NERC
NORTH AMERICAN ELECTRIC
RELIABILITY CORPORATION

White Paper

Zero Trust Security for Electric Operations Technology
June 2023

Executive Summary

Zero trust (ZT) offers the electric industry a clear direction forward for continual improvement to securing our critical infrastructure against emerging threats to operations technology (OT)—including ransomware and the proliferation of industrial control system malware tools, such as Pipedream. ZT is a collection of concepts intended to drive least privilege further, building upon and enhancing historical controls and perimeter-based security models rather than tearing them down. Industry needs to continue to develop equipment and software as well as people, processes, policies, and governance capable of delivering on ZT principles. Entities should invest in staff training for ZT, develop OT security programs, design roadmaps based on a ZT maturity model for the development of ZT architecture (ZTA) at the right pace for their organization. Additionally, using a thoughtful implementation process will allow organizations to incorporate ZTA incrementally and should be done in collaboration with OT integrators and vendors. OT networks and legacy devices may create constraints that require hybrid approaches to solve. No single product or tool on the marketplace provides a complete ZTA, and organizations may already have infrastructure and controls in place that qualify as components of a ZTA. Finally, entities are encouraged to stage rollouts of ZTA starting with information technology (IT) networks and demilitarized zones (DMZ) to build familiarity with the complexities, challenges, and impacts of the controls and technology before implementing in the OT space.

Introduction

The purpose of this white paper is to inform the electricity sector about ZT concepts and to provide considerations and recommendations regarding the adoption of ZT controls in OT and industrial control system (ICS) environments. This paper describes some of the key differences between OT and IT environments; however, this paper's focus is specifically on the implementation considerations in the OT environment. This paper also leverages the concepts of ZT maturity models for varying levels of implementation by registered entities. Lastly, this paper describes considerations regarding ZT adoption by registered entities and the NERC Critical Infrastructure Protection (CIP) standards. One threat today that has driven the need for ZT principles is ransomware. When ransomware compromises one device inside a typical network perimeter, it then uses the inherent trust of network peers to spread to other devices. In the case of enterprise networks, this may compromise dozens, hundreds, or tens of thousands of network peers through the inherent trust inside the perimeter. Another example is known as "pivoting," where an attacker may gain access to one device through a legitimate communication allowed through a network perimeter but then launches attacks from that device and compromises other more critical systems within the perimeter. ZT aims to strengthen security with controls better able to detect, mitigate, or prevent such threats.

What is ZT?

Computer networks have been traditionally designed to follow a “bastion” model wherein strong, multilayered defenses are utilized to mitigate intrusion. Defenses inside the bastion are typically far less robust, and the average process or user can traverse a network, system, or application to access those resources they desire once admitted. Internal security controls and monitoring are potentially less robust and assume that a running process, service, or authenticated user is “trusted” so their actions receive less scrutiny than they do at the network boundary. ZT principles are designed to mitigate these types of “inherent trust” issues. The concept of ZT shifts cyber security control design philosophy from the old adage of “trust, but verify” to “never trust, constantly verify.” No device gains inherent trust from its network location even if it is a local network peer. In a ZTA, no user or device is implicitly trusted and undergoes access and authorization tests continually. There is no default visibility or access on networks; any and all access must be enabled via policy. Policy, not topology, governs visibility and access between devices. Additionally, ZT principles drive access decisions to be much more granular, granting access to individual resources, services, or limited network access with potentially different tests for each access request and ongoing tests to maintain access. As ZTAs mature, access can be granted based on user ID, method of authentication, state of the user’s device, protocol security level, and numerous other variables to build deeper, fuller trust that the access being granted is currently authorized. ZTA aims to fill internal control gaps and reinforce perimeters with additional or more effective controls rather than tear them down.

ZT as Defined by the National Institute of Standards and Technology

The National Institute of Standards and Technology (NIST) defines ZT as a collection of concepts and ideas designed to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services in the face of a network viewed as compromised.[1] The basic premise of ZT is that there is no implicit trust granted to user or systems based on their physical or network location because there is no trust of any network, user, or device.[2]

NIST further defines ZTA as an enterprise cyber security plan that utilizes ZT concepts and encompasses component relationships, workflow planning, and access policies. Therefore, a ZT enterprise is the network infrastructure (physical and virtual) and operational policies that are in place for an enterprise as a product of a ZTA plan.[3]

Tenets of ZT from NIST

ZTA is designed and deployed with adherence to the following ZT basic tenets:

- ***All data sources and computing services are considered resources.*** A network may be composed of multiple classes of devices. A network may also include small footprint devices, such as sensors and collectors, that send data to aggregators/storage, software-as-a-service applications, and more. Also, an enterprise may decide to classify personally owned devices as resources if they can access enterprise-owned resources.
- **All communication is secured regardless of network location.** Network location alone does not imply trust. Access requests originating from trusted devices and trusted network locations (e.g., located on enterprise-owned network infrastructure) should be held to the same minimum level of

[1] https://csrc.nist.gov/glossary/term/zero_trust

[2] https://www.nerc.com/pa/Stand/Project_201602_Modifications_to_CIP_Standards_RF/2016-02_CIP-005_and_Zero_Trust_Webinar_Slides_02192020.pdf

[3] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

security scrutiny as requests originating from anywhere else, including untrusted devices, public internet, partner WANs, etc. In other words, trust should not be automatically granted based on the device being on an enterprise network infrastructure. All communication is done in a secure manner, protecting confidentiality and integrity and providing source authentication.

- **Access to individual enterprise resources is granted on a per-session basis.** Trust in the requester is evaluated before the access is granted; access should also be granted with the least privileges needed to complete the task. Additionally, authentication and authorization to one resource will not automatically grant access to a different resource.
- **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**. An organization protects resources by defining what resources it has, who its members are (or the ability to authenticate users from a federated community), and what access to resources those members need. For ZT, client identity can include the user account (or service identity) and any associated attributes assigned by the enterprise to that account or artifacts to authenticate automated tasks. Requesting asset state can include device characteristics like software versions installed, network location, time/date of request, previously observed behavior, and installed credentials. Behavioral attributes include, but are not limited to, automated subject analytics, device analytics, and measured deviations from observed usage patterns. Policy is the set of access rules based on attributes that an organization assigns to a subject, data asset, or application. Environmental attributes may include factors like requestor network location, time, reported active attacks, etc. These rules and attributes are based on the needs of the business process and acceptable level of risk. Resource access and action permission policies can vary based on the sensitivity of the resource/data. Least-privilege principles are applied to restrict both visibility and accessibility.
- **The enterprise monitors and measures the integrity and security posture of all owned and associated assets.** No asset is inherently trusted. The enterprise evaluates the security posture of the asset when evaluating a resource request. An enterprise implementing ZTA should establish continuous diagnostics and mitigation or a similar system to monitor the state of devices and applications and should apply patches/fixes as needed. Assets that are discovered to be subverted, have known vulnerabilities, and/or are not managed by the enterprise may be treated differently (including denial of all connections to enterprise resources) than devices owned by or associated with the enterprise that are deemed to be in their most secure state. This may also apply to associated devices (e.g., personally owned devices) that may be allowed to access some resources but not others; this also requires a robust monitoring and reporting system in place to provide actionable data about the current state of enterprise resources.
- **Resource authentication and authorization are dynamic and strictly enforced before access is allowed.** This is a constant cycle of obtaining access, scanning and assessing threats, adapting, and continually re-evaluating trust in ongoing communication. An enterprise implementing a ZTA would need to have identity, credential, and access management as well as asset management systems in place. This includes the use of multifactor authentication for access to some or all enterprise resources. Continual monitoring, with possible re-authentication and reauthorization, occurs throughout user transactions as defined and enforced by policy, (e.g., time-based, new resource

requested, resource modification, anomalous subject activity detected) that strives to achieve a balance of security, availability, usability, and cost-efficiency.

- **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.** An enterprise should collect data about asset security posture, network traffic, and access requests; process that data; and use any insight gained to improve policy creation and enforcement. This data can also be used to provide context for access requests.

To summarize, an authenticated user or process accesses data through the intermediation of applications in a ZT world. Network based access, such as that conferred by virtual private networks, is avoided while still potentially used, is no longer relied on as the sole source of authentication and encryption; instead, further bolstering those same protections that are already necessarily included in the ZT system and software. Identity management and access controls (e.g., conditional, role based) are enforced at the application. The world of ZT thus resembles modern mobile applications, and many of the services in use today use this model, such as Office 365. ZT embeds comprehensive security monitoring; granular, dynamic, and risk-based access controls; and system security automation in a coordinated manner throughout all aspects of the infrastructure in order to focus specifically on protecting critical assets (data) in real-time within a dynamic threat environment. This data-centric security model allows the concept of least privileged access to be applied for every access decision where the answers to the questions of who, what, when, where, and how are critical for appropriately allowing or denying access to resources.

ZT Maturity Model

ZT should be considered a forward-thinking strategy for control design. However, while real-time, trustless, and policy-based algorithmically-driven access decisions across an organization’s technology footprint are a pinnacle to be strived for, it is best to realize a sliding scale maturity model from implicit trust controls to less trust and then to trustless. Across all industries, ZT is a paradigm shift that is accompanied by a roadmap to gradually implement least-trust controls. It is important to understand that available technology solutions can provide incremental steps forward on an entity’s ZT maturity roadmap, but no single product, tool, or policy will achieve a complete ZTA. Furthermore, vendors may or may not advertise their products as ZT, as ZT is more a collection of concepts for designing controls that more effectively enforce least privilege security. Carefully planning a transition into ZT controls may allow an entity to manage the risks of difficult design challenges, including choosing on-device deployment (endpoint protection) of ZT controls versus stand-alone solutions like internal network security monitoring solutions. Additionally, granting the time and resources to rollout endpoint protection systems carefully and gradually is paramount to avoiding potentially dangerous operational impacts.

When examining the application of ZT controls in an OT environment, entities may find that not all systems or networks are viable for implementation of all ZT controls. Products and tools designed for OT may lean towards improving detective capabilities over prevention to enable compatibility with sensitive OT requirements and legacy assets. However, limitations can still arise, and hybrid design approaches may represent an optimal solution.

The Cybersecurity and Infrastructure Security Agency (CISA) has defined a ZT maturity model with five distinct pillars (see **Figure 1**). Each pillar may be advanced independently, but an organization is likely to see cross-pillar interoperability and dependencies that require process and technology coordination as they reach advanced maturity.

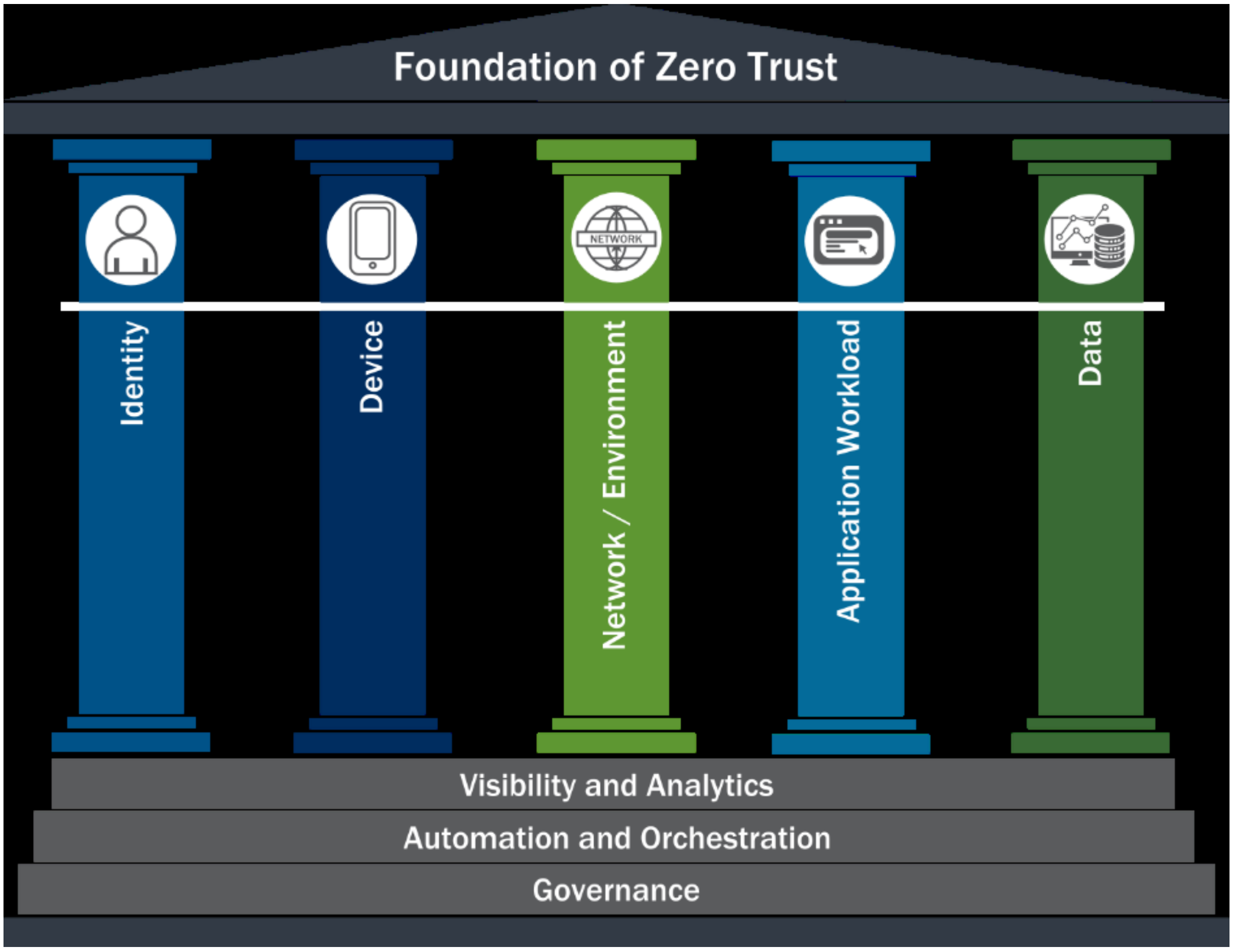

Figure 1: CISA's Foundation for ZT

CISA's maturity model shown in **Figure 2** further develops these pillars across three levels of maturity: traditional, advanced, and optimal.

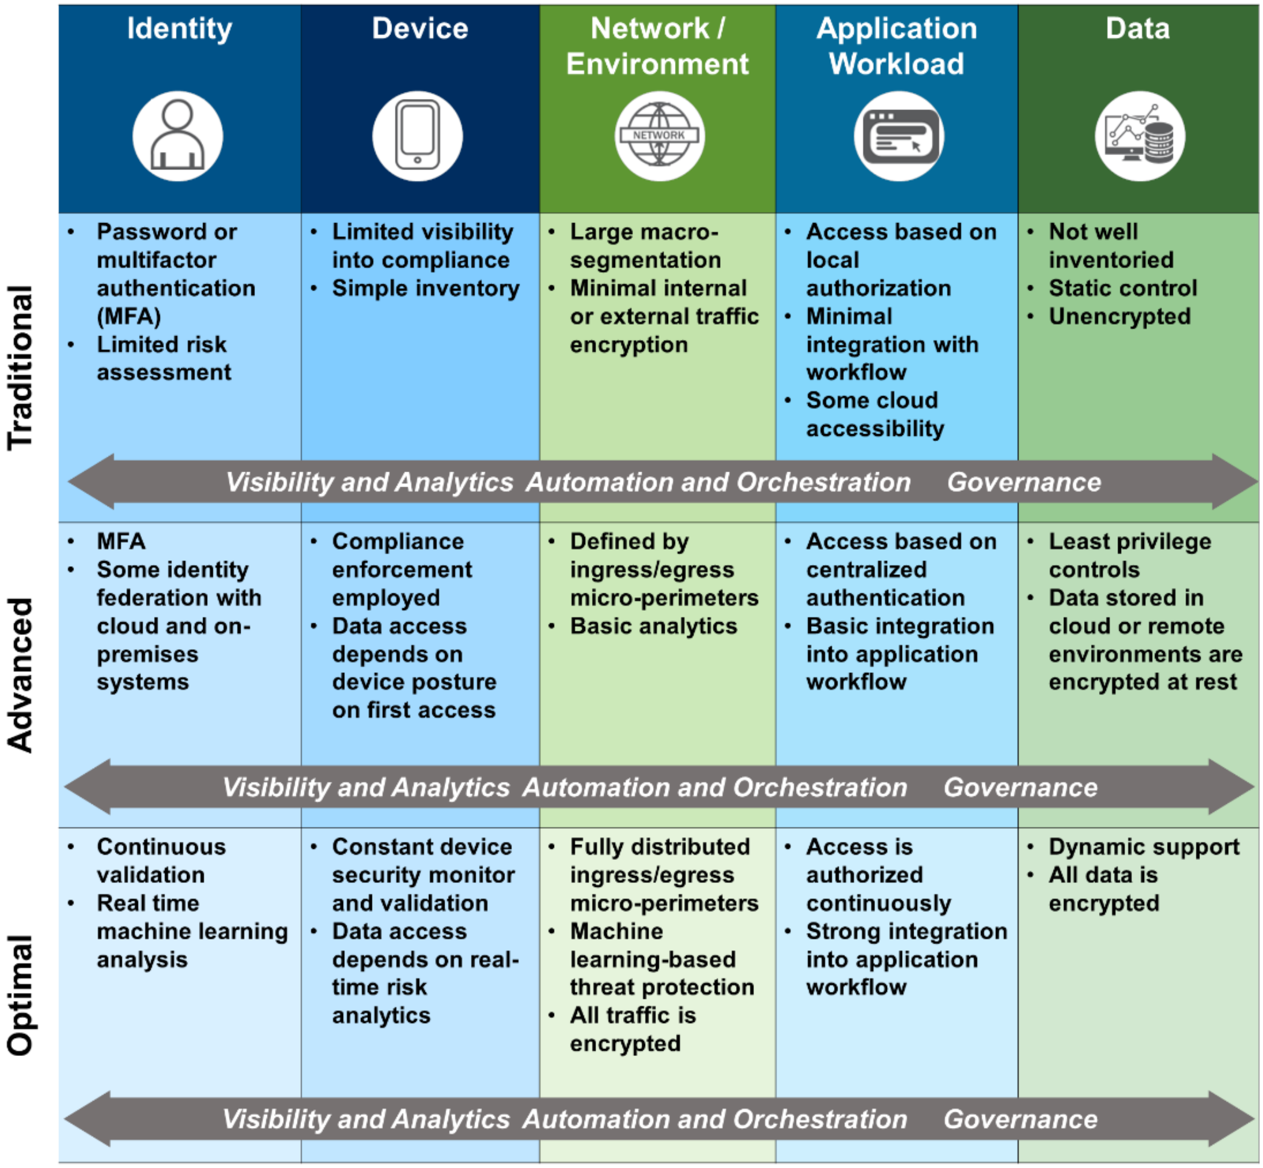

	Identity	Device	Network / Environment	Application Workload	Data
Traditional	• Password or multifactor authentication (MFA) • Limited risk assessment	• Limited visibility into compliance • Simple inventory	• Large macro-segmentation • Minimal internal or external traffic encryption	• Access based on local authorization • Minimal integration with workflow • Some cloud accessibility	• Not well inventoried • Static control • Unencrypted
	Visibility and Analytics		*Automation and Orchestration*		*Governance*
Advanced	• MFA • Some identity federation with cloud and on-premises systems	• Compliance enforcement employed • Data access depends on device posture on first access	• Defined by ingress/egress micro-perimeters • Basic analytics	• Access based on centralized authentication • Basic integration into application workflow	• Least privilege controls • Data stored in cloud or remote environments are encrypted at rest
	Visibility and Analytics		*Automation and Orchestration*		*Governance*
Optimal	• Continuous validation • Real time machine learning analysis	• Constant device security monitor and validation • Data access depends on real-time risk analytics	• Fully distributed ingress/egress micro-perimeters • Machine learning-based threat protection • All traffic is encrypted	• Access is authorized continuously • Strong integration into application workflow	• Dynamic support • All data is encrypted
	Visibility and Analytics		*Automation and Orchestration*		*Governance*

Figure 2: CISA's ZT maturity model

Application of Principles

Trustless Example

What does it mean to be trustless? Consider an example of a system operator logging into a supervisory control and data acquisition (SCADA) human-machine interface or workstation and opening their application client. First, it is taken for granted that the workstation has network access; however, under a ZTA implementation, trustless network access controls would replace this assumed or implied workstation network access authorization. A ZT policy decision engine would perform an algorithmic evaluation of a number of risk factors, such as the workstation's current security patch levels, completed anti-malware scans status, mac address validation, security certificate validation, and/or access authorization to the specific network subnet or virtual local area network (VLAN), such as the control center operator VLAN. Passing all checks results in the workstation being granted network access, but failing one or more tests could result in a quarantining action whereby the network connection is reassigned to a remediation VLAN that limits connectivity to only what is necessary for the system to communicate with patching, anti-virus, and other system management servers. Additionally, it is important to understand that in a truly trustless design, the access decision is one that is continually re-evaluated over time. This ensures adherence to security polices and doesn't allow perpetual access based off prior access decision outcomes.

Similarly, when an operator attempts to authenticate into their SCADA client, additional ZT policies are evaluated against the policy engine: Does the operator have the proper roles or group membership assignments necessary to be authorized for the SCADA application? Does this access request fall within normal operating hours as defined within the policy? Does this login match with the user's previous system usage behavior?

These factors are combined with the source system (the workstation) evaluations previously mentioned. What is the real-time security risk state of the organization at this time? For example, has malware recently been detected? All of these real-time and dynamic evaluations determine if access is granted and to what extent. Dependent on the measured and evaluated risk of the request, access could be denied, granted, or granted with lesser privileges until remediation is achieved. Alternatively, the authentication could be elevated to a multi-factor authentication prompt to address elevated risk.

The example controls given above may seem excessive, or they may be perceived as creating undo operational risk. These are fair concerns and they emphasize the need for utilities to approach ZT with their own roadmap to maturity in collaboration with their OT vendor(s). Controls at the upper end of ZT maturity, especially preventative controls, come with an equal cost of technical complexity and administrative burden due to the system and communication knowledge required to design and upkeep ZT policies that will not cause operational disruptions.

ZT in OT/ICS Environment

In OT/ICS environments, it is important to consider ZT in terms of securing and protecting critical processes, not just data. In other words, when implementing ZT in OT/ICS one must not only consider access and authorization to the data hosted by a data source but managing access to the data source device itself.

Within many (most) OT/ICS environments, a distinction has to be made between the different subsystem capability and functional (Purdue model[4]) levels when considering the implementation of ZT (see **Figure 3**).

[4] Purdue Enterprise Reference Architecture - Wikipedia

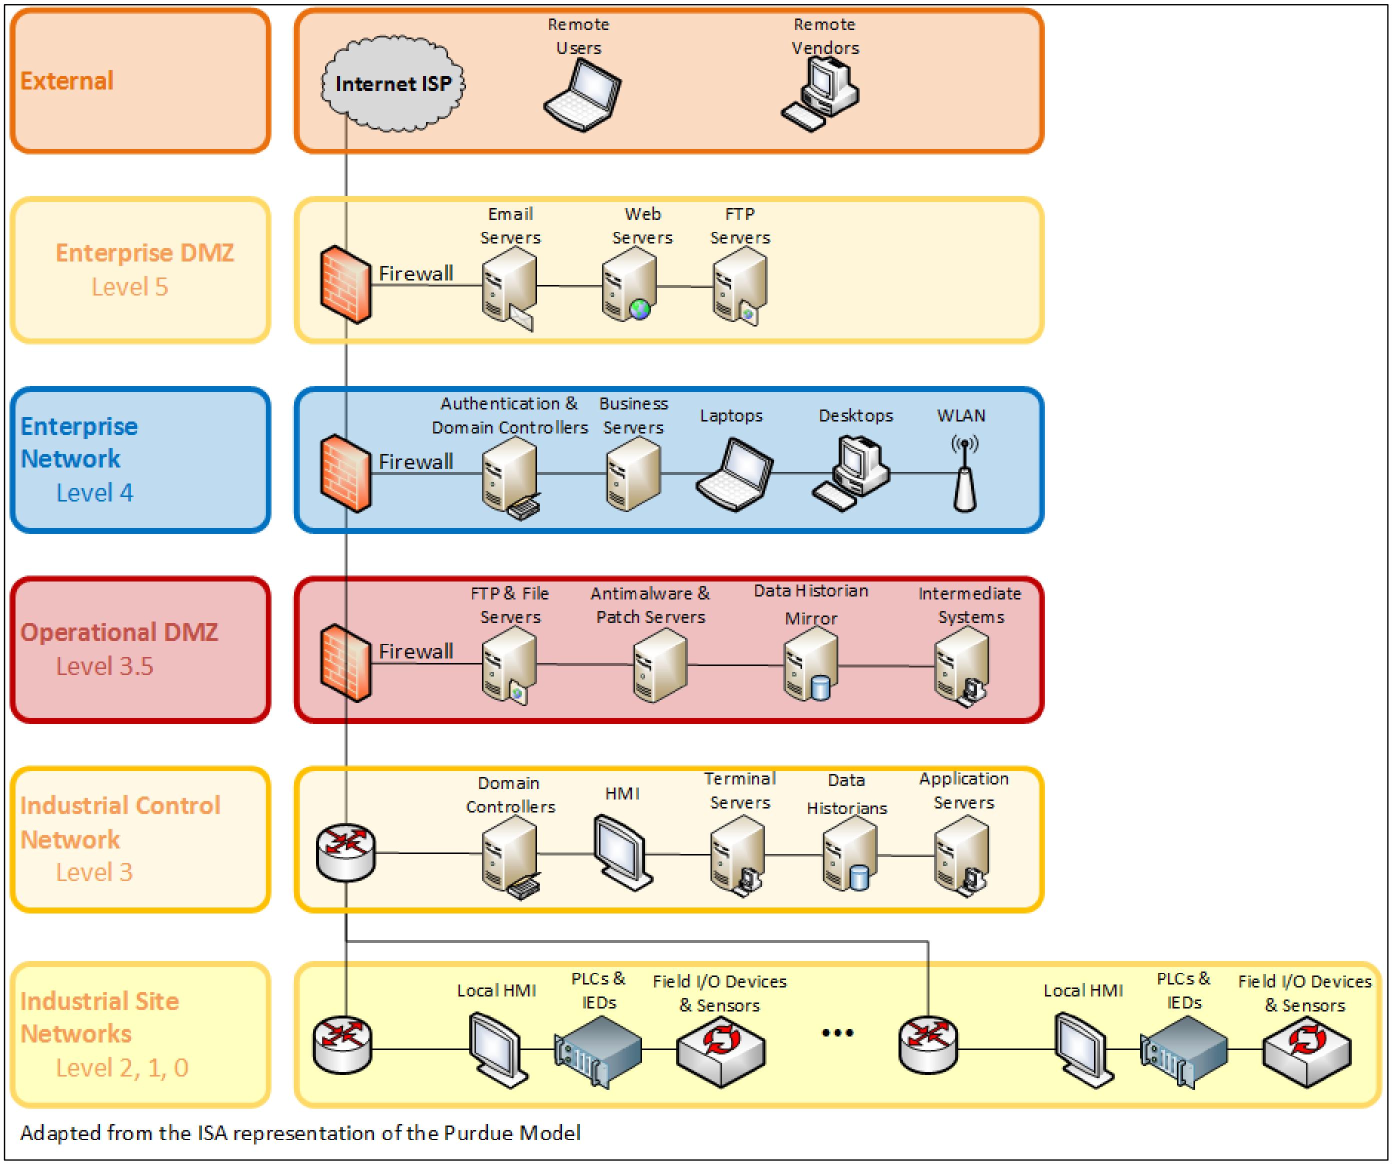

Adapted from the ISA representation of the Purdue Model

Figure 3: ICS Network by Purdue Model

Lower level (Purdue level 0–1) ICS systems and devices (e.g., IEDs, PLCs, sensors) lack the capability of granular access controls on the device itself and instead rely on perimeter, gateway, and/or front end systems to implement ZT controls.

However, many OT support and control systems (Purdue Level 3 and up), such as historians, human-machine interfaces, PACS, and EACMS, are built on platforms that allow for on-device deployment of ZT controls through granular access management via built-in capabilities or through the use of add-ons, such as endpoint security applications.

The Purdue model is useful for understanding concepts like network segmentation and grouping devices based on function and/or criticality. However, since ZT does not base authorization and trust on the physical or network location of devices and systems, an alternate approach to the Purdue model needs to be considered for devices and systems that are not capable of deploying ZT access controls. The ISA/IEC 62443 model of security zones and conduits (see Figure 4) offers a more granular approach for identifying appropriate and applicable ZT controls that can be implemented within an ICS/OT environment.

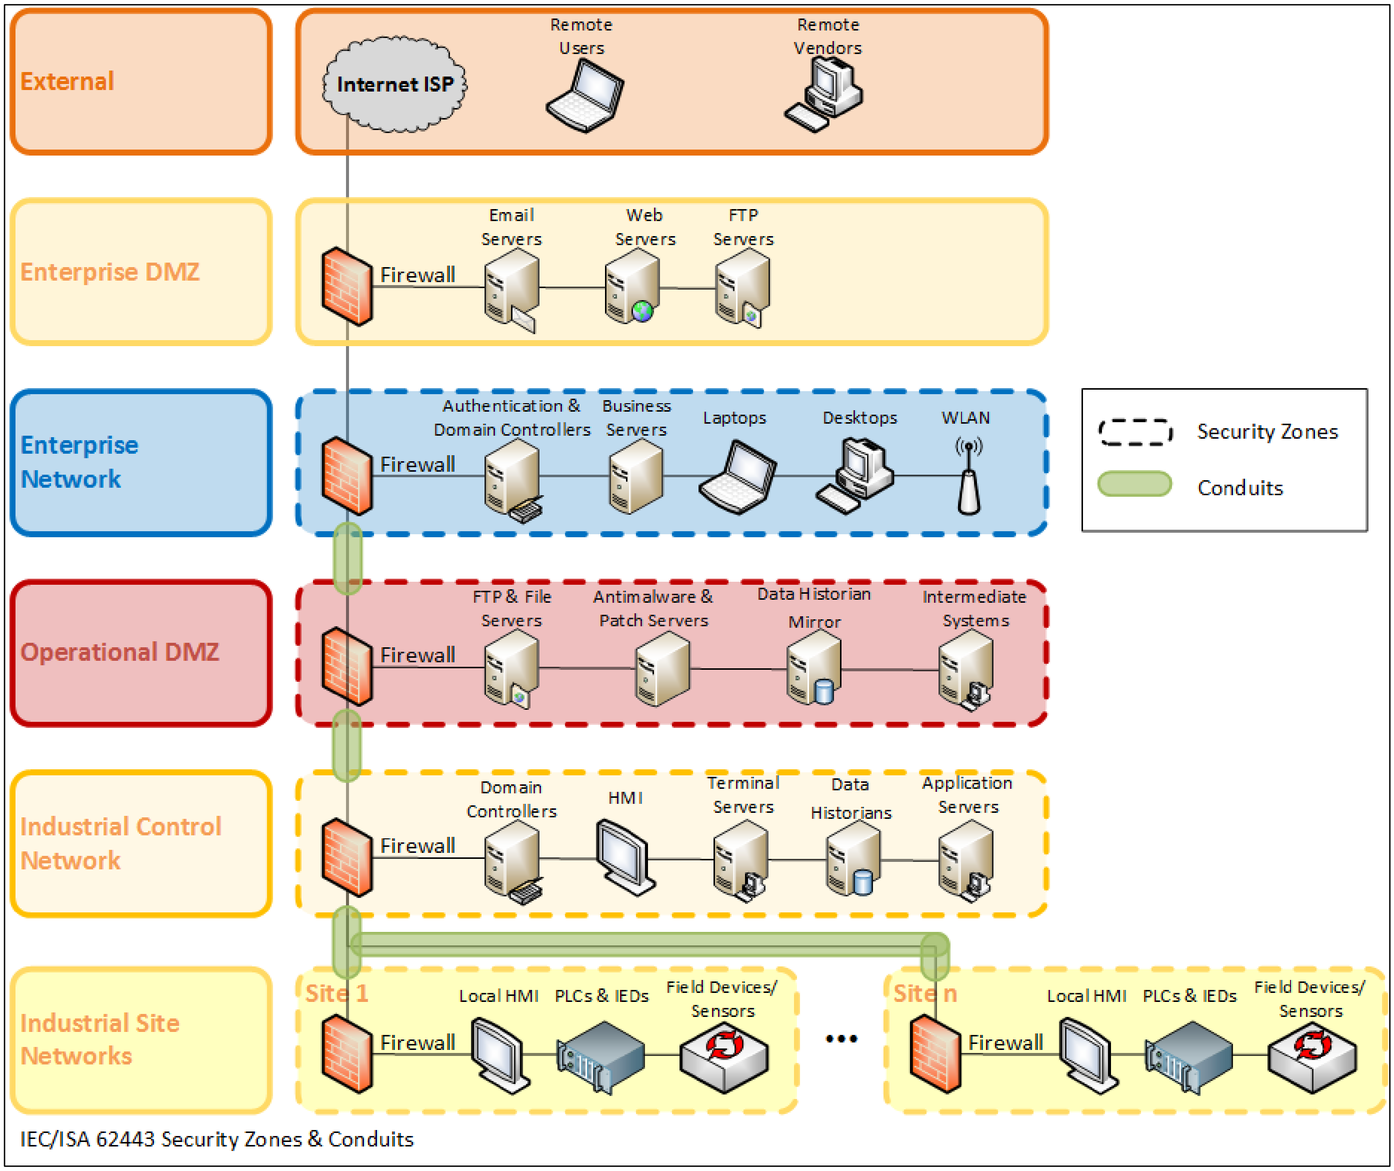

Figure 4: ICS Network by IEC/ISA 62443 Model

Grouping systems and devices that are on the same level within the Purdue model into different security zones allows for establishing ZT access controls even between peer systems on the same level, thus achieving a hybrid design. Security zones can be defined by facility, location, or subsystem within a facility. For example, a utility can define each substation as a single security zone or create separate zones within each substation for an approach that parallels establishing NERC CIP electronic security perimeters (see Figure 5).

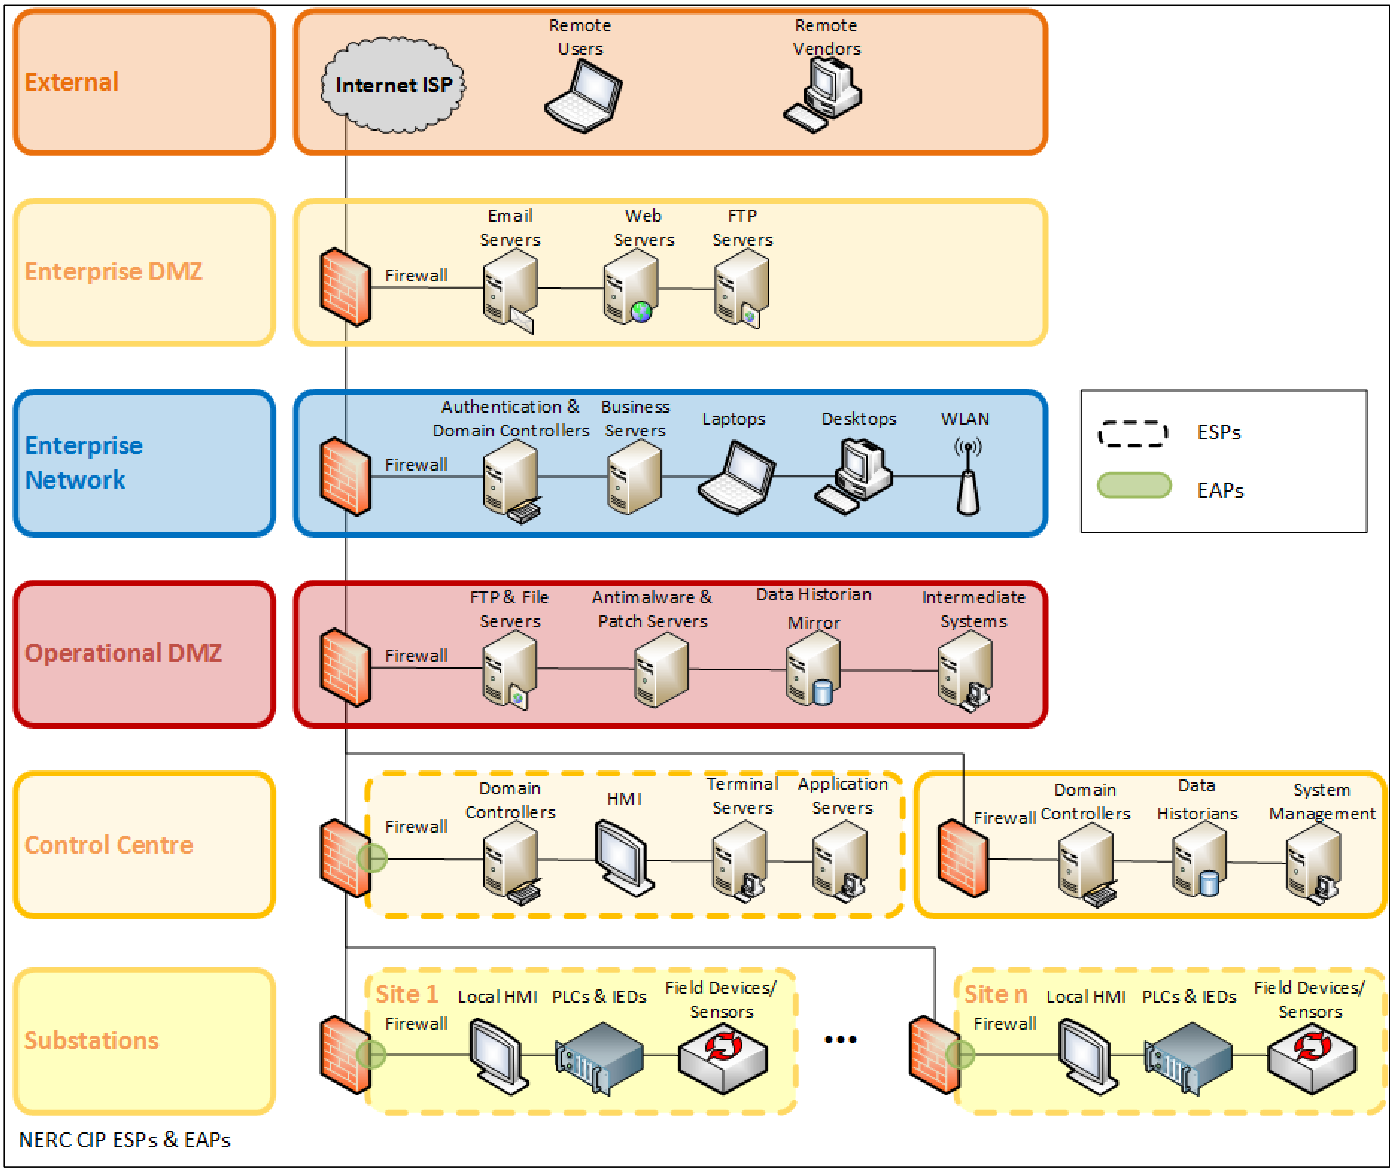

Figure 5: ICS Network by NERC CIP Electronic Security Perimeters

This hybrid approach to ZT can be implemented at a zone level in the areas where the devices within that zone are not capable of implementing host-based security controls (e.g., a substation security zone firewall filtering external inbound/outbound traffic at an application-level) and more granular controls in security zones where the devices are capable of host-based security controls (e.g., a server in the ICS demilitarized zones (DMZ) that filters all connections to the services it hosts). See Figure 6.

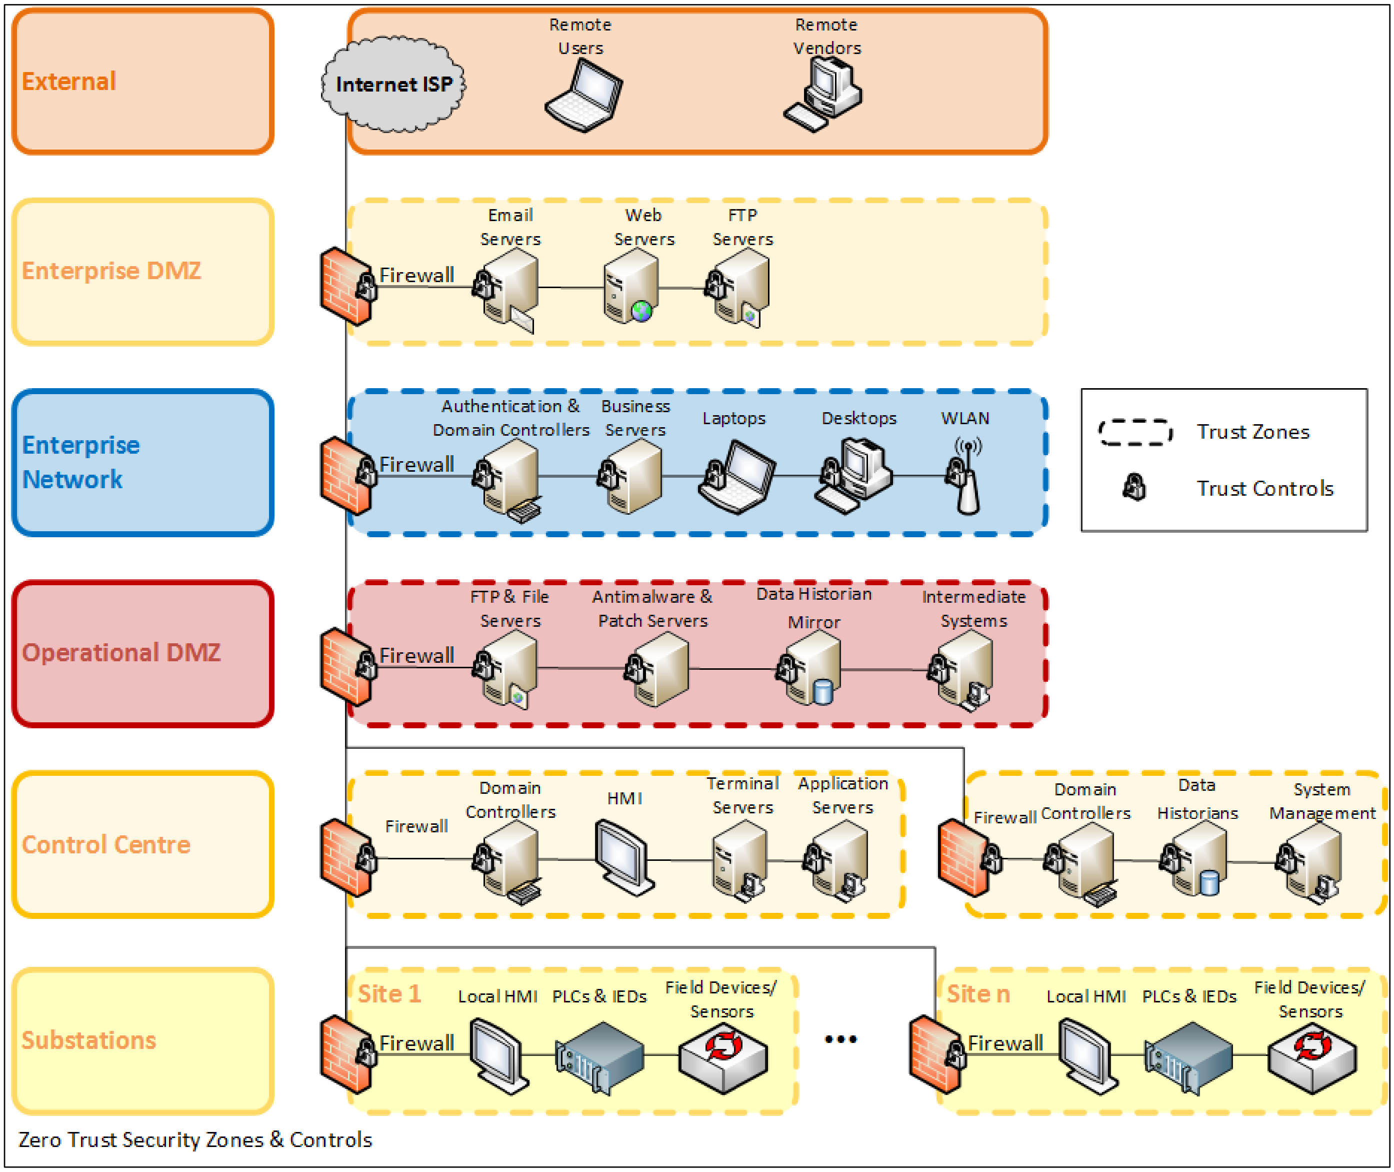

Figure 6: ICS Network by Trust Zones

In light of these mixed capability environments and other factors, there is no “one size fits all” approach that can implement ZT across an organization’s entire OT/ICS environment. Rather, the components of ZTA need to be separated and applied where they are capable of being deployed. The ability and extent to which ZTA components can be deployed must be assessed on a site, facility, and subsystem basis. However, assets planned for the future are alleviated of some of these constraints, and ZTA should be part of the design and planning phases moving forward.

Another important consideration when implementing ZT is exception and failure handling. In most IT environments, it is safe and appropriate to block access (fail-closed) when authentication and authorization cannot be definitively established. In an OT/ICS environment, ZT cannot be deployed in a fail-closed implementation for every subsystem or resource. There are always critical subsystems and resources that have to fail-open and be able to communicate and coordinate, otherwise the system itself may fail or cease to function and/or lead to cyber-physical impacts, including the potential for loss of life.

Benefits and Challenges for OT/ICS

While the need to secure the bulk power system's critical infrastructure is greater than ever before, and the paradigm shift to ZT is the obvious direction for the future of cyber security, there is still a clear need to approach the use case of ZTA for electric OT systems with caution and careful analysis. Additionally, smaller utilities must be wary of advancing too quickly into the cutting edge and taking on too great an administrative or technological burden without appropriately evolving their governance processes and support staff to achieve ZT maturity. The practical use-case for electric OT is explored through benefits, challenges, and recommendations.

Benefits

- Reduced threat surface and associated risk reduction due to increased monitoring controls
- Maximized use of authentication
- Increased visibility tools for security operations into all resource activity (e.g., users, processes, services)
- The ability to dynamically provide access based on real-time assessments
- Reduce an attacker's ability to move laterally and improve the capability to detect lateral movement
- Reduce or improve detection of data exfiltration
- Protection against both internal and external threats
- Improved overall security posture
- Potential for earlier detection of threats providing for quicker response and mitigation
- Tools provide granular details assisting with compliance requirements (e.g., patch management, ports and services, identity management, configuration management, incident response)
- Automated and dynamic processes and controls to quickly adapt to changing threats
- Improves the probability of early detection for malicious or unauthorized access

Challenges

- Early adopters may face challenges with the lack of standardization and lose the benefit of lessons learned from other early adopters.
- High-end ZT implements preventative controls alongside detective controls. Improperly implemented preventative controls could pose additional operational risks for electric OT systems. Comprehensive testing is required including scenarios for handling false positives.
- Diversity of vendor technology offerings may lead to an incomplete approach that may create control gaps.
- Integrating ZT with legacy devices is challenging due to incompatibility with many control solutions using agent/server implementations, SSL certificates, or secure protocols.
- Integrating ZT in networks with extremely low latency requirements (e.g., teleprotection communication) is challenging due to the constraints on viable ZT security solutions.

- There are increased administrative and technical burdens due to the overall complexities associated with ZTA.
- ZT models rely on strictly defined permissions within policies. People, roles, locations, and hardware assets may change, so ZT policies require upkeep and maintenance to be effective.
- At the higher ends of ZTA, deep knowledge of the system, OT devices, and all OT communications is necessary to configure policy allowance of data flows.

Recommendations

- Invest in staff training to better understand ZT concepts, technology offerings, and implementation risks
- Develop or improve cyber security governance when addressing remediation efforts identified in cyber vulnerability assessments, risk assessments, incident response activities, or other internal assessments by taking advantage of these opportunities to advance organizational ZT maturity
- Prioritize establishing an OT cyber security program or improving existing programs
- Perform a cyber-risk assessment of electric OT systems
- Take an asset inventory or validate existing inventories
- Perform a comprehensive controls assessment of the electric OT systems to identify ZT control improvement opportunities.
- Develop a ZT roadmap in order to transition to a ZTA (utilize existing models like CISA's ZT maturity model):
    - Within the roadmap, define maturity transition steps for OT (independently from IT if necessary)
    - Prioritize specific controls that build a foundation to move to ZTA in a hybrid manner (e.g., micro-segmentation, SDN, identity management, MFA)

Implementation Effort

For OT/ICS environments, implementing ZT is an evolutionary process that requires coordination between multiple business units and disciplines. Vertically integrated utilities have multiple groups that hold responsibility for different areas and components of their OT/ICS environment, such as field operations, substations, control centers, engineering, and IT and security. Regardless of the organizational model, leadership buy-in and direction is critical to these undertakings.

Making changes to site network infrastructure as well as access management processes and controls may likely only be feasible when a facility is new, is undergoing major upgrades, or during large scheduled maintenance outages and must be carefully planned, deployed, and validated to help ensure no negative impact to operations. Because of this, it will likely take years with careful planning and full support from all operational areas and leadership to implement ZT in stages across an organization's entire OT/ICS environment. However, some organizations may find that legacy systems and facilities may not be feasibly updateable to ZTA. These entities will need to account for any residual risks from such facilities if they deem ZT controls are necessary for risk mitigation.

For many organizations, the first steps in staging and applying ZT in OT/ICSs will follow after implementing ZT in their IT environments and then most probably targeting the IT/OT DMZs and operational control centers. These areas typically utilize more modern and flexible digital platforms with multipurpose commodity operating system-based servers and workstations as well as advanced network infrastructures that undergo more frequent refresh cycles and upgrades. These comparatively abbreviated refresh cycles allow for more opportunities to move toward ZTA and are more likely to provide full-scale redundancy to mitigate unforeseen negative impacts of a staged ZTA rollout. Deploying ZT within the IT/OT DMZs and control centers also provides the best cost/benefit return as it addresses a majority of the concerning attack surface and threat landscape and has the potential for the least impact to the system in the event the controls block access for a service and/or user.

As previously discussed, due to the large range of OT/ICS system device capabilities and associated limitations, it may not be realistic to consider a device level approach for implementing ZT beyond the DMZs and control centers. Instead, an organization may need to consider hybrid approaches to bring ZT benefits into OT networks.

However, advances in technology and enhanced industry needs in the face of evolving and sophisticated cyber threats means OT equipment manufacturers are increasingly offering more robust cyber security capabilities in their product lines that will help facilitate industry wide movement to systems designed with and capable of ZTA. These more modern systems include authentication and cryptographic mechanisms (among other things) and are less of a technical challenge to implement. Entities should consider these newer technology offerings when creating their ZT roadmaps.

Compliance Consideration

General Approach

Entities must maintain their current compliance programs and responsibilities regardless of adopting any new cyber security controls and associated architectures, such as ZT. Consideration of purposed implementations must be evaluated against applicable CIP standards as would any change to the environments subject to CIP jurisdiction.

Identity and Resource Management

With ZT controls, processing an authorization request for system access may be enhanced to include additional evaluation criteria, such as device security posture, time-based behavioral data, and current organizational risk. But these are in addition to system privileges or permissions pre-mapped to roles, groups, or other identity criteria of accounts that form a strong basis of any authorization control. An organization’s processes that govern baseline system privileges or permissions are designed with business justification, and they follow role-based access control (RBAC) per best practices and are likely best suited as evidence for the control objectives associated to electronic access authorization.

Authentication

Some ZTAs may utilize a service gateway to intercept all incoming requests for resource (system or data) access and present the point of authentication at that gateway. Other solutions may include certificate management, including for public key infrastructure and reliance on external root authority servers. In the design of these ZT solutions, the same CIP compliance considerations must be given to the new technological components as is given to existing applicable systems, such as electronic access control and monitoring systems and BES cyber assets.

Software Defined Networking

Organizations should carefully consider how to align dynamic and policy-based software defined networking **(**SDN) with CIP's use of logical network access and defined ESP's. When employing ZT policies with SDN, both the network location of individual systems and allowed inbound or outbound communication at network boundaries access control lists can be dynamic. The policy rules established at the SDN controller may offer criteria to redirect or disable communication (causing dynamic update to ACL's) as well as relocate or quarantine systems (causing VLAN change). However, there is often a resulting "base state" of configuration for the network through these policies and then a "deviation state" due a higher state of security or reliability need. As a single system example, SDN policy may result in an assignment of a virtual desktop to a particular VLAN. It may then deviate from that assignment when an additional policy-based evaluation identifies that the workstation is missing recent security patches. This moves the system out of its base state assigned VLAN within an electronic security perimeter and into a deviation state quarantine VLAN (likely a DMZ, potentially outside the ESP) that only allows the necessary ACL-restricted communication to serve security patching services. To aid in evidencing requirements for ESP's and inbound/outbound communication, it is recommended to orient written control processes to maintain CIP compliance around the use of SDN's policies by clearly explaining base state versus deviation states. Such efforts require collaboration and input from subject matter experts within the organization, including compliance, security, and network engineers.

ZT Controls Guidance

Among the types of controls making up ZTA, some are more suited than others for deployment across electric OT and IT-centric systems. Table 1 provides general guidance on control compatibility for various environments. Furthermore, design guidance is provided for common ZT controls.

Table 1: Control Compatibility

	Control Center & OT DMZ	Substations	Generation DCS
Network Segmentation and Software Defined Networks	👍👍	👍👍	👍👍
Application Layer (Deep Packet) Inspection) Gateways	👍👍	👍	👍
Secure Remote Access	👍👍	👍	👍
Secure Protocols	👍👍	👎	👎
Endpoint Protection	👍👍	👎	👎
Enhanced Identity Access Management	👍👍	👍	👍

Legend: 👍👍*: Multiple/widely supported options; granular device-specific controls can be implemented*

👍*: Limited/complex options, dependent on system-specific architecture; likely only system/site/network level controls can be implemented*

👎*: Very limited, if any, options, dependent on system-specific architecture; may not be feasible to deploy controls (cost/benefit, operational impacts)*

Network Segmentation and Software Defined Networks

Network segmentation allows entities to limit attack surfaces and disrupt or detect lateral network movement of attackers. It is a critical component and an early maturity step for ZT roadmaps. SDN allows organizations to create network segmentation faster and easier through automatic configuration of firewalls, switches, and routers. It offers more agile and flexible approaches to isolate or segment both VLANs and application layer network traffic over traditional tools through the use of policy-based configuration and security to establish least trust. The following are aspects of SDN and network segmentation as part of successful ZTA implementations:

- Establish security zones with application layer inspection gateways between zones
- Network segmentation of group related workloads for the purpose of establishing granular VLANs
- Separate management traffic from operational traffic even within single facility/site
- Use secure logical overlay networks to establish SDN and build software defined perimeters
- Provides means to implement ZT for devices not capable of deploying on-device security

- Network Access Controls–Provides conditional network access upon policy-based security assessment of end point device configurations and/or behaviors

Application Layer (Deep Packet) Inspection Gateways

A variety of devices are capable of traffic monitoring or control, contributing to ZT maturity by inspecting network packets up to the application layer. These devices work hand-in-hand with SDNs. Different solutions may be deployed either internally or at the perimeter of networks. Some may include additional features that enable network and data flow mapping, asset and configuration inventories, and intrusion detection or intrusion prevention capabilities. Examples of these technologies and their feature sets include the following:

Next Generation Firewalls

- Deployed at perimeters (north/south traffic) or internally if software-based (east/west traffic) to establish security zones
- Application-level access control for inbound/outbound traffic
- Malicious code detection
- Support for OT protocols–provides packet-level ability to allow/deny protocol-specific messages
- SSL decryption for packet inspection

Data-diodes

- Deployed at perimeters or internally
- Enforcement of one-way communications for strict data flow control

Passive Security Monitoring

- Ability to monitor OT/ICS traffic protocols, flows, and time analysis attributes providing insights into normal network conditions and detection of anomalous conditions
- Support for OT protocols–provides packet-level ability to recognize protocol-specific messages

Secure Remote Access

Secure remote access includes solutions that provide access to applications and services that utilize connection brokering, encryption, and intermediate systems. Depending on individual solution capabilities, additional features may include trustless policy-based identity and access management to grant conditional access. Examples of secure remote access solutions include service gateways and terminal servers with virtual application delivery or virtual desktop availability deployed at a network perimeter within a secure DMZ. Other features may include the following:

- Support for multifactor authentication and single sign-on
- Session policy controls: re-authentications and session timeouts
- Session monitoring and enhanced logging
- Data loss prevention

- Bandwidth control
- Inline malware prevention

Secure Protocols

An important aspect of integrating security into an entity's technology footprint with an emphasis on ZT is to standardize use of secure protocols over legacy and unsecure protocols. It is crucial that industry continues to push their original equipment manufacturers to support and build in functionality for leading protocol innovations. Likewise, it is the responsibility of entities to ensure that the selection and procurement of new technology prioritizes compatibility with the newest protocols. Furthermore, entities should consider that implementation and configuration includes architecting the ability to turn on or switch to an updated protocol later when all integrations and endpoints are fully compatible. If this designed-in approach is not taken, it is much more likely that a legacy/unsecure protocol will continue to be utilized due to the inconvenience and technological burden of change. Finally, it should be noted that intermediate technology, such as port servers, proxies, or gateways, may be necessary to facilitate secure protocol use in OT networks due to the presence of legacy devices. Below are examples of secure protocols being evaluated for use in the electricity OT field:

- mTLS
- IPSEC
- DNP3-SA v5/v6
- IEC62351

Endpoint Protection

Endpoint protection solutions include endpoint detection and response with signature-based and heuristic analysis to continuously monitor, detect, and respond to cyber threats (like ransomware and malware) as well as active intrusions by threat actors. Other solutions specialize in detection through configuration baseline monitoring of file integrity, software, services, and logical network ports. Additional features may include the following:

- Policy-based application whitelisting
- Unified endpoint management
- Host based software firewalls
- Endpoint security auditing
- Domain and URL web filtering

Enhanced Identity and Access Management

The concept of least privilege is not new, but it is brought forth with renewed vigor in the paradigm shift of a ZT controls philosophy where the achievement of the "least" is examined in greater detail. Therefore, identity and access management under a ZT maturity roadmap can provide technology solutions to further scrutinize the how, what, and when for authorization, authentication, and access to applications and data. Newer technologies incorporate sophisticated policy based intelligence to support both RBAC or attribute-based access control (ABAC) strategies while enabling risk-based decisions, such as raising authentication from single factor to multi factor and granting reduced privilege to a resource. For example, instead of outright success or deny of access, dynamic access may be granted while locking out some features, specific categories of data, or by simply reducing privileges to read-only over full edit/full control.

The implementation between ABAC and RBAC are significantly different with more complexity and automated control being delivered with ABAC and easier implementation but less granular controls with RBAC. Both are well worth exploring for implementation to support ZT practices. Additionally, entities may consider requiring out of band approval for system management access. This is a best practice implementation to remediate the ability of an attacker to approve elevated system access on a node they have successfully infiltrated. The out of band approval process restricts request and access granting to systems and networks that are not accessible by the grantor systems and networks.

Public key infrastructure (PKI) can be used to support identity and access management controls by providing a robust framework for secure authentication, authorization, and encryption. PKI employs digital certificates, which are issued and validated by trusted certificate authorities, to establish trust between parties. Through the use of asymmetric cryptography, PKI enables the secure exchange of digital signatures and encrypted data, ensuring that only authorized individuals or entities can access specific resources.

Conclusion

ZT offers the electric industry a clear direction forward for continual improvement to securing our critical infrastructure. ZT is a paradigm shift that builds on and enhances existing controls and capabilities of cyber security plans. Security policy enforcement becomes data-centric (what data requires protection) instead of network-centric or device-centric. The emphasis is on entity identity and context over location within a perimeter. Research and testing must be completed to successfully transition with minimal disruption.

Industry also needs to continue to develop equipment and software as well as people, processes, policies, and governance capable of delivering on ZT principles. Advanced applications (e.g., real time contingency applications) and support applications (e.g., historians) offer likely paths for testing of implementations of ZT controls. Engineering access to equipment also offers a possible avenue to enable and test these concepts. Entities can collaborate and assist one another through memberships in various organizational groups. Government can provide tax incentives for infrastructure investments, grants for industry organizations promoting cyber security, and funding to assist less capable smaller entities with the process of moving to a more defensible electric infrastructure.

ZT implementation requires attention, focus, and planning. Stakeholder buy-in and executive support at the highest levels are essential for success. Developing a ZT environment in the OT space will take time and deliberate action. Some organizations have not started, some already have existing network infrastructure or controls in place that may classify at part of a ZTA while some may have already begun the transition to ZT. Regardless of where an entity is currently, all organizations should take the necessary steps to assess the value of ZT to their IT and OT security programs in support of BPS infrastructure and develop a roadmap to mature technology and controls towards ZTA with an emphasis on realistic time lines and resources to move themselves forward on the maturity scale. A well thought out implementation process will allow an organization to incorporate ZT incrementally and in collaboration with OT integrators and vendors as appropriate. It is crucial for the industry to take these steps of maturity to ensure resilience of the BPS against cyber threats and protect the critical function of providing secure and reliable electricity.

Appendix A: References and Resources

Control Design

- NIST SP 800-207 - Zero Trust Architecture
- NSA - Segment Networks and Deploy Application-Aware Defenses
- NIST SP 800-162 - Guide to Attribute Based Access Control (ABAC) Definition and Considerations