Planning for a Zero Trust Architecture:

A Starting Guide for Administrators

Scott Rose
Advanced Network Technologies Division
Information Technology Laboratory

August 4, 2021

This publication is available free of charge from:
https://doi.org/10.6028/NIST.CSWP.08042021-draft

National Institute of
Standards and Technology
U.S. Department of Commerce

Abstract

Zero trust is a set of cybersecurity principles used when planning and implementing an enterprise architecture. Input and cooperation from various stakeholders in an enterprise is needed in order for a zero trust architecture to succeed in improving the enterprise security posture. Some of these stakeholders may not be familiar with risk analysis and management. This document provides a quick overview of the NIST Risk Management Framework (NIST RMF) and how the NIST RMF can help in developing and implementing a zero trust architecture.

Keywords

architecture; information technology; risk; zero trust.

Disclaimer

Any mention of commercial products or reference to commercial organizations is for information only; it does not imply recommendation or endorsement by NIST, nor does it imply that the products mentioned are necessarily the best available for the purpose.

Additional Information

For additional information on NIST's Cybersecurity programs, projects and publications, visit the Computer Security Resource Center. Information on other efforts at NIST and in the Information Technology Laboratory (ITL) is also available.

Zero trust related information is also found on the zero trust topic page.

Public comment period: *August 4, 2021 through September 3, 2021*

National Institute of Standards and Technology
Attn: Advanced Network Technologies Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8920) Gaithersburg, MD 20899-8920
Email: zerotrust-arch@nist.gov

All comments are subject to release under the Freedom of Information Act (FOIA).

Acknowledgments

The author would like to thank the members of the NIST Risk Management Framework team and the Zero Trust Architecture project team for their input and review.

Audience

This document was written to help enterprise administrators and system operators understand how the various roles and tasks in the NIST Risk Management Framework (RMF) can be used when moving to a zero trust architecture. This document briefly introduces zero trust, and how the RMF process can be used in a zero trust migration process. It is assumed that the reader is familiar with the concepts of zero trust as described in NIST SP 800-207 and has had exposure to federal information security practices.

Trademark Information

63 **Table of Contents**

70

1 Zero Trust

Zero trust (ZT) is the set of principles upon which information technology architectures are planned, deployed, and operated [1]. ZT uses a holistic view that considers all potential risks to a given mission or business process and how they are mitigated. As such, there is no single specific infrastructure implementation or architecture, but it depends on the workflow (i.e., part of the enterprise mission) being analyzed and the resources that are used in performing that workflow. Zero trust strategic thinking can be used to plan and implement an enterprise IT infrastructure, which then could be said to be a zero trust architecture (ZTA).

Enterprise administrators and system operators need to be involved in the planning and deployment for a ZTA to be successful. ZTA planning requires input and analysis from system and workflow owners as well as professional security architects. Zero trust cannot be imposed from above onto an existing workflow but needs to be integrated into all aspects of the enterprise. This paper introduces some of the concepts in the NIST Risk Management Framework (RMF) to administrators and operators. The RMF lays out a set of processes and tasks that is integrated into enterprise risk analysis, planning, development, and operations. Administrators who may normally not perform the tasks detailed in the RMF may find that they will need to become familiar with them as they migrate to a ZTA.

NIST Special Publication 800-207 [1] gives a conceptual framework for zero trust. While not comprehensive to all information technology it can be used as a tool to understand and develop a ZTA for an enterprise. NIST SP 800-207 also provides an abstract logical architecture that can be used to map solutions and gaps upon. The abstract architecture is repeated in figure 1 below.

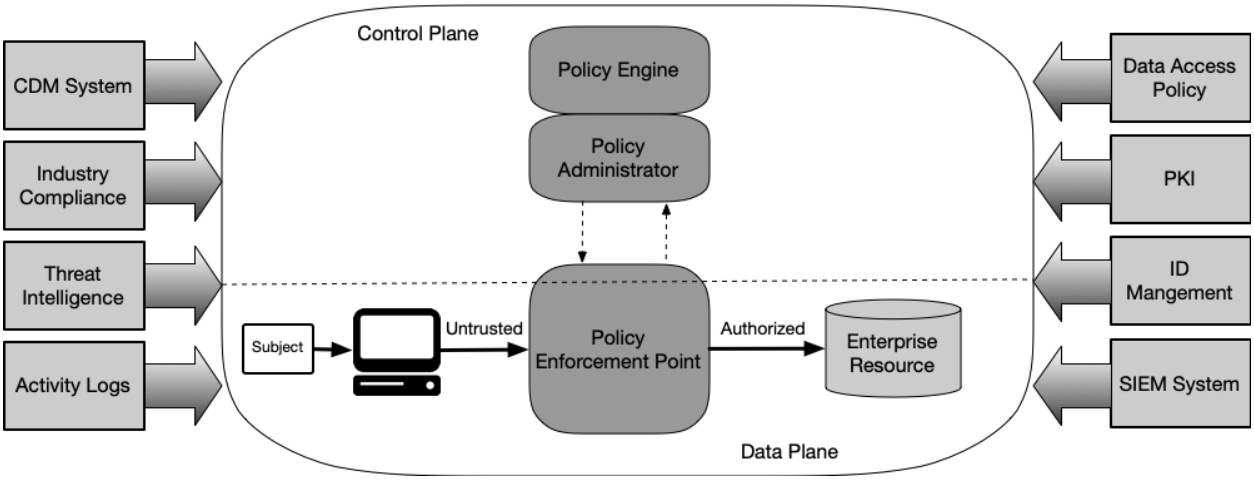

Figure 1: Abstract Zero Trust Logical Architecture

In this diagram, the components are listed as their logical function, and thus do not necessarily represent a single operational system. It is possible that multiple components may serve one logical function in a distributed manner, or a single solution may fulfill multiple logical roles. The roles are described in the SP, but to summarize:

- **Policy Engine (PE):** The "brain" of a ZTA implementation and the components that ultimately evaluate resource access requests. The PE relies on information from the

various data sources (access logs, threat intelligence, device health and network ID authentication checks, etc.)

- **Policy Administrator (PA):** The executor function of the PE. The PA's role is to establish, maintain and ultimately terminate sessions in the data plane. The PA, PE and PEP communicate on a logically (or physically) separate set of channels called the control plane. The control plane is used to establish and configure the channels used to send application traffic (i.e. the data plane).
- **Policy Enforcement Point (PEP):** The component that applications, devices, etc. will interact with to be granted access permission to a resource. The PEP is responsible for gathering information for the PE and following the instructions issued by the PA to establish and terminate communication sessions. All data plane communications (i.e. all workflow application traffic) between enterprise resources must go through a PEP.
- **Information Feeds:** This is the set of policies, identity and device attributes, environmental factors and historical data used by the PE to generate resource access decisions.

1.1 Tenets of Zero Trust

Zero trust could be summarized as a set of principles (or tenets) used to plan and implement an IT architecture. The tenets below were originally defined in NIST SP 800-207 [1] but are repeated here and grouped as tenets relating to network identity, device health, or data flows. Some discussion of the tenets is included, and some considerations that planners should keep in mind when developing a zero trust architecture.

1.1.1 Tenets that Deal with Network Identity Governance

I. **All resource authentication and authorization are dynamic and strictly enforced before access is allowed.** A typical enterprise has a wide collection of network identities: end users, service accounts, etc. Some end users may have multiple network identities, and some identities may only be used by hardware/software components. The enterprise needs to have a governance policy and structure in place so that only authorized operations are performed, and only when the identity has properly authenticated itself. The enterprise needs to consider if their current identity governance policies are mature enough and where and how are authentication and authorization checks currently performed.

1.1.2 Tenets that Deal with End Devices

I. **All data sources and computing services are considered resources.** An enterprise relies on different resources to perform its mission: mobile devices, data stores, compute resources (including virtual), remote sensors/actuators, etc. All of these components need to be considered in a ZTA. Some components (e.g. IoT sensors) may not be able to support some solutions such as configuration agents, app sandboxing, etc. so alternatives that use the underlying network infrastructure may be needed. If the resource lack certain security capabilities, the enterprise may need to add a PEP component to provide that functionality.

II. **The enterprise monitors and measures the integrity and security posture of all owned and associated assets.** This tenet deals with the aspects of cyber hygiene: configuration, patching, application loading, etc. The state of resources should be monitored and appropriate action taken when new information such as a new vulnerability or attack is reported or observed. The confidentiality and integrity of data on the resource should be protected. This requires enterprise admins to know how resources are configured, maintained, and monitored.

1.1.3 Tenets that Apply to Data Flows

I. **All communication is secured regardless of network location.** In zero trust, the network is always considered contested. There should be an assumption that an attacker is present on the network and could observe/modify communications. Appropriate safeguards should be in place to protect the confidentiality and integrity of data in transit. If the resources cannot provide this functionality natively, a separate PEP component may be necessary.

II. **Access to individual enterprise resources is granted on a per-session basis.** In an ideal zero trust architecture, every unique operation would undergo authentication and authorized before it is performed. For example, a delete operation following a read operation to a database should trigger an authentication and authorization check. This is may not always possible and other mitigating solutions such as logging and backups may be needed to detect and recover from unauthorized operations. Enterprise administrators will need to learn how to enforce fine grain access policies on individual resources. If the current set of tools do not allow this, other solutions such as logging, versioning tools, or backups may help mitigate risk.

III. **Access to resources is determined by dynamic policy—including the observable state of client identity, application/service, and the requesting asset—and may include other behavioral and environmental attributes.** In zero trust, the default behavior for all resources is to deny all connections with an allow list. The members of this allow list must authenticate themselves and prove they meet the enterprise policy to be granted the session. This may include meeting requirements such as client software versions, patch level, geolocation, historical request patterns, etc. Note that it may not be possible to perform all check immediately prior to the access request, but some may be performed recently (e.g. daily software versioning checks).

IV. **The enterprise collects as much information as possible about the current state of assets, network infrastructure and communications and uses it to improve its security posture.** Zero trust adds a dynamic response factor that was lacking (or not possible) in previous perimeter based architectures. System logs and threat intelligence are used to refine or change policy in response to new information. For example, a new vulnerability in a software component in use in the enterprise is announced. A zero trust enterprise would move quickly to quarantine the affected resources until they can be

patched or modified to mitigate the newly discovered vulnerability. Enterprise admins will need to set up and maintain a comprehensive monitoring and patching program for the enterprise and should consider how automated tools could assist in responding to newly discovered threats.

2 Getting Started on the Journey

Moving to a zero trust architecture will likely never start from scratch, but will involve a series of upgrades and changes over time. Some changes may be simple configuration changes, and some may involve the purchase and deployment of new infrastructure; it all depends on what is currently used and available to the enterprise.

The process of migrating to a ZTA is not a unique process and is similar to other cybersecurity upgrades, improvements, etc. Existing frameworks such as the NIST Risk Management Framework (RMF) [2] and Cybersecurity Framework (CSF) [3] can help an enterprise discuss, develop, and implement a ZTA. In the following sections, the RMF will be used to describe a series of steps and processes that could be used to migrate a workflow to a ZTA.

Additionally, there is the Federal CIO Handbook [4] that provides information and links to relevant policies, mandates and programs that apply to federal agencies. This includes programs like the DHS Continuous Diagnostics and Mitigation (CDM) program and the Trusted Internet Connection (TIC) policy that can provide additional guidance and tools for federal agency administrators as well as planners and managers.

2.1 The Process

NIST SP 800-37, Revision 2 [2] describes the Risk Management Framework methodology and its seven steps:

- Organizational and system preparation (PREPARE step)
- System categorization (CATEGORIZE step)
- Control selection (SELECT step)
- Control implementation (IMPLEMENT step)
- Control assessment (ASSESS step)
- System authorization (AUTHORIZE step)
- Control monitoring (MONITOR step)

While the steps are described in order, after initial implementation, they may be carried out or revisited in any sequence. The individual tasks that make up the seven steps could be conducted and revisited as needed, and possibly in parallel with other steps/tasks. The transitions between steps can be fluid (see figure 2). This is true when developing and implementing a ZTA, as the dynamic nature of zero trust may require a reiteration or rapid transitions in the RMF steps to respond to new information or technology changes. The details of the individual steps are documented in NIST SP 800-37r2 [2] and the accompanying Quick Start Guide [5].

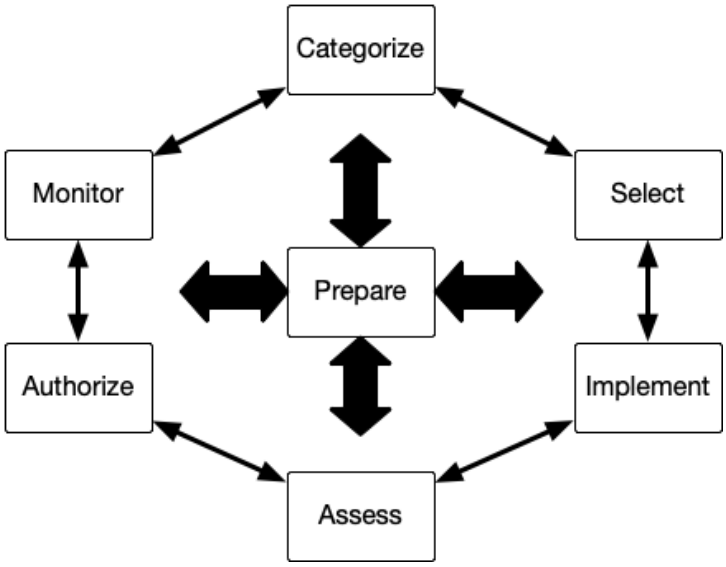

Figure 2: RMF State Machine

For an initial migration, the steps are usually followed in order (but it is not necessary). The RMF steps are very similar to the high-level steps developed for zero trust by John Kindervag [11-12] and are partially mapped below. This process assumes the authorization boundary has been created and the system components used in the workflow are known (i.e. the PREPARE step has been performed and data collected). The is no explicit CATEGORIZE step as this high level description was not developed with federal agencies in mind.

1. Map the attack surface of the resource and identify the key parts that would be targeted by a malicious actor. These will be covered by the tasks in the SELECT step.
2. From the PREPARE step (tasks P-12 and P-13), the data flows should be identified and mapped.
3. The IMPLEMENT step: Focus on implementing the controls from the SELECT phase on the resource and related PEP. The PEP may be a separate software component from the resource itself and is used to meet authentication/authorization related controls. The underlying network should not be considered trusted, so links between individual resources must pass through a PEP.
4. The ASSESS Step: Make sure all access policies developed and put in place during the IMPLEMENT step are implemented and operating as intended. This would conclude with the AUTHORIZE step, where the system and workflow is considered in a state to begin actual operation.
5. The MONITOR step: Implement the monitoring and management process for the resource (and its security posture).

2.1.1 Prepare

The first step in the RMF process is the PREPARE. When starting the zero trust transition, this step will may be the longest as a full inventory of roles and enterprise resources is the foundation of ZT. The Prepare step includes steps and tasks applicable to the Organization and

mission/business levels and at the system level. System architects, administrators and operators will likely only focus system level-based tasks in the PREPARE step but may have valuable input to the mission/business level tasks. The PREPARE step is primarily focused on preparing the organization to manage its security and privacy risks using the NIST RMF, and setting up essential activities at the organization, mission and business process, and system levels.

The enterprise architecture team should focus on identifying relevant business processes (workflows) and systems at the RMF mission/business level. A risk analysis should be done on each workflow. The owners and key personnel involved in the workflow should be identified and have input in the analysis, as they may have knowledge and experience about the workflows that deviate from existing workflow or system documentation. This maps to the organization level tasks (P-3 to P-7) of the PREPARE step [2].

System administrators and operators should focus on identifying the resources that are used to conduct the identified business processes. These map to the system level tasks (P-8 through P-18) of the PREPARE step [2]. This covers:

- Resources involved in each workflow that will be the subject of the security plan. Resources could fall under two different categories:
    - Workflow specific resources that are used to directly support the given workflow. Examples would include a single purpose report database and cloud-based application used to submit reports to that database.
    - General infrastructure resources that are shared by several (or all) workflows. Examples include network infrastructure (switches, wireless network access points, etc.), DNS, email, etc.
- Network identities and governance tools used within the organization. This is not just a list of end user accounts, but includes service accounts used by software components, device IDs, etc.
- Any data classification programs and procedures used within the organization.
- The current state of monitoring of enterprise resources. One of the foundations of zero trust is knowledge of data flows in the enterprise. It is vital that an enterprise have a solid continuous monitoring plan and toolset that can be leveraged before implementing a zero trust architecture.

Once the foundational work of identifying unique workflows and enterprise resources has been done, the authorization boundaries can be produced (task P-11). Architects and security advisors should “draw” the boundaries to include only the core required components of the system. The authorization boundary will likely include any PEP component that provides security capabilities. Connection between resources within the authorization boundary must also be secure and not implicitly trusted. Zero trust principles consider the network contested and so connections between resources within the authorization boundary are subject to the same controls as connections crossing the authorization boundary (i.e., from outside to within the boundary and vice versa). Controls that are covered by PEP components may be reusable in other systems if the same PEP solution is used with other resources, such as some cloud secure access broker (CASB) or similar solutions when used to provide the PEP component for multiple different resources (see NIST SP 800-37r2 Appendix G).

2.1.2 Categorize

This step does not change in a ZT planning process. FIPS 199 [5] and FIPS 200 [7] are used to place resources in a LOW, MODERATE or HIGH category based on its confidentiality, integrity, and availability requirements in the workflow. The owners of the resource and workflows that use the resource can be valuable input in this set of tasks.

2.1.3 Select

This step also does not change in a ZT planning process. The baseline controls for LOW, MODERATE and HIGH-impact systems are listed in NIST SP 800-53B [8]. Additional controls may be added or removed as part of control tailoring, adjusting the controls to manage risk to the resource and its position in the workflow. The use of overlays[1] may assist in this, but the overlay should not be considered immutable, but may need to be adjusted for the unique resource. The planners should also consider what controls will be met by the PEP, and what may need to be implemented in the resource itself. As with the CATEGORIZE step, the resource owners and owners of the workflows that use the resource may provide valuable input in this step. As zero trust places importance on continual monitoring and updating of security postures, cybersecurity architects and administrators need to develop a comprehensive monitoring process that can handle the volume of data needed for the dynamic nature of ZT.

In addition to NIST SP 800-53 [9] and SP 800-53B [10], enterprise architects and administrators may wish to consult other resources as necessary such as the CIO Handbook [4] and TIC 3.0 documents and use cases [9] for other requirements. In particular the TIC 3.0 use case documents may provide a high level, initial playbook for a potential architecture. These documents may help in developing the desired set of requirements and security properties for the resource.

2.1.4 Implement

The IMPLEMENT step, like the two previous steps, does not have any ZT specific concerns. However, as with the RMF and ZT, future monitoring/maintenance operations should be kept in mind. Administrators may want to avoid solutions that involve frequent human required actions or do not easily fit into monitoring systems. ZT encourages automation to have dynamic responses to changing security concerns and manual changes may not be able to keep up with frequent changes.

2.1.5 Assess

In a zero trust architecture, the assessment of security controls should be continual in the face of a changing environment. Modern IT environments and trends like DevOps/DevSecOps mean that a snapshot in time assessment of a system quickly becomes outdates as improvements and

[1] An overlay offers organizations additional customization options for control baselines and may be a fully specified set of controls, control enhancements, and other supporting information (e.g., parameter values) derived from the application of tailoring guidance. Overlays also provide an opportunity to build consensus across communities of interest and develop a starting point of controls that have broad-based support for very specific circumstances, situations, and/or conditions.

configuration changes are done to mitigate newly discovered threats or changes to the enterprise infrastructure.

In response, the ASSESS step should be thought of as comprising two assessment processes: continual assessment of the system, and one of the processes used to manage the system. The process must be assessed as the dynamic nature of zero trust means that the system will likely change quicker than a human performed assessment program can manage at scale. This assessment takes factors like the change process into consideration to assess how the system is modified.

The assessment of the system itself should have a continual assessment component based on a monitoring program [13]. Frequent automated checks or scans should be conducted to detect changes in the system. Logging data should be used to detect possible malicious behavior that requires further investigations or remediation. This assessment may also include active processes such as red team testing of the system as input into the assessments.

2.1.6 Authorize

This step may evolve interpretation in a zero trust architecture (as in the ASSESS step above), but the goal remains the same. As a ZTA is built to be more dynamic and fluid to respond to changing network conditions, authorizations should not be viewed as to a static system, but the system and its processes for changes or updates.

2.1.7 Monitor

As stated previously, zero trust requires the enterprise to monitor the resources used to conduct its primary mission(s). Exactly how this is done depends on the technology solutions in place in the enterprise. However, regardless of the technology, the enterprise should have policies in place to trigger actions based on behaviors seen in monitoring. This may include reacting to security events or tied to a DevOps process to modify or improve the system.

In addition to monitoring the current activity and state of enterprise resources, cybersecurity planers should consider how external threat intelligence can help in pre-emptive responses to new conditions. A tool like the .GOVCAR [14] may be useful in prioritizing threats to be addressed. For federal agencies there are also additional monitoring programs that may assist such as DHS CDM dashboards [15] and the AWARE [16] program.

2.1.8 RMF Operational Loops

Zero trust lends itself to the use of more dynamic DevOps and DevSecOps style operations. The cycles of security updates and reviews could be described as involving a subset of the RMF process. For example, a DevOps cycle for the cybersecurity posture could be expressed as figure 3 below:

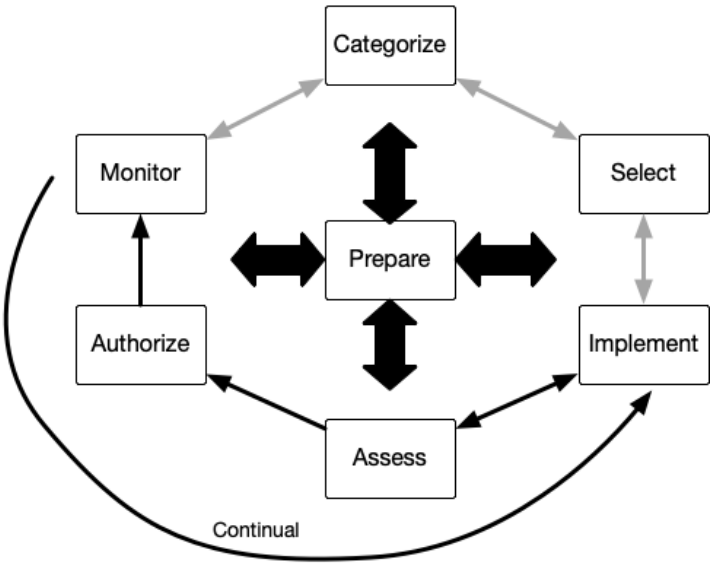

362

363 **Figure 3: DevOps Cycle**

364 In this loop, the data collected in the MONITOR step then feed back into the
365 IMPLEMENTATION step as improvements and refinements are implemented, they are then
366 assessed and follow the continual AUTHORIZE step to enter operations. If necessary, the
367 DevSecOps team may even fall back to the SELECT step if new information leads to new
368 controls to be added or existing controls to be removed.

369 Even in a more static IT operational environment (i.e., no DevOps), a zero trust model could be
370 seen as a loop of only operations. In this loop, there is no DevOps component so the ASSESS
371 and AUTHORIZE steps are continually cycled as new information is gathered from system logs,
372 threat intelligence, etc. This may lead to new configuration changes or policy updates. Larger
373 changes to the operations will be less frequent and involve a longer cycle as other steps outside
374 of the loop are performed if new information requires a larger change.

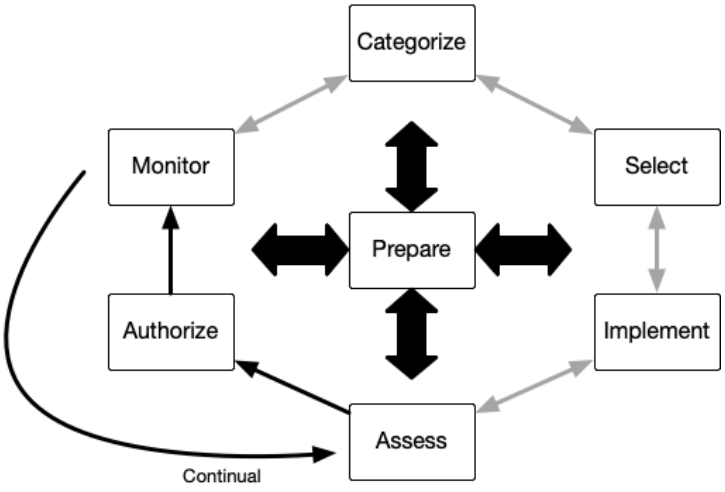

375

376 **Figure 4: Operations Cycle**

3 Conclusion

Zero trust is not a single technology solution, but a larger cybersecurity strategy and operational practice. A successful zero trust architecture requires the cooperation of cybersecurity planners, management, and administration/operations. Zero trust also requires the involvement of system, data, and process owners who may not traditionally provide input on the risks to their charges. This input is vital; zero trust is a holistic approach to enterprise cybersecurity and therefor needs support from every individual in the enterprise.

The NIST Risk Management Framework provides a toolset developed to help those who conduct risk assessments. However, it can also help administrators and operators and others that do not primarily focus on cybersecurity. This white paper provides a quick overview of the NIST RMF and provides links and pointers on how administrator and operators can begin understanding the steps of RMF and how these steps support zero trust. The goal is to provide pointers to IT staff to help them understand how their roles may evolve in a ZTA and where risk management staff need to bring in other IT staff to assist in their analysis.

References

[1] Rose S, Borchert O, Connelly S, Mitchel S. Zero Trust Architecture. (National Institute of Standards and Technology, Gaithersburg MD), NIST Special Publication (SP) 800-207. https://doi.org/10.6028/NIST.SP.800-207

[2] Joint Task Force (2018) Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-37, Rev. 2. https://doi.org/10.6028/NIST.SP.800-37r2

[3] National Institute of Standards and Technology (2018) Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1. (National Institute of Standards and Technology, Gaithersburg, MD). https://doi.org/10.6028/NIST.CSWP.04162018

[4] Chief Information Officers Council (2021), *CIO Handbook*. Available at https://www.cio.gov/cio-handbook/

[5] National Institute of Standards and Technology (2021) *About the NIST Risk Management Framework (RMF)*. Available at https://csrc.nist.gov/projects/risk-management/about-rmf

[6] National Institute of Standards and Technology (2004), Standards for Security Categorization of Federal Information and Information Systems. (U.S. Department of Commerce, Washington, DC), Federal Information Processing Standards Publication (FIPS) 199, February 01, 2004. https://doi.org/10.6028/NIST.FIPS.199

[7] National Institute of Standards and Technology (2006), Minimum Security Requirements for Federal Information and Information Systems. (U.S. Department of Commerce, Washington, DC), Federal Information Processing Standards Publication (FIPS) 200, March 01, 2006. https://doi.org/10.6028/NIST.FIPS.200

[8] Joint Task Force (2020) Security and Privacy Controls for Information Systems and Organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-53, Rev. 5. Includes updates as of December 10, 2020. https://doi.org/10.6028/NIST.SP.800-53r5

[9] Department of Homeland Security, Cybersecurity & Infrastructure Security Agency (2021). *TIC 3.0 Core Guidance Documents*. Available at https://www.cisa.gov/publication/tic-30-core-guidance-documents

[10] ON2IT (2020) *A hands-on-approach to Zero Trust implementation*. Available at https://on2it.net/wp-content/uploads/2020/01/hands-on-approach-zero-trust-implementation.pdf

[11] Kindervag J (2017) 'Zero Trust': The Way Forward in Cybersecurity (DarkReading). Available at https://www.darkreading.com/attacks-breaches/zero-trust-the-way-forward-in-cybersecurity/a/d-id/1327827

[12] Dempsey K, Chawla NS, Johnson A, Johnston R, Jones AC, Orebaugh A, Scholl M, and Stein K. Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations. (National Institute of Standards and Technology, Gaithersburg MD), NIST SP 800-137. https://doi.org/10.6028/NIST.SP.800-137

[13] Department of Homeland Security, Cybersecurity & Infrastructure Security Agency (2020) *CDM Program: What is .govCAR? Fact Sheet*. Available at https://www.cisa.gov/publication/cdm-program-what-govcar

[14] Department of Homeland Security, Cybersecurity & Infrastructure Security Agency (2020) *CDM Program: Dashboard Ecosystem Fact Sheet*. Available at https://www.cisa.gov/publication/cdm-program-dashboard-ecosystem

[15] Department of Homeland Security, Cybersecurity & Infrastructure Security Agency (2020) *CDM Program: AWARE Scoring Fact Sheet*. Available at https://www.cisa.gov/publication/cdm-program-aware-scoring-fact-sheet