PROJECT DESCRIPTION

IMPLEMENTING A ZERO TRUST ARCHITECTURE

Alper Kerman

National Cybersecurity Center of Excellence
National Institute of Standards and Technology

Oliver Borchert
Scott Rose

Advanced Network Technologies Division
National Institute of Standards and Technology

Eileen Division
Allen Tan

The MITRE Corporation

October 2020

nccoe-zta-project@list.nist.gov

This revision incorporates comments from the public.

The National Cybersecurity Center of Excellence (NCCoE), a part of the National Institute of Standards and Technology (NIST), is a collaborative hub where industry organizations, government agencies, and academic institutions work together to address businesses' most pressing cybersecurity challenges. Through this collaboration, the NCCoE develops modular, adaptable example cybersecurity solutions demonstrating how to apply standards and best practices by using commercially available technology. To learn more about the NCCoE, visit https://www.nccoe.nist.gov. To learn more about NIST, visit https://www.nist.gov.

This document describes a challenge that is relevant to many industry sectors. NCCoE cybersecurity experts will address this challenge through collaboration with a Community of Interest, including vendors of cybersecurity solutions. The resulting reference design will detail one or more approaches that can be incorporated across multiple industry sectors.

ABSTRACT

The proliferation of cloud computing, mobile device use, and the Internet of Things has dissolved conventional network boundaries. The workforce is more distributed, with remote workers who need access to resources anytime, anywhere, and on any device, to support the mission. Enterprises must evolve to provide secure access to company resources from any location and asset, protect interactions with business partners, and shield client-server as well as inter-server communications.

A zero trust cybersecurity approach removes the assumption of trust typically given to devices, subjects (i.e., the people and things that request information from resources), and networks. It focuses on accessing resources in a secure manner, regardless of network location, subject, and asset, and enforcing risk-based access controls while continually inspecting, monitoring, and logging interactions. This requires device health attestation, data-level protections, a robust identity architecture, and strategic micro-segmentation to create granular trust zones around an organization's digital resources. Zero trust evaluates access requests and communication behaviors in real time over the length of open connections, while continually and consistently recalibrating access to the organization's resources. Designing for zero trust enables enterprises to securely accommodate the complexity of a diverse set of business cases by informing virtually all access decisions and interactions between systems and resources.

This NCCoE project will show a standards-based implementation of a zero trust architecture (ZTA). Publication of this project description begins a process that will further identify project requirements and scope, as well as the hardware and software components to develop demonstrations. The NCCoE will build a modular, end-to-end example ZTA(s) using commercially available technology that will address a set of cybersecurity challenges aligned to the NIST Cybersecurity Framework. This project will result in a freely available NIST Cybersecurity Practice Guide.

KEYWORDS

cybersecurity; enterprise; identity and access management; network security; remote access; zero trust; zero trust architecture

DISCLAIMER

Certain commercial entities, equipment, products, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by NIST or NCCoE, nor

is it intended to imply that the entities, equipment, products, or materials are necessarily the best available for the purpose.

TABLE OF CONTENTS

1 Executive Summary

Purpose

Conventional network security has focused on perimeter defenses—once inside the network perimeter, subjects (i.e., end users, applications, and other non-person entities that request information from resources) are often given broad access to multiple corporate resources. If the subjects are compromised, malicious actors—through impersonation and escalation—can gain access to the resources from inside or outside the network. Moreover, the growth in cloud computing, Internet of Things (IoT), business partners, and the growing number of remote workers raises the complexity of protecting an organization's digital resources, because more points of entry, exit, and data access exist than ever before.

Organizations are rethinking the conventional network security perimeter. A zero trust architecture (ZTA) addresses this trend by focusing on protecting resources, not network perimeters, as the network location is no longer viewed as the prime component to the security posture necessary for a resource.

Zero trust is a set of cybersecurity principles used to create a strategy that focuses on moving network defenses from wide, static network perimeters to focusing more narrowly on subjects, enterprise assets (i.e., devices, infrastructure components, applications, virtual and cloud components), and individual or small groups of resources. A ZTA uses zero trust principles to plan and protect an enterprise infrastructure and workflows. By design, a ZTA environment embraces the notion of no implicit trust toward assets and subjects, regardless of their physical or network locations (i.e., local area networks versus the internet). Hence, a ZTA never grants access to resources until a subject, asset, or workload are verified by reliable authentication and authorization.

This document defines a National Cybersecurity Center of Excellence (NCCoE) project to help organizations design for zero trust. This project will produce an example implementation(s) of a ZTA, using commercially available technology designed and deployed according to the concepts and tenets documented in National Institute of Standards and Technology (NIST) Special Publication (SP) 800-207, *Zero Trust Architecture* [1]. The primary objective of this project is to demonstrate a proposed architecture(s) that brings into play different enterprise resources (e.g., data sources, computing services, and IoT devices) that are spread across on-premises and cloud environments that inherit the ZTA solution characteristics outlined in NIST SP 800-207.

Another objective of this project is to document the impacts on administrator and end-user experience because of employing a ZTA strategy.

This project will result in a publicly available NIST Cybersecurity Practice Guide, a detailed implementation guide of the practical steps needed to implement a cybersecurity reference design that addresses the project goals.

Scope

The scope of this project is limited to implementing a ZTA for a conventional, general purpose enterprise information technology (IT) infrastructure that combines users (including employees, contractors, guests, and non-person entities); assets; and enterprise resources. Resources could be hosted and managed—by the corporation itself or a third-party provider—on premise, in the cloud, at the edge, or some combination of these. There may also be branch or partner offices, teleworkers, and bring-your-own-device (BYOD) usage.

This project will focus primarily on access to enterprise resources. More specifically, the focus will be on behaviors of enterprise employees, contractors, and guests accessing enterprise resources while connected from the corporate (or enterprise headquarters) network, a branch office, or the public internet. Access requests can occur over both the enterprise-owned part of the infrastructure as well as the public/non-enterprise-owned part of the infrastructure. This requires that all access requests be secure, authorized, and verified before access is enforced, regardless of where the request is initiated or where the resources are located.

ZTAs for industrial control systems and operational technology (OT) environments are explicitly out of scope for this project. However, the project seeks to provide an approach and security principles for a ZTA that could potentially be extended to OT environments.

Challenges

Many organizations are looking to build for zero trust, but challenges exist. Current challenges to implementing a ZTA include:

- Maturity of vendor products to support a ZTA.
- Organization's ability/willingness to migrate to a ZTA because of:
    - heavy investment in other (legacy) technologies
    - absence of, or deficiency in, identity governance
    - lack of ability/resources to develop a transition plan, pilot, or proof of concept
- Security concerns such as:
    - compromise of the zero trust control plane
    - ability to recognize attacks and detect malicious insiders
- Interoperability considerations of ZTA products/solutions with legacy technologies such as:
    - standard versus proprietary interfaces
    - ability to interact with enterprise and cloud services
- User experience. To date, there has been no detailed examination of how a ZTA would or could affect end-user experience and behavior. The goal of a ZTA should be to enhance security in a way that is transparent to the end user.

This practice guide aims to mitigate these challenges, using the solutions and collaborators selected for the demonstration project.

Background

Historically, the perimeter-based network security model has been the dominant model for information security. It assumes users inside the corporate network perimeter are "trusted" and anyone on the outside is "untrusted." For several decades, this view of trust has served as the basis for determining what resources a subject/asset can access.

Several high-profile cyber attacks in recent years, including the Office of Personnel Management breach in 2015, have undermined the case for the perimeter-based model [2]. Moreover, the perimeter is becoming less relevant due to several factors, including the growth of cloud computing, mobility, and changes in the modern workforce. It is with this backdrop that the Federal Chief Information Officer (CIO) Council [3] engaged the NIST NCCoE in 2018 to help federal agencies coalesce around a definition for ZTA and understand the benefits and

limitations of a ZTA. The interagency collaboration resulted in publication of NIST SP 800-207, *Zero Trust Architecture*.

This NCCoE project description builds on the work with federal agencies and the Federal CIO Council as we seek to build and document one or more demonstrable ZTAs, using commercially available products that align to the concepts and principles in NIST SP 800-207.

2 SCENARIOS

Responses from industry organizations that express interest in taking part in this project will affect the potential scenario-set in terms of the composition and number of scenarios demonstrated. These scenarios encapsulate the notion of providing subjects access to corporate resources hosted on premise or in the cloud. Access requests may come from within the enterprise network or the public internet, in the case of teleworkers. It is assumed the enterprise is implementing a ZTA within an existing typical corporate environment.

Scenario 1: Employee Access to Corporate Resources

An employee is looking for easy and secure access to corporate resources, from any work location. This scenario will demonstrate a specific user experience where an employee attempts to access corporate services such as the corporate intranet, a time-and-attendance system, and other human resources systems by using either an enterprise-managed device or a personally owned device. The ZTA solution implemented in this project will enforce the associated access request, dynamically and in near real-time. The employee will be able to perform the following:

- Access on-premise corporate resources while connected from the corporate intranet.
- Access corporate resources in the cloud while connected directly from the corporate intranet.
- Access on-premise corporate resources while connected from a branch office.
- Access corporate resources in the cloud while connected from a branch office.
- Access on-premise corporate resources from the public internet while teleworking.
- Access corporate resources in the cloud from the public internet while teleworking.

Scenario 2: Employee Access to Internet Resources

An employee is trying to access the public internet to accomplish some tasks. This scenario will show a specific user experience where an employee attempts to access an enterprise-sanctioned, web-based service on the internet by using an enterprise-managed device. Although the web-based service is not owned and managed by the enterprise, the associated access request for that resource will still be enforced, dynamically and in real time, by a ZTA solution implemented in this project. The solution will manage the employee's access, regardless of location. That is, the employee can access the internet while connected inside the corporate intranet, a branch office, or the public internet by using an enterprise-managed device.

If an employee is allowed by corporate policy to access non-enterprise-managed resources and services in the public internet by using enterprise-managed devices, the ZTA solution will allow the enterprise to determine the extent of this access.

Examples of access restrictions in the above paragraph could include:

- Access to social media sites is not sanctioned.

- Access to an internet search engine is permitted, and the associated access request for this resource does not need to be granted in real time through the corporate network when an employee is working at a branch office or while teleworking (e.g., coffee shop or airport).
- Mission-critical services on the public internet (e.g., GitHub) can be accessed directly by the employee.

Scenario 3: Contractor Access to Corporate and Internet Resources

A contractor is trying to access certain corporate resources and the internet. This scenario will show a specific user experience where a contractor attempts to access certain corporate resources and the internet to perform the planned service for the organization. The corporate resources can be on premise or in the cloud, and the contractor will be able to access corporate resources while on premise or from the public internet, using an enterprise-managed device given to the contractor, a contractor-owned and managed device, or a BYOD scenario. The ZTA solution implemented in this project will enforce, dynamically and in near real time, the associated access requests for resources by the contractor.

Scenario 4: Inter-server Communication Within the Enterprise

Corporate services often have different servers communicating with each other. For example, a web server communicates with an application server. The application server communicates with a database to retrieve data back to the web server. This scenario will demonstrate examples of inter-server interactions within the enterprise, which will include servers that are on premise, in the cloud, or between servers that are on premise and in the cloud. The ZTA solution implemented in this project will enforce, dynamically and in near real time, the associated network communications among designated servers that interact with one another.

Scenario 5: Cross-Enterprise Collaboration with Business Partners

Two enterprises (Enterprise A and Enterprise B) may collaborate on a project where resources are shared. In this scenario, the ZTA solution implemented in this project will enable users from one enterprise to securely access specific resources from the other enterprise, and vice versa. For example, Enterprise A users will be able to access a specific application from Enterprise B, while Enterprise B users will be able to access a specific database from Enterprise A.

Scenario 6: Develop Trust Score/Confidence Level with Corporate Resources

Enterprises have monitoring systems, security information and event management (SIEM) systems, and other resources that can provide data to support security analytics to a policy engine to create a more granular trust score/confidence level for access to corporate resources and promote strict access based on the confidence level. In this scenario, a ZTA solution will integrate these monitoring and SIEM systems with the policy engine to produce more precise calculation of trust scores/confidence levels in near real time.

Note: The scenarios above may be created and demonstrated in different phases throughout the project.

3 HIGH-LEVEL ARCHITECTURE

Figure 1 illustrates a high-level, notional architecture of the logical and functional components that could make up a ZTA for a typical IT enterprise.

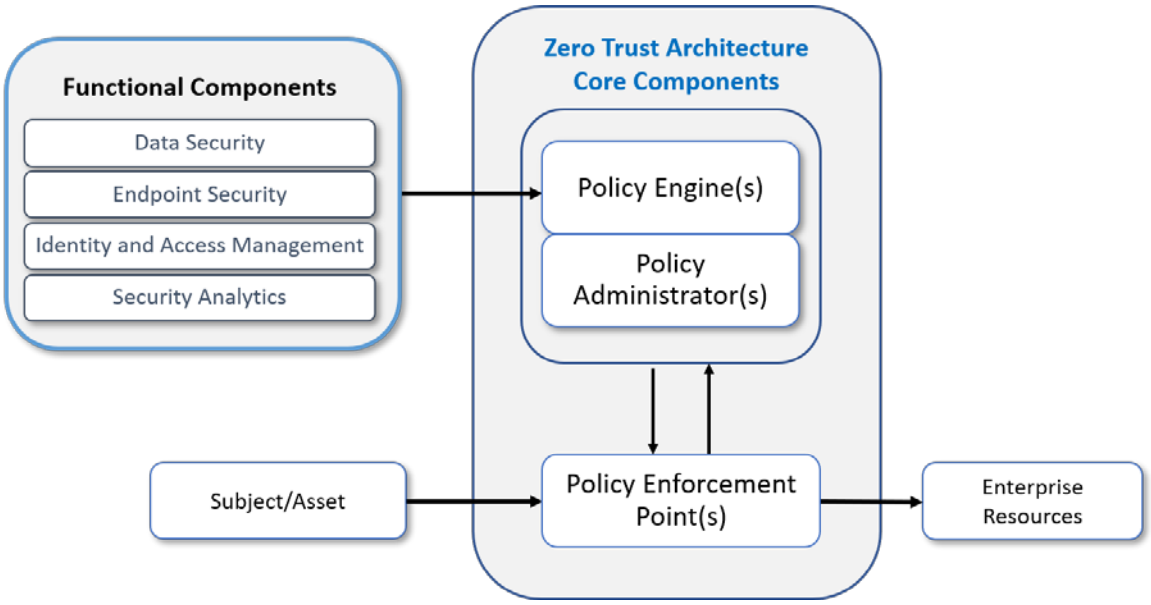

Figure 1. ZTA High-Level Architecture

Component List

The technical components required of the ZTA solution(s) for this project include but are not limited to:

Core Components:

- The policy engine handles the ultimate decision to grant, deny, or revoke access to a resource for a given subject. The policy engine calculates the trust scores/confidence levels and ultimate access decisions.
- The policy administrator is responsible for establishing/terminating the transaction between a subject and a resource. It generates any session-specific authentication and authentication token or credential used by a client to access an enterprise resource. It is closely tied to the policy engine and relies on its decision to ultimately allow or deny a session.
- The policy enforcement point handles enabling, monitoring, and eventually terminating connections between a subject and an enterprise resource.

Functional Components:

- The data security component includes all the data access policies and rules that an enterprise develops to secure its information, and the means to protect data at rest and in transit.
- The endpoint security component encompasses the strategy, technology, and governance to protect endpoints (e.g., servers, desktops, mobile phones, IoT devices) from threats and attacks, as well as protect the enterprise from threats from managed and unmanaged devices.
- The identity and access management component includes the strategy, technology, and governance for creating, storing, and managing enterprise user (i.e., subject) accounts and identity records and their access to enterprise resources.

- The security analytics component encompasses all the threat intelligence feeds and traffic/activity monitoring for an IT enterprise. It gathers security and behavior analytics about the current state of enterprise assets and continuously monitors those assets to actively respond to threats or malicious activity. This information could feed the policy engine to help make dynamic access decisions.

Devices and Network Infrastructure Components:

- Assets include the devices/endpoints, such as laptops, tablets, and other mobile or IoT devices, that connect to the enterprise.
- Enterprise resources include data and compute resources as well as applications/services hosted and managed on premise, in the cloud, at the edge, or some combination of these.
- Network infrastructure components encompass network resources a medium or large enterprise might typically deploy in its environment. It is assumed that the ZTA core and functional components and devices are connected via, or integrated into, the network infrastructure. Note: The network infrastructure is not depicted in Figure 1. The NCCoE will provide these components as part of its internal lab infrastructure.

Desired Security Characteristics and Properties

This project seeks to develop a reference design and implementation, using commercially available technology that meets the following characteristics:

- All interactions throughout the proposed architecture are achieved in the most secure manner available, with emphasis on protecting confidentiality and integrity through a consistent identification, authentication, and authorization scheme.
- All interactions throughout the proposed architecture are continually reassessed with possible reauthentication and reauthorization as necessary to mitigate unauthorized access to enterprise resources.
- Access to an enterprise resource is assessed on a per-session basis and authorized specifically for that enterprise resource.
- Access requests are evaluated dynamically based on organizational policies and rules for accessing enterprise resources, including the observable state of:
    - subject identity (e.g., user account or service identity with associated attributes)
    - requesting asset (e.g., laptop, mobile device, server) device characteristics such as the software version installed, security posture, network location, time/date of request, previously observed behavior, and installed credentials
    - requested resource (e.g., server, application, service) characteristics
- Enterprise assets and resources are continuously monitored and reassessed to maintain them in their most secure states possible.
- Log and event data generated about the current state of enterprise assets, resources, and interactions throughout the proposed architecture are collected and leveraged for better policy alignment and enforcement to increase the enterprise's overall security posture.
- Secure access to corporate resources, hosted either on premise or within a cloud environment, as well as to non-corporate resources on the internet are provided

without the use of conventional network and network perimeter access and security solutions.

- Integration with various directory protocols and identity management services (e.g., Lightweight Directory Access Protocol [LDAP], OAuth 2.0, Active Directory, OpenLDAP, Security Assertion Markup Language) is demonstrated.
- Integration with SIEM tools through common application programming interfaces is demonstrated.
- Desired enterprise device security characteristics are demonstrated, including:
    - maintaining data protection at rest and in transit
    - remediating device vulnerabilities that could result in unauthorized access to data stored on or accessed by the device, and misuse of the device
    - mitigating malware execution on the device that could result in unauthorized access to data stored on or accessed by the device, and misuse of the device
    - mitigating the risk of data loss through accidental, deliberate, or malicious deletion or obfuscation of data stored on the device
    - maintaining awareness of and responding to suspicious or malicious activities within and against the device to prevent or detect a compromise of the device

4 Relevant Standards and Guidance

The references, standards, and guidelines that apply to this project are listed below.

- NIST Cybersecurity Framework v.1.1, Framework for Improving Critical Infrastructure Cybersecurity
  https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf
- NIST SP 800-30 Revision 1, Guide for Conducting Risk Assessments
  https://doi.org/10.6028/NIST.SP.800-30r1
- NIST SP 800-37 Revision 2, Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach For Security and Privacy
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-37r2.pdf
- NIST SP 800-40 Revision 3, Guide to Enterprise Patch Management Technologies
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-40r3.pdf
- NIST SP 800-46 Revision 2, Guide to Enterprise Telework, Remote Access, and Bring Your Own Device (BYOD) Security
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-46r2.pdf
- NIST SP 800-53 Revision 4, Security and Privacy Controls for Federal Information Systems and Organizations
  https://csrc.nist.gov/csrc/media/publications/sp/800-53/rev-4/archive/2013-04-30/documents/sp800-53-rev4-ipd.pdf
- NIST SP 800-57 Part 1 Revision 4, Recommendation for Key Management: Part 1: General
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-57pt1r4.pdf
- NIST SP 800-61 Revision 2, Computer Security Incident Handling Guide
  http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf

- NIST SP 800-63 Revision 3, Digital Identity Guidelines
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-3.pdf
- NIST SP 800-92, Guide to Computer Security Log Management
  http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-92.pdf
- NIST SP 800-114 Revision 1, User's Guide to Telework and Bring Your Own Device (BYOD) Security
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-114r1.pdf
- NIST SP 800-122, Guide to Protecting the Confidentiality of Personally Identifiable Information (PII)
  https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-122.pdf
- NIST SP 800-124 Revision 2 (Draft), Guidelines for Managing the Security of Mobile Devices in the Enterprise
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-124r2-draft.pdf
- NIST SP 800-160 Vol. 2, Developing Cyber Resilient Systems: A Systems Security Engineering Approach
  https://csrc.nist.gov/publications/detail/sp/800-160/vol-2/final
- NIST SP 800-162, Guide to Attribute Based Access Control (ABAC) Definition and Considerations
  https://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.sp.800-162.pdf
- NIST SP 800-175B, Guideline for Using Cryptographic Standards in the Federal Government: Cryptographic Mechanisms
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-175b.pdf
- NIST SP 800-171 Revision 2, Protecting Controlled Unclassified Information in Nonfederal Information Systems and Organizations
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-171r2.pdf
- NIST SP 800-205, Attribute Considerations for Access Control Systems
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-205.pdf
- NIST SP 800-207 (Second Draft), Zero Trust Architecture
  https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207-draft2.pdf
- NIST SP 1800-3, Attribute Based Access Control
  https://www.nccoe.nist.gov/sites/default/files/library/sp1800/abac-nist-sp1800-3-draft-v2.pdf
- Cloud Security Alliance, Software Defined Perimeter Working Group, SDP Specification 1.0
  https://downloads.cloudsecurityalliance.org/initiatives/sdp/SDP_Specification_1.0.pdf
- ISO/IEC 27001, Information Technology–Security Techniques–Information Security Management Systems
- American Council for Technology-Industry Advisory Council, Zero Trust Cybersecurity Current Trends
  https://www.actiac.org/system/files/ACT-IAC%20Zero%20Trust%20Project%20Report%2004182019.pdf
- Federal Information Processing Standards 140-3, Security Requirements for Cryptographic Modules
  https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.140-3.pdf

5 SECURITY CONTROL MAP

This table maps the characteristics of the commercial products that the NCCoE will apply to this cybersecurity challenge to the applicable standards and recommended practices described in the *Framework for Improving Critical Infrastructure Cybersecurity* and to other NIST guidance. This information represents an approach to document the applicability of standards, guidelines, and recommended practices to the security characteristics of the solution, but it does not imply that products and services will meet an industry's requirements for regulatory approval or accreditation.

Table 1: Security Control Map

Cybersecurity Framework v1.1		
Function	**Category**	**Subcategory**
IDENTIFY (ID)	Asset Management (ID.AM)	ID.AM-1: Physical devices and systems within the organization are inventoried.
		ID.AM-2: Software platforms and applications within the organization are inventoried.
		ID.AM-5: Resources (e.g., hardware, devices, data, time, personnel, and software) are prioritized based on their classification, criticality, and business value.
	Risk Assessment (ID.RA)	ID.RA-1: Asset vulnerabilities are identified and documented.
		ID.RA-3: Threats, both internal and external, are identified and documented.
PROTECT (PR)	Identity Management, Authentication, and Access Control (PR.AC)	PR.AC-1 Identities and credentials are issued, managed, verified, revoked, and audited for authorized devices, users, and processes.
		PR.AC-3 Remote access is managed.
		PR.AC-4 Access permissions and authorizations are managed, incorporating the principles of least privilege and separation of duties.
		PR.AC-5: Network integrity is protected (e.g., network segregation, network segmentation).
		PR.AC-6: Identities are proofed and bound to credentials and asserted in interactions.

Cybersecurity Framework v1.1		
Function	**Category**	**Subcategory**
		PR.AC-7 Users, devices, and other assets are authenticated (e.g., single-factor, multifactor) commensurate with the risk of the transaction (e.g., individuals' security and privacy risks and other organizational risks).
	Data Security (PR.DS)	PR.DS-2 Data in transit is protected.
		PR.DS-5: Protections against data leaks are implemented.
		PR.DS-6 Integrity-checking mechanisms are used to verify software, firmware, and information integrity.
		PR.DS-8: Integrity-checking mechanisms are used to verify hardware integrity.
	Information Protection Processes and Procedures (PR.IP)	PR.IP-1: A baseline configuration of IT/industrial control systems is created and maintained, incorporating security principles (e.g., concept of least functionality).
		PR.IP-3: Configuration change control processes are in place.
	Protective Technology (PR.PT)	PR.PT-3 The principle of least functionality is incorporated by configuring systems to provide only essential capabilities.
		PR.PT-4: Communications and control networks are protected.

Cybersecurity Framework v1.1		
Function	**Category**	**Subcategory**
DETECT	Anomalies and Events (DE.AE)	DE.AE-2: Detected events are analyzed to understand attack targets and methods.
		DE.AE-3: Event data are collected and correlated from multiple sources and sensors.
		DE.AE-5: Incident alert thresholds are established.
	Security Continuous Monitoring (DE.CM)	DE.CM-1: The network is monitored to detect potential cybersecurity events.
		DE.CM-2: The physical environment is monitored to detect potential cybersecurity events.
		DE.CM-4: Malicious code is detected.
		DE.CM-5: Unauthorized mobile code is detected.
		DE.CM-6: External service provider activity is monitored to detect potential cybersecurity events.
		DE.CM-7 Monitoring for unauthorized personnel, connections, devices, and software is performed.
		DE.CM-8: Vulnerability scans are performed.
	Detection Processes (DE.DP)	DE.DP-5: Detection processes are continuously improved.
RESPOND	Mitigation (RS.MI)	RS.MI-1: Incidents are contained.
		RS.MI-2: Incidents are mitigated.

APPENDIX A REFERENCES

[1] S. Rose et al., *Zero Trust Architecture,* National Institute of Standards and Technology (NIST) Special Publication 800-207, Gaithersburg, Md., August 2020. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

[2] J. Chaffettz, *Adopting a zero trust cyber model in government*, Federal News Network, September 19, 2016. Available: https://federalnewsnetwork.com/commentary/2016/09/adopting-zero-trust-cyber-model-government/

[3] The Federal Chief Information Officer Council: https://www.cio.gov/