zero Trust for the Federal Workforce*, which comprises an agency’s users and their devices and how they access applications.



The Zero Trust Approach

The zero trust principles share much in common with security fundamentals. Like default deny, zero trust begins with no access until trust is demonstrated and established. As with least privilege, zero trust relies on just enough trust and seeks to minimize excessive trust. Zero trust builds upon these fundamentals with following concepts:

Visibility informs policy. Provide as much intelligence and insight as possible to the people administering the technology, in order to produce informed policies.

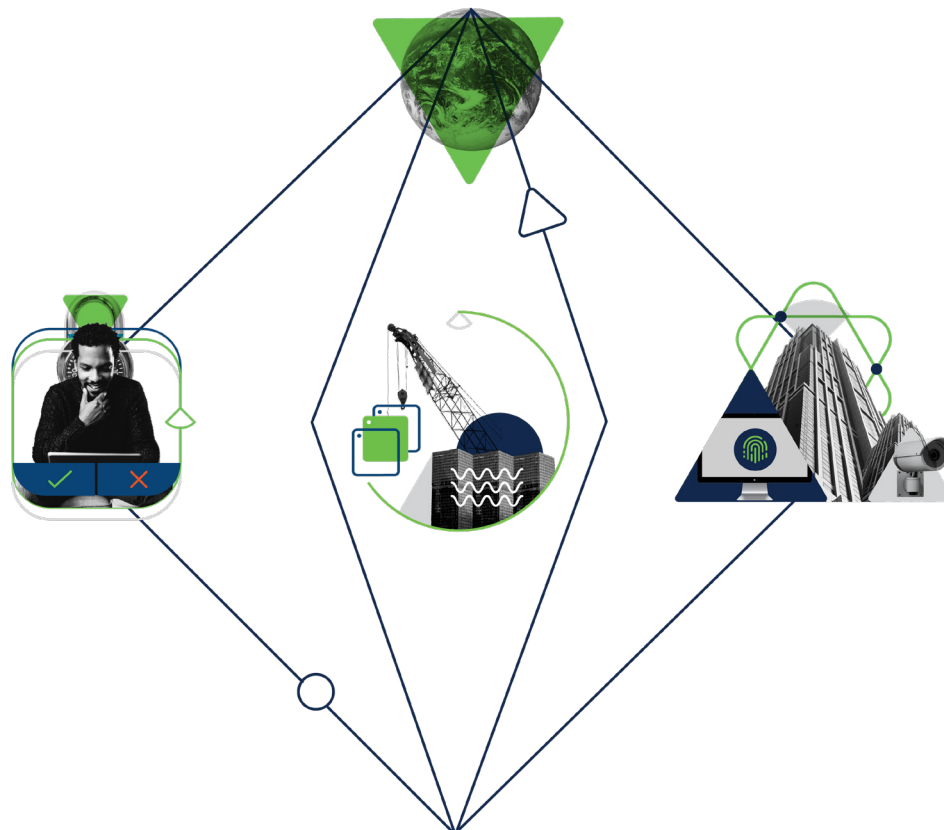
Trust is neither binary nor permanent. Continually reassess the posture of users, devices and applications and adjust your trust accordingly. Be prepared to respond to events that raise the risk level by containing newly discovered threats and vulnerabilities.

Ownership is not a control. Validate and extend trust to devices, applications and networks that you don't own or manage, from BYOD (bring your own device) and IoT (Internet of Things) devices to SaaS and public cloud.

The perimeter is any place where you make an access control decision. Choose the layers and process points that work for your environment, whether it's at the network layer, the application layer, at the point of identity verification or during a transaction workflow.

Access decisions are based on re-establishing trust every time. Membership within a group, an application service within a tier or a device connected to a network location, are not enough on their own to authorize activity.

Containment. Combine least privilege and segmentation with response capabilities to monitor for threat activity and limit its spread by default.



Introducing the Three Pillars of Zero Trust

Zero Trust for the Workforce: People such as employees, contractors, partners and vendors accessing work applications using their personal or agency-managed devices. This pillar ensures only the right users and secure devices can access applications.


Zero Trust for Workloads: Applications running in the cloud, in datacenters and other virtualized environments that interact with one another. This pillar focuses on secure access when an API, a microservice or a container is accessing a database within an application.

Zero Trust for the Workplace: This pillar focuses on secure access for any and all equipment connecting to enterprise networks; such as user endpoints, physical and virtual servers, printers, cameras, HVAC systems, kiosks, infusion pumps, industrial control systems and more.

WORKFORCE



WORKLOADS



WORKPLACE



Using this Guide to Zero Trust for the Workforce

This guide recommends an iterative approach for the journey to Zero Trust for the Workforce. Tightly scope one aspect of the organization, proceed with that scope through the five phases of the journey and then integrate that scope into the organization's zero trust architecture. The approach means each initiative is a self-contained project within the larger transformation. To use this guide, within the scope of each initiative, use the following sections for each phase of the journey.

Description and Objectives. For each of the five phases in the journey to zero trust, an overview is provided along with the objectives we must meet to complete this phase. These objectives are scoped to the zero trust initiative, not to the overall organization. For example, establishing user trust and device trust is specific to the people and their devices within the portion of the organization we are moving to the zero trust architecture.

Transformation. The beginning of each phase should include a workshop to gain consensus, support and identify next steps. Suggested attendees are stakeholders from the security function, the IT operations and support function and the business units within the initiative's scope. Questions are provided along the strategic, management and operational disciplines. By scoring these with a 1 through 5 assessment, the team can determine the organization's maturity for the given phase.

Components and Challenges. Successful transformations involve integrating technology while managing potential pitfalls. The components section includes recommended technologies for the given phase of the Zero Trust for Workforce initiative. Under challenges, this paper provides frequently seen concerns and potential solutions.

Metrics. Metrics should be implemented for guiding action and tracking success along the transformation. In this section, metrics are suggested for risk management, security, IT support and IT operations. Each specific scoped initiative can use these metrics to progress through the phases. Once the scoped initiative is completed, the metrics can continue to be collected to measure the efficacy of the overall zero trust architecture.

Phase 1: Establish User Trust

Ensure you have the right mechanisms and processes to ensure only authorized users are attempting to access your resources. This can be achieved in a number of ways. Multi-factor authentication (MFA) is a commonly used technology. An emerging option is passwordless: a single strong authentication factor.

Description

The perimeter is any place where you make an access control decision. With Zero Trust for the Workforce, the perimeter gets established when a person accesses an application. It is a seemingly simple moment. A click, a fleeting pause, and the application opens. But within that moment, many tasks and verifications are being performed. Many more are possible. For it is within the authentication workflow we have a control point to evaluate and enforce trust. The first step in establishing a zero trust architecture is gaining control over this identity verification.

Privileges are what people can do with the IT we protect. Trust is whether we have confidence that people will use those privileges accordingly. Much like the principle of least privilege emphasizes we only assign the privileges needed and no more, the principle of zero trust emphasizes we only trust people as much as needed and no more.

We need to take stock of the areas we trust people and the means in which people can establish or lose trust. We can approach this by type of users; for example, those with privileged access or those with remote access. Alternatively, we can approach this by activity; for example, a critical process or a frequently targeted process. Scope the zero trust effort to one area and assess the people and technology involved.

After scoping the user trust effort, three work streams are required: the first workstream is to increase security; the second is to manage usability; and the third is to help people through the change by communicating and socializing what's to come. For usability, identity verification must be quick, convenient and nearly invisible under most circumstances. But for security, when the circumstances are suspicious, the identification verification offers a control point to prevent access. Balancing the two is a success factor in establishing user trust.

Ensuring the security system can trust the user is who they claim to be is the first step in establishing a zero trust architecture. User authentication forms the foundation. In addition to username and password, two-factor authentication (2FA) or multi-factor authentication (MFA) provides a stronger assurance that the user is who the user claims to be. When trust is in question, such as because of user behavior or when additional trust is required, such as when an action puts the organization at risk, the security system must be able to challenge the user to re-authenticate or to produce additional forms of authentication. This establishes and maintains strong identity across a variety of situations in a zero trust architecture.



Ensuring the security system can trust the user is who they claim to be is the first step in establishing a zero trust architecture.



Objectives

- Risk-rank people by role and access to information
- Go beyond employees to include contractors, temps and other hard to secure groups
- Expand the use of MFA
- Evaluate the use of passwordless for strong authentication
- Expand the number of people using the zero trust architecture
- Ensure every user is continuously validated

Transformation

Strategic

The objective here is to determine whether there is an established direction for the organization and may lead to a broad discussion.

Is there a clear identity strategy in your organization?

- Determine whether there is a clear vision and direction, and who owns this strategy and how is it governed within the organization
- Publish the strategy so that it is clear to the organization
- Determine the extent and coverage of the strategy and whether it covers third parties, partners and the use of external resources such as cloud providers
- Is the strategy effective and are clear measurements provided?
- By which methods are change and improvements measured, identified and implemented?

Management

The objective is to define how well the identity and access management (IAM) requirements for the organization are understood and managed on a daily basis.

Do you have a clearly defined IAM function?

- Is there a named group which manages the IAM requirements for the organization?
- What is the level of hygiene of the IAM systems being used?
- How does it integrate with other parts of the business to ensure full joiner, mover and leaver consistency?
- How clear is the reporting from the IAM function?
- Determine the speed and consistency of changes implemented and whether they are in line with business change

The objective is to identify the progress being made in the implementation of an MFA program, which will provide an assurance of identity that is segregated from the IAM function.

Has an MFA solution been implemented?

- Has a solution has been identified and implemented?
- To which extent has the MFA solution been integrated with the IAM?
- Is there feedback and reporting on the implementation?

Operational

The objective is to identify the extent to which any solution has been implemented across the user base.

What percentage of users enrolled are using MFA?

- Has a solution been implemented?
- Is there a clear metric which shows the number of users enrolled?
- Is there a dependable metric on the user base?
- Is there a metric that breaks down the implementation by user, group or business area?
- Plan to ensure a full deployment is in place, measured and monitored

Is MFA integrated into other functions?

- Has a solution been implemented?
- Identify potential areas of security and IT in particular that can benefit from the MFA capability
- What are the outputs from an MFA solution?
- Which method and manner are outputs shared with other functions, automated or otherwise?
- Create a benefits statement showing how the MFA solution improves delivery in other functions

Components

Identity Database. To keep the information and attributes about your users, and to group them where necessary according to organizational, geographical or other aspects. This is generally provided by the **Identity Provider (IdP)** service.

Strong Authentication. 2FA, or MFA, is vital to the strategy of making it harder to compromise an account. In addition, making a request to complete an MFA challenge is a means of re-asserting trust.

Challenges

Our end users are very hesitant to change.

Past security changes have made it harder for end users to get work done, making them resistant to future changes. The journey to zero trust must minimize the impact on people, while communicating and socializing the need for change.

We're focused on securing our remote access.

Remote access is a logical starting point for a zero trust initiative. With more apps going to the cloud, the definition of remote is expanding and the perimeter is blurring. Build momentum with remote access and continue on to other workflows.

We can't interrupt workflows across the organization.

Nothing will stop a security initiative faster than halting the organization's work. Take the journey to zero trust one workflow at a time with careful planning, testing and an emphasis on usability.

Metrics

Risk

- Overall risk level
- Risk register – issue mitigation
- Audit compliance – issue mitigation

Security

- Incident reduction
- Account takeover (ATO)
- Business email compromise (BEC)

Support

- FTE usage
- User MFA support tickets
- User MFA support satisfaction (NPS)

Operations

- Apps configured with MFA
- Users enrolled in MFA

Phase 2: Device and Activity Visibility

Which endpoint or device is being used with every access request? What is its current security state, and from where is the request originating? This is a key stage for detecting account takeover attempts and other risks.

Description

The Zero Trust for Workforce implementation rests on the combination of validated users and endpoint devices. In the first phase, the emphasis is on expanding the users who are within the scope of a zero trust architecture and deepening the means of validating trust in those users. In the second phase, the emphasis shifts to expanding the devices and applications within the scope. This completes the visibility into the workflow of a user on a device accessing an application. The workflow sets the foundation for what is to come.

Access decisions are based on re-establishing trust every time people perform an activity. An activity is a set of applications used by a specific set of users completed in support of an organization’s function. Our focus remains on the scope previously defined. We begin with inventorying, risk-ranking and prioritizing the applications utilized by the target user population. To move an application into the zero trust architecture, three conditions must be met: integrate with strong authentication (typically with Radius, SAML or AD FS); ensure anywhere access to the applications (typically with a network gateway or VPN); and unify how users launch the application (typically with a single sign-on portal). As organizations often have hundreds or thousands of applications, and as activities change over time, the process of identifying and integrating applications with the zero trust architecture is ongoing.

Activities are completed by validated users on devices. Under zero trust, devices fall into two categories: authentication devices and access devices. Authentication devices are those that an end user uses to perform strong authentication to establish and maintain trust. The access device is where the end user launches and interacts with the application. Users often use a number of devices to complete work, including organizationally owned and personally owned devices. Criminals may exploit security vulnerabilities to hold the devices for ransom, gain access to sensitive data or otherwise disrupt the organization. The zero trust security system can evaluate trust in the device using several measures including:

- **Have we seen this device before?**
- **Is the device’s operating system and software up to date?**
- **Are there any indications of tampering?**

In this phase, the objective is to gain visibility into devices and associated trust factors.

Ensuring the security system can trust the devices authenticating and accessing applications is the second step in establishing a zero trust architecture. We begin by inventorying and prioritizing applications and devices. The security goal is to establish a perimeter, and the perimeter is where the security system makes an access control decision about a user for a given application. Therefore, one ongoing workstream is integrating the applications with the zero trust authentication and access controls. The perimeter is only as good as the hardware it is running on. The other ongoing workstream is increasing visibility into the access and authentication devices and their accompanying security posture. Together, this builds the device and activity portfolio which we can use later to make policy and enforcement decisions in the zero trust architecture.



Objectives

- Risk-rank and prioritize applications
- Gain visibility into the devices and applications in use
- Go beyond on-premises applications to include securing cloud apps
- Expand the applications in the zero trust architecture
- Ensure continuous validation of access to every application

Transformation

Strategic

To identify if there is a clear strategy to provide device visibility on managed or unmanaged devices.

Is there a clear delineation between managed and unmanaged devices?

- Identify whether there is a policy or standpoint on the use of non-agency devices to access resources
- Ascertain whether the resources and assets are assigned different levels of risk so distinct access policies may be developed

Management

The objective is to define how well the device inventory and device status are understood and managed operationally.

Is access to applications fully controlled for users?

- Ascertain whether there is a clear inventory of applications
- Have applications been assessed for potential risk if compromised?
- Is there a clear distinction as to the extent of internal and external users?
- Are policies in place to control the external access to applications and clearly restrict the use of any other assets?
- Is there a governance process to manage the policy update process?

Is the device inventory up to date and regularly reviewed?

- Determine who owns the asset inventory and is responsible for its maintenance
- What are the methods of updating the inventory?
- Identify whether the use of device identification at the point of access is implemented
- Is device identification at the point of access linked to the asset inventory in place?
- Is there constant feedback and review between the inventory status and the access control?

Has a device hygiene check been implemented?

- Determine whether a solution to check the hygiene of a device is in place and in operation
- To which extent has solution been implemented within the organization?
- Has the solution been integrated with the other security and IT function?
- Is progress clearly measured and reported?
- Is there a governance process to manage the solution?

Operational

The objective is to identify the extent to which any solution has been implemented across the user base and how effectively it contributes to the rest of the security function.

What percentage of applications are managed or unmanaged?

- Determine whether there is a clear inventory of cloud-based and network-based applications
- Determine whether there is an integration with any security solutions and applications
- Ascertain whether the end user will see all applications in a similar fashion
- Ascertain whether there is measurement of user expectations and approval rating
- Identify whether improvements and exceptions are reported to a clear owner

What percentage of devices are managed or unmanaged?

- Determine if there is a clearly defined process for identifying managed or unmanaged devices
- Determine whether there is a centralized inventory for devices
- Is there a process for maintaining the inventory on a periodic basis?
- Ascertain the reporting on devices
- Are reports audited or otherwise validated?

Is device visibility integrated into other functions?

- Determine whether any solution providing visibility over devices has been implemented
- Identify potential areas of security and IT in particular that can benefit from device visibility
- What are the outputs from such a solution?
- In what method and manner are outputs shared with other functions, automated or otherwise?
- Create a benefits statement showing how enhanced device visibility improves delivery in other functions

Components

Application Inventory Database. A listing of applications and activities maintained in an up-to-date repository, including type, purpose, risk ranking and application owner.

Access Proxy. This service carries out the connections and policy enforcement. More complex Access Proxies can provide traffic load balancing, Transport Layer Security (TLS), authentication, access control list (ACL) evaluation, authorization and self-service for users.

Single Sign-On (SSO). Make it much easier on your users by providing one portal for access to all of their applications and systems.

Challenges

We lack the resources to cover all applications and devices. Many teams are stretched thin and increasing the apps or devices could increase the administrative burden. Fortunately, it's not all or nothing. Select a zero trust approach that scales and is easy to operate. Steadily increase coverage in a prioritized and systematic way.

Only these apps require 2FA for compliance. Mission accomplished. When the initial business case for zero trust includes compliance, it is possible for the initiative to stall once those requirements are met. Use each new application workflow to re-establish the organization's commitment to zero trust.

Metrics

Risk:

App risk reduction – core apps

Security:

Unmanaged apps

Unmanaged devices

Support:

App integration MFA support tickets

App integration MFA support satisfaction (NPS)

Operations:

Core apps configured with MFA

Core apps configured with SSO

Auth devices – known devices

Phase 3: Device Trust

Whether it's agency-owned or not, and whether it's managed or unmanaged, an organization can mark as trusted the devices it has registered and expects to see associated with that particular user.

Description

In the classic environments, before zero trust, device trust tended to be based on location. Because the device was on the network, we assumed it was supposed to be there, and it got access to anything it asked for. Both of these “tests” failed for a number of reasons, from stolen passwords to spoofed network addresses to compromised endpoints. In a zero trust architecture, the path to trust needs more checkpoints, such as authentication factors and conditions placed on the device. One of these conditions can be whether it's a managed, agency-owned endpoint.



A managed endpoint is presumably owned by the organization, or at least known. Device visibility in the earlier phase will illuminate the size of the endpoint population. Managed devices may be tracked as part of an inventory, enrolled in a configuration and patch management program and monitored for security events. For this reason, we may choose to trust it more than we would trust an unmanaged, personal device. Many organizations have the policy that only the endpoints they own and assign to staff can be used to access agency data. However, this policy can be difficult to enforce, especially if there's no way to check.

There are different ways to bring a device under management. If the endpoint has the VPN client installed, it's assumed to be an approved and managed asset, so whoever is using it will be allowed to access the internal network from the outside (say, at home, or from a hotel or coffee shop). With common port-based network access control (NAC), if the endpoint has an 802.1x certificate installed, it's assumed to be an approved and managed asset, so whoever is using it will be allowed to connect to the internal network from inside the building. Finally, enrolling devices into a mobile device management (MDM) system allows us to enforce configuration policies by installing an agent.

In each of these cases, we've marked the endpoint as trusted by installing something on it (or given it a second factor, “something it has”). If we can't manage an endpoint, it's generally more difficult to convince that endpoint owner to let us install something on it. A certificate or other method of fingerprinting is lightweight, and may be more acceptable than installing running software. Still, the key requirement is to make that marking unforgeable and prevent it from being copied to another device. The important point is that we've seen the device before and expect to grant it access, as opposed to endpoints that are trying to access our applications that we've never seen before that may be used by attackers.

Since we will be making trust decisions based on the marking's presence or absence, it functions as yet another authentication factor and it needs protection in the same way we must protect the primary user credentials (username and password) and the second factor (such as a one-time password, U2F device, or push-based authentication). Certificates offer a way to identify the device as managed. We can take it a step further by including device and user data in the certificate, tying them together so neither one's credentials can be leveraged alone.

The net result is a spectrum of device management options, each providing telemetry to establish trust. We can have a fully managed device provided by the organization, with a full set of agents and visibility into every app and action. We can have a traditional BYOD model, with a personally owned device running a single agent, such as MDM, with visibility and control over well-defined areas. Or we can have a light-weight control such as a cookie or certificate, with visibility over what's available from the web browser user agent or user hints at the time of authentication. Managing devices as a portfolio, with specific levels trusted to complete specific activities, simplifies the management and support costs.

In this phase, the objective is to establish controls over devices and the trust factors available given the device management telemetry.

From a defense perspective, consider the situation when a user loses credentials to an attacker. With device trust, the attacker still needs to use a valid endpoint belonging to that user to get into an application. It's not enough to have the username and password with a different agency device. Trusting the devices only if they're with the right user is the next step towards tighter security that the zero trust architecture makes possible.

Objectives

- Risk-rank and prioritize devices
- Go beyond organizationally-owned devices to include BYOD
- Establish a portfolio of technologies for trusting devices, including MDM and VPN
- Expand the trusted end-user devices in the zero trust architecture
- Ensure every access and authentication device is continuously validated

Transformation

Strategic

To identify if there is a clear strategy on policy driven assessments of devices at the point of access.

Is there an existing strategy on trusted devices?

- Identify whether there is a strategy or program for the assessment and control of trusted devices
- Establish an approach for agency-owned and personally-owned device management
- Ascertain whether risk-based policies are applied to identify the level of trust to be granted to an endpoint device

Management

The objective is to define how well any policy driven assessment of user devices is updated, reviewed and enforced.

Are policies constantly reviewed against known vulnerabilities?

- If policies exist, are they the responsibility of a clearly named owner?
- Ascertain whether there is a link between policy development and vulnerability identification
- Ensure that policy amendments are based on vulnerability identification
- Identify where policy amendments are tracked and recorded
- Identify how any asset inventory is amended with policy updates

Are devices constantly checked at the login stage?

- Determine whether there is a solution to check policy status at the login stage
- To what extent has a solution been implemented within the organization?
- Determine whether there is an integration of the solution with the other security and IT function
- Determine whether progress is clearly measured and reported
- Determine whether there is a governance process to manage the solution

Operational

The objective is to identify the extent to which any solution has been implemented so as to involve the end users in updating their devices, and how effectively it contributes to the rest of the security function.

Are users prompted to update devices to ensure access is maintained?

- Determine if there is a clearly defined process for identifying policy non-compliance
- Determine whether non-compliance is communicated to the end user
- Ascertain whether users are guided to update their devices to bring them within policy
- Determine whether the non-compliance is reported by users

Is device trustworthiness integrated into other functions?

- Determine whether any solution providing policy control over devices has been implemented
- Identify potential areas of security and IT in particular that can benefit from device trustworthiness
- What are the outputs from such a solution?
- In what method and manner are outputs shared with other functions, automated or otherwise?
- Create a benefits statement showing how enhanced device trustworthiness improves delivery in other functions

Components

Device Inventory Database. An up-to-date repository for information on all devices you allow to access the network, including type, purpose, network addresses, asset tags, components, configuration and responsible user or maintainer.

Managed Devices. Implement with a spectrum of toolsets and approaches. Agency-owned devices can be provided with a number of agents, including endpoint detection and response (EDR). BYOD can be provided with mobile device management (MDM). For other trusted but not fully managed devices, evaluate device trust and health.

Certificate Issuer. This is used to mark your managed or otherwise approved devices with a client-side certificate. Depending on which types of certificates you plan to use, the public key infrastructure (PKI) for this may already be part of another security product.

Challenges

Our end users are resistant to installing software on their devices.

For a variety of reasons, from concerns over privacy to local regulations, people dislike installing agents on their phones and devices. Select a technology to verify security hygiene and enabled security features during authentication without introducing privacy concerns.

We lack the resources to manage all devices.

Device management is often time intensive. For a zero trust architecture, the priority is ensuring the access devices and authentication devices are evaluated for trust. Steadily increase coverage in a scalable and light-touch way.

We already have a mobile device management (MDM) platform.

A full-featured platform may already be in place for device management. A portfolio of technologies for trusting devices, including MDM and VPN, will likely be used to establish your zero trust architecture. Begin with the investments that you have already made.

Metrics

Risk

- App risk reduction – critical apps

Security:

- Incident reduction
 - Compromised devices
- Auth devices – vulnerable devices
- Access devices – vulnerable devices

Support:


- Device MFA support tickets
- Device MFA support satisfaction (NPS)

Operations:

- Critical apps configured with MFA
- Critical apps configured with SSO
- Access devices – trusted devices

Phase 4: Adaptive Policies

Implement requirements for access based on the sensitivity of the resources and the known security state to manage risk levels appropriately. These policies can range from allowing only agency-managed devices to requiring certain versions of patched software, encryption or step-up authentication based on user behavior.



Description

We can set policies requiring people use known and approved endpoints to access the most critical data and applications (for example, privileged users must use an agency-owned device). The access proxy takes on the role of enforcing access to organizational resources, regardless of whether they're outside or inside the traditional perimeter. Enforcement strategy is one way we express risk tolerance; right-sizing those policies depends on factors such as sensitivity, threat, user community, regulatory requirements and other considerations.

A major drawback to the classic network perimeter security model was that organizations tended to have one level of trust everywhere on the inside. Building in different tiers required network segmentation that was often too complex to implement. The zero trust approach separates out levels of trust at the application layer, which is why the phased approach emphasizes determining where critical and sensitive data and applications are early on. In order to access them, users and their devices may need a higher level of trust, which means they need to pass more tests and comply with stricter requirements.

Start with a baseline level of trust for all users and all devices regardless of what they're accessing, and then add more to reach the level of risk management we need for access to the most sensitive tiers. We can determine the trust level of users by whether or not it's the first time we're seeing the person, what factors they used to authenticate, and if there are signs of impersonation or malicious activity. We can determine the trust level of devices by whether or not it's the first time we're seeing the device; whether or not it's managed; if it's running an up-to-date operating system and web browser; if security features like screen lock and encryption are enabled; and if there are signs of infection or tampering. Using this information, zero trust access policies must be flexible to allow access, allow access but nudge for remediation, require remediation to improve trust or altogether block access as untrustworthy.

The most important thing is to carve away at the devices, software, sources and behaviors you know you don't want to allow, thereby reducing your exposure to attacks. Changing the security lifestyle of an organization takes dedicated work, but once you have the controls fit more closely to where they belong – the users, their devices and the applications – you'll be addressing the gaps in today's traditional security paradigm and moving beyond it.

Objectives

- Establish an enforcement strategy to express our risk tolerance
- Apply policies to authentication workflows based on trust indicators
- Take action for the continuously validated people and devices

Transformation

Strategic

The objective is to identify and establish a clear strategy on policies based on assets and users controlling access to assets within the organization.

Have policies been implemented to control users' access based upon device and user trust?

- Is there a clear strategy to ensure that access to assets is controlled by policies based on a risk-based decision derived from the context of user and device?
- Identify ownership and responsibility for such a strategy
- Ascertain whether these policies are being applied in the organization

Management

The objective is to define how well any adaptive policy driven is managed and updated.

Are users managed consistently according to policies?

- If policies exist, are they the responsibility of a clearly named owner?
- Ensure that policy controls are clearly defined
- Identify processes by which policies are implemented
- Identify any measurements for the implementation of policies

Is there a process in place to update user privileges?

- Ascertain whether there is a link between policy development and policies managing devices and users
- Is there a clear set of policies to ensure changes to device and user policies are reflected in the controls on access to assets?
- Determine whether the output of risk assessments on assets, especially applications, is linked to any policy update process
- Are policy updates clearly measured and reported?
- Is there a governance process to manage the policy update process?

Operational

The objective is to identify whether there is visibility over devices and users, and processes to manage any exceptions.

Are out of policy users and devices identified and reported?

- Determine if there is a clearly defined process for identifying policy non-compliance of both users and devices
- Determine whether there is clear reporting on any non-compliance
- Ascertain whether the policy update process takes into account exception reporting
- Ascertain whether there is a clear ownership of exception reporting and actions required
- Identify whether exceptions are reporting to other IT and security functions

Is there an exception workflow in place to manage out-of-policy users?

- Determine whether exceptions are recognized and acted upon
- Identify how the output is integrated into operational actions
- Is that integration automated or manual?
- Are the processes measured and reported on?
- Is there a clear owner for the process and its continued development?

Components

Access Control Engine. The repository of all your access policies, such as “only this group of users, together with their up-to-date, assigned and managed devices, may use this sensitive application.”

Trust Inference. Deciding what conditions will cause you to place or lose trust in a given device (such as hardware changes). The trust inferrer will rely on a steady input of data from the sources you choose. Examples include checking to see whether the device is encrypted; whether it has all management agents working; whether its software is up to date; and whether all of the information about that device is current.

Challenges

We can't block users across the organization. Nothing will stop a security initiative faster than halting the organization's work. Start small with policies and perform policy modeling and testing to get the appropriate results.

Our risk management isn't well-defined and we haven't agreed on a risk tolerance. Rather than talk broadly about risk reduction, point to specific risks that policy addresses. For example, users will not be able to access email from jailbroken devices. Effective policy is explicit in the threat scenarios it prevents.

IT, security and the business unit are at odds over the policy. Different stakeholders often have different understandings and different priorities. Given the importance of policy decisions, be sure to include these stakeholders early on in this phase of the zero trust initiative.

Enforcing policies increases our help desk call volume. A key support metric to track is support tickets related to zero trust projects. Use this metric as feedback to lower the help desk impact on future projects in the journey to zero trust.

Metrics

Risk:

- Policy coverage for apps
- Policy coverage for devices

Security:

- Incident reduction

Support:

- Policy exceptions (workflow)
- Policy support tickets
- Policy support satisfaction (NPS)

Operations:

- Policies enforced
- Policy compliance
- Accounts - inactive
- Access devices - inactive
- Auth devices - inactive
- Auth devices - shared

Phase 5: Zero Trust for The Federal Workforce

By this point, all applications and systems within scope are covered by the previously listed stages. Monitoring and response to threat scenarios are ongoing. People have a consistent experience across devices and activities. It is time to optimize the management.

Description

The iterative approach detailed in this guide results in transforming one aspect of the organization's IT into a zero trust model. The final phase is integrating that aspect into the broader organization-wide zero trust architecture. The integration and move to continuous improvement must result in a consistent experience for end users, IT administration and security operations.

This final step begins with a look backwards. Ask these questions:

- **Were the success criteria met?**
- **Did the zero trust initiative produce the expected results?**
- **Did the initiative align with the organization's strategy and design principles?**

The gaps between the expectations and results – both positive and negative – point to areas to explore for lessons learned. By measuring and reporting benefits and results in a consistent way, we can compare the individual initiative to the collection of initiatives in the transformation, providing another area to explore for lessons learned. By looking backwards, we can better prepare for future initiatives, thus accelerating the overall transformation process.

Next, we integrate the initiative with the overall architecture. The incremental approach makes the transformation manageable and scalable. But it also means that, at any given moment, portions of the environment are not zero trust, portions are zero trust and portions are in transition. Without a clear integration strategy and plan, an organization can end up with several siloed environments. This, in turn, can drive up costs, overhead and complexity. Therefore, it's important to integrate the technologies, the IT processes and the security processes. The overall architecture must provide a consistent experience along a common set of capabilities, which each individual zero trust initiative snaps into.

Finally, we look forward to ongoing operations and future iterations. Tracking the metrics along the entirety of the zero trust architecture provides visibility into what is working and where improvements are needed. The individual initiative we tracked through the five phases is at an end, and it is handed off to operations for ongoing support and improvement. Future initiatives will continue the momentum, moving more of the organization's people, devices and applications onto the zero trust architecture.

Objectives

- Integrate the zero trust initiative into the organization-wide zero trust architecture
- Shift towards continuous improvement
- Capture lessons learned and momentum to drive future zero trust initiatives

Transformation

Strategic

The objective is to ascertain whether a zero trust strategy being implemented achieves the benefits, and to identify if there is a clear alignment of a zero trust strategy with the IT and business units.

Is there a clear statement of benefits that zero trust delivers to the overall IT and business strategy?

- Identify whether there is a clear zero trust strategy with a set of principles
- Does the strategy have a clear statement of what benefits are expected as a result of that strategy?
- Is there a clear link between the strategy and the organization's IT and business strategy?
- Identify ownership and responsibility for the continued development of such a strategy
- Are benefits measured and reported in a systematic manner?

Management

The objective is to define how well any zero trust strategy has been implemented.

Has a clear zero trust architecture been defined and implemented?

- Does the strategic architecture cover the organization's total IT environment?
- Discover the extent to which the architecture has been implemented
- Identify processes by which architecture is reviewed and improvements or adaptations identified
- Are the architecture's benefits measured and reported in a systematic manner?

Operational

The objective is to identify whether the zero trust strategy has been implemented to improve end-user experience while reducing security risk.

Has there been an integration into the security function?

- Ensure that all elements within the security function or impacted by the security function have been clearly identified
- Have the changes been documented with regard to responsibilities and process change?
- Has the integration been measured and exceptions identified?
- Is there a clear program for continuous improvement?
- Have the benefits and improvements in the security function been identified?

Challenges

We have legacy technology. Many organizations have to support technology that's years or decades old. Not everything has to plug into the zero trust architecture. The objective is mindfully moving workflows into the security model to reduce risk and improve security where it makes the most sense.

We won't be ready for a zero trust model for 3 to 5 years. Adopting new patterns and practices takes significant time and planning. The journey to zero trust covered in this paper uses a series of projects to secure workflows, allowing you to move as quickly or as deliberately as you need.

Metrics

- Number of workflows upgraded into the zero trust architecture

Summary

This guide has laid out the journey for a zero trust transformation. One key to success is specificity. Scope the initiative around a specific activity, that is, a set of applications used by a specific set of users, completed in support of an organization’s function. Be specific in the threat scenarios that we’re avoiding by requiring people and devices to establish trust. Another key to success is iterating. Launch an initiative to transform one activity onto zero trust architecture. Learn along the way, and then repeat.

For two decades, we have discussed eliminating excess trust. We have debated telemetry and monitoring techniques to continuously evaluate trust. Finally, in recent years, technology has caught up with the philosophy. By following the steps in this guide, your organization can implement these ideas with off-the-shelf software at a sustainable pace. The zero trust revolution is well under way. Join us.



About Duo Security

Duo Security, now part of Cisco, is the leading multi-factor authentication (MFA) and secure access provider. Duo comprises a key pillar of Cisco Secure’s Zero Trust offering, the most comprehensive approach to securing access across IT applications and environments, from any user, device, and location. Duo is a trusted partner to more than 25,000 customers globally, including Bird, Facebook, Lyft, University of Michigan, Yelp, Zillow and more. Founded in Ann Arbor, Michigan, Duo also has offices in Austin, Texas; San Francisco, California; and London. Try it for free at signup.duo.com

Duo’s FedRAMP authorized editions: duo.com/editions-and-pricing/duo-federal-editions

Cisco Case Study: duo.com/solutions/customer-stories/cisco