SAND2023-04588C

Application of Zero Trust Architectures for Nuclear Power Plants: Benefits and Challenges to Implementation

Benjamin Karch[1*], Andrew Hahn[1], Alexandria Haddad[1], Christopher C. Lamb[1]

[1]Sandia National Laboratories, Albuquerque, NM

[leave space for DOI, which will be inserted by ANS]

ABSTRACT

In this report, we look at Zero Trust Architecture (ZTA) principles and outline where and which tenets are applicable to nuclear power control systems, both for current generation systems and potential future Small/Modular and Advanced systems. ZTA approaches are becoming more popular in IT systems and are recommended approaches for building new systems. We have also seen some partial ZTA solutions in place for industrial systems, but nothing with the rigor required of nuclear power systems. We first define ZTA and discuss multiple current implementations in IT systems, cloud computing systems, and finally industrial systems. With this context, we then discuss where ZTA techniques can be applied in current and future systems based on current standards and regulatory guidance. We close the report with a summary of technical challenges that need to be addressed for ZTA to be useful, and where in nuclear systems ZTA can have the most impact on system security.

Keywords: zero trust architecture, ZTA, cybersecurity, advanced control systems, remote operation

1. INTRODUCTION

In the past decade, the world has witnessed an explosion of cyber-attacks on industrial systems and critical infrastructure. Starting in the 2010s, we have had attacks on critical infrastructure ranging from water treatment plants in Florida, to power systems in Ukraine, to industrial furnace facilities in Germany. Due to the deteriorating global political climate today, these attacks show no sign of becoming anything other than more frequent.

We have seen two attacks on or adjacent to nuclear systems. One, in the United States, was malware installed on a laptop in a nuclear facility business local area network (LAN) via spearfishing, watering hole attacks and exploit kits. The other, in India, did not breach control system protections, but did show a deep understanding of the attacked systems and resulted in large amounts of data being successfully exfiltrated from the facility. In fact, the first attack more broadly, and successfully, compromised facilities and personnel across the energy sector in the United States, including at federal agencies. The threat to our energy and nuclear systems is here, today, and likely to become stronger in the coming years.

Furthermore, costs associated with implementing cybersecurity controls is escalating across the nuclear sector. In an energy sector where we depend on carbon-free energy production that operates on razor thin

*brkarch@sandia.gov

margins, this increase in cybersecurity costs leads to an increase on cost per unit of power generated, making nuclear less competitive with other carbon-emitting energy production methods.

This report is an initial examination of using Zero Trust Architecture (ZTA) techniques to secure nuclear control systems. We briefly examine the history of ZTA to set the context and understand the motivations around creating it, look at the state of ZTA today via a group of case studies, and examine how we could apply ZTA to control systems in nuclear plants. When looking at potential nuclear plant application, we look over the regulatory and standards landscape to see how ZTA needs to adapt to this environment. We also look at key differences between Information Technology (IT) and operational technology (OT) systems that an OT flavor of ZTA would need to accommodate to be applied in nuclear systems.

2. DEFINING ZERO TRUST ARCHITECTURE

Most recently (2020), NIST provides abstract definitions for ZTA in Special Publication 800-207 [1]. This document defines a set of tenets that can be used to describe a network implementing ZTA, like those commandments set forth by the Jericho Forum [2].The tenets defined by NIST are as follows:

- All data sources and computing services are considered resources
- All communication is secured regardless of network location
- Access to individual enterprise resources is granted on a per-session basis
- 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
- The enterprise monitors and measures the integrity and security posture of all owned and associated assets
- All resource authentication and authorization are dynamic and strictly enforced before access is allowed
- 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

These tenets are purposely vague, so that determinations and design decisions can be made abstractly before assessing specifics for procedures such as authentication, log gathering, and identity management. However, some of these tenets prove to be difficult to implement due to their ambiguity. For example, one organization's implementation of a dynamic access policy will look quite different to another's, and there may be discrepancies in their efficacy.

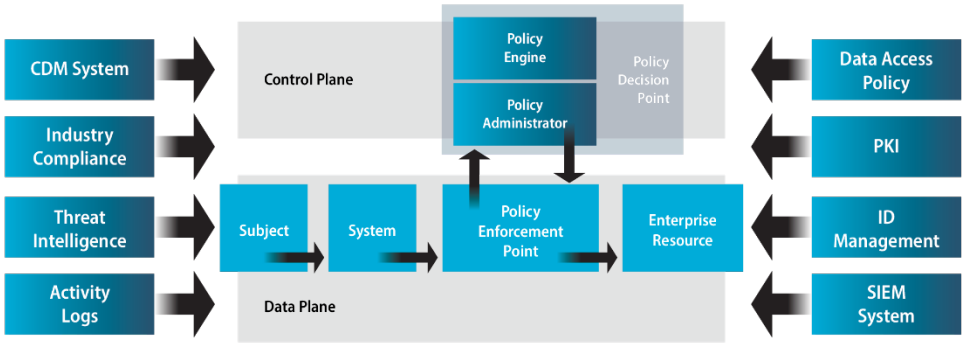

Figure 1. ZTA Components [1]

NIST improves on the general understanding of ZTA provided by the Jericho Forum by defining the logical components (depicted in Figure 1) found in a ZTA deployment. These components are used in determining access to resources by subjects. Each request must always pass through the Policy Enforcement Point (PEP), which will receive the request and negotiate with the Policy Decision Point (PDP). The PDP determines whether the system can be considered trusted and allowed to access the resource in question. The PEP is responsible for carrying out the decisions made at the PDP, enabling, monitoring, and terminating connections between systems and resources. The PDP is composed of two logical components: the Policy Administrator (PA) and Policy Engine (PE). The PE is the component which will utilize access policies, identity information, trust algorithms, etc., to decide if the system is allowed access. The PA will then establish the communications path (e.g., session-specific authentication tokens) between the system and resource, and issue relevant commands to the PEP. These components are distinct in their logical responsibilities but may be implemented in the same physical component.

The PE contains the trust algorithm [1]. The trust algorithm has several inputs to consider when making a decision regarding an access request. These include the access request itself, subject database, asset database, resource policy requirements, and any related threat intelligence and logs. It is important to draw a distinction between the subject and the asset in this process. The subject can be considered the person or process that is requesting access to the service, and the asset is the device that the subject's request originates from. The subject database contains a set of identity information / attributes that can be used to establish a level of confidence in the subject's purported identity. The asset database contains known status of the device: operating system (OS), location, installed software, etc. Establishing this type of information has been explored in IT-based systems but lacks research in OT/ICS systems. Instrumentation and Control (I&C) devices operate without a human present most of the time, meaning that the requestor for many access requests would be a process running on the I&C device. This process will typically be running with a high level of access to the asset's OS, making establishing trust in the subject and asset a challenge at the PE.

NIST also goes on to describe a few use cases and scenarios for which an organization would benefit from a ZTA implementation [1]. These scenarios are useful for outlining the benefits that are gained from a ZTA implementation. These include situations such as an enterprise with satellite facilities, cloud-to-cloud enterprises, and enterprises with contracted services and/or nonemployee access. These highlight the benefits of ZTA, mainly showing that ZTA offers organizations a cybersecurity solution in which data can stream between organizations, services, devices, etc. without the need for overbearing restrictions on whether that data flow may be visible to untrusted individuals or devices. This is because all data is protected by default and each endpoint is validated at the beginning of each session. The scenario most like a Nuclear Power Plant (NPP) would be the enterprise with contracted services. Vendors or maintenance personnel may be required to access NPPs physically, which poses a cybersecurity threat given their current cybersecurity practices. ZTA in this case allows those contractors to have access to resources and devices that they need to and be denied access by default to any other part of the network. Additionally, identity management components of ZTA mean that the contractors can be verified to be members of the organization that they represent and have their access levels immediately and automatically assigned.

2.1. NPP Cybersecurity Landscape

Typically, innovations in cybersecurity and computation in general are adopted first in Enterprise environments, and later may be integrated into OT environments. ZTA implementations are growing in popularity in the Enterprise sector, but just beginning to emerge for OT applications. For these OT applications, because the nuclear industry is highly risk averse, standard practices tend to stay commonplace for very long periods of time, and the process for implementing modern techniques may put a NPP in violation of standards.

Currently, the common practice for cybersecurity in a NPP involves strict perimeterization. The current security approach involves network segmentation according to the security posture of the physical area [3]. Each boundary implies a division in the network, often many individual networks exist within the boundaries of these areas. Within these networks, there is often very little or no cybersecurity measures, and cybersecurity is assumed because of stringent physical access requirements.

3. EXISTING OT IMPLEMENTATIONS

Very few well-documented implementations of ZTA for OT networks exist. In 2021 Deloitte released a document describing a project that claims to have established a ZTA in a large ICS spread across multiple facilities of a chemical manufacturer [4]. Using the Purdue model of ICS architecture as a basis, Deloitte integrated principals of ZTA into the design of this ICS network. Implementing micro-segmentation, data flow restrictions and controls, access control, VPNs and more secure jump servers, and monitoring capabilities with a SIEM has certainly made a more secure environment. These security principals laid over the Purdue model are proposed as a reference architecture for ZTA in OT networks.

While the Deloitte implementation does have many qualities of ZTA, not all tenets are fully implemented in accordance with NIST 800-207 [1]. The described implementation contains many areas that have implicit trust between devices and is highly reliant on perimeterization of these implicit trust zones. This architecture is more focused on the interconnections to and above level 3 of the Purdue model, leaving the entire OT system from production servers and databases down to the PLCs and physical process controls as implicit trust areas of the network. Though the network is segmented, individual resource access and communications in the OT network are not secured or controlled.

The architecture that Deloitte developed is a step in the right direction for OT network security, but it does not fully qualify as a ZTA by the fundamental tenets of Zero Trust. The level of control over communications and resources within OT networks that ZTA demands is difficult to retrofit into systems that were never designed to provide these resources. The reasons that there are so few attempts at ZTA implementation in OT could be explained by the current state of the technology available and a cost benefit imbalance for system owners.

4. ANALYSIS OF ZTA FOR NUCLEAR POWER PLANTS

Typical IT implementations do not include custom identity management; generally existing authentication methods like Kerberos are utilized. This vastly constrains the type of devices that can be used, and compatible devices do not include ICS devices like PLCs, HMIs, or other typical industrial systems. Given currently available literature, ZTA can be implemented in the IT environment that exists at the plant before the historian, like the Deloitte implementation discussed above. In other words, a NPP can implement ZTA for the business portion of their network, but there is no available literature or tools for extending that environment into the operational portion of the network. If a NPP was to extend this capability into the OT network, it would violate regulation by setting an accessible path to safety-related ICS devices from outside of the network, i.e. deperimeterized.

Furthermore, industrial systems must be viewed through an operational lens, where availability and integrity are more important than confidentiality. This is in direct contrast to business environments, where confidentiality is of great concern and brief disturbances in availability do not result in potential safety impacts (and are often expected). Recognizing all computing services and data sources and sinks as resources, we can re-interpret ZTA guidance regarding secure communications. Specifically, communication must be secured by focusing on availability and integrity first and foremost. This specifically impacts NIST ZTA tenets (1) and (2) and is a distinctly different approach to how communications are secured in IT systems, where confidentiality and integrity are typically more important attributes than availability. Not to imply that IT communications are not concerned about system availability, rather that availability is much more important to industrial systems and NPP as a lack of data availability at the wrong time can lead to significant physical consequences.

In addition, granting access to resources within a control system is more frequently system-to-system use, not user-to-system use; this renders modern techniques used to strongly authenticate users like multi-factor authentication much more difficult. Authenticating a system based on something that system has, like a certificate, is possible, but evaluating something a system knows is not, as that is trivially accessible

if a system is compromised. This makes controlling system access on a per-session basis more difficult. Certainly, systems can re-present certificates, but this may be of little value as those certificates are present on compromised systems as well as uncompromised ones. In ZTA implementations, systems are identified using a device database. For example, Google's BeyondCorp [5] uses a device inventory database and requires each user device to contain a Trusted Platform Module (TPM). Modern mobile devices like iPhones use a Secure Enclave Processor (SEP), separate and inaccessible to the application processor to store sensitive information like certificates or biometric data. Google uses trusted computing support to identify systems in tandem with their device database. Other approaches use weaker methods to identify systems like MAC addresses to authenticate systems for network access, but these kinds of methods can result in compromise through MAC spoofing attacks. Overall, ICS systems were not historically designed to provide these kinds of strong authentication services. This impacts NIST ZTA tenet (3).

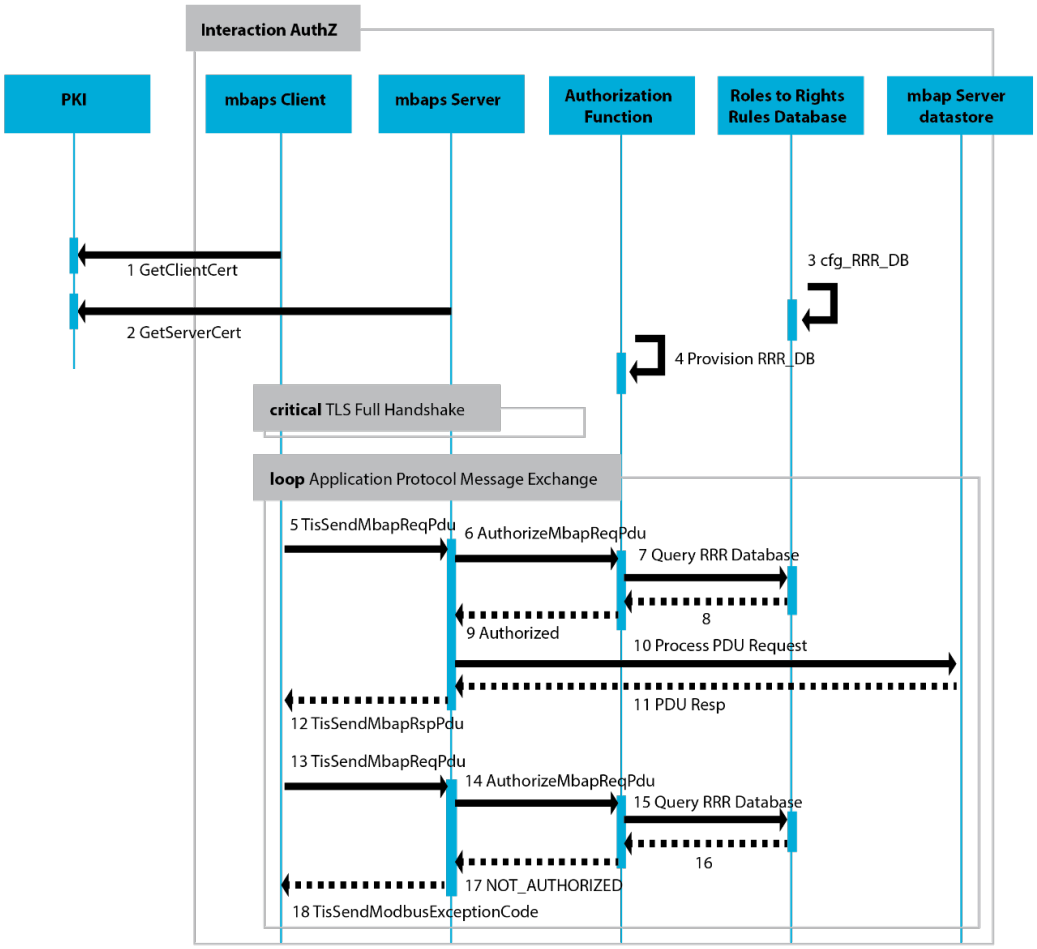

Figure 2. Modbus/TCP Security

NIST ZTA tenets (4) and (6) require dynamic authentication policy application after authentication material is presented. This requires fusing state information with policies describing system-to-system access. Typically, this is applied at PEPs and PDPs within authentication systems and require policy storage and access. Even if this kind of evaluation is engineered to be extraordinarily performant, it still creates dependencies on authenticating systems. This leads to additional potential points of failure (which must be very carefully managed), as well as some unavoidable additional latency. As connections

between resources is managed on a per-session basis, the amount of introduced latency is heavily variable, but the system dependencies will always exist in some form.

NPP licensees typically apply NIST ZTA tenets (5) and (7), already. This is done by closely monitoring system behavior for pending failures and overall performance, as well as system security, when designed and built following international standards and national regulatory guidance.

NIST ZTA tenets (1) and (2) have a different focus in NPP systems. Most protocols used to secure communication provide confidentiality and integrity protections, but not necessarily availability protections. Typically, availability is provided by other protocols in each network stack. TLS and IPSEC are both common communication protection protocols, but session control, including retransmission, is provided by TCP. ZTA tenets (3), (4), and (6) have significant barriers to implementation, potentially increasing system risk by increasing the risks of single point and common cause failures as well as increasing communication latency in unpredictable ways. An additional consideration for these tenets is the development, purchase, and integration of OT-focused solutions for operators. Finally, tenets (5) and (7) are typically adhered to within NPPs today.

5. ZTA for NPP Systems

Advantages. Many advantages can be achieved through implementing ZTA in an organization, and those advantages also apply to NPPs. Current cybersecurity practices in NPPs are lacking in maturity; these practices have not kept pace with advancements in cybersecurity that have become commonplace in IT and enterprise environments and make many assumptions about access to resources that may not hold true. If a NPP were to implement ZTA across its entire network architecture, its cybersecurity posture would be much more hardened than that of the typical NPP. This is due to the various tenets and requirements of ZTA implementations, for example that data should be protected in use, transfer, and storage. Many OT environments utilize legacy serial communications protocols that do not offer modern cybersecurity protections like encryption, message authentication, or non-repudiation.

ZTA puts a strong emphasis on federation and identity management for threat intelligence purposes, as defined by tenet 4 of NIST SP 800-207 [1]. Policy must be set not only according to identities of assets as users, but also environmental factors and behavioral attributes. A proper ZTA implementation takes measures to record information relevant to these attributes, which can be used for threat intelligence and continuous monitoring, adhering to tenet 5. Federation and Identity Management are further defined in NIST SP 800-63c [6]. These requirements will result in any ZTA-implementing NPP's maintenance of rigorous databases on relevant information to circumstances surrounding a user's access to important NPP operational data. A Licensee may utilize a Federation Authority to normalize data across multiple plants or sites. It is also possible for a Federation Authority to be developed such that multiple separate Licensees may benefit from the recorded behavioral characteristics from another Licensee's plant. This federation would provide a strong basis for real-time threat detection, as the Federation Authority can revoke access levels of a party to all its managed environments after malicious or abnormal activity is detected at one. This directly lowers risk for Licensees by reducing the potential impact of a malicious actor.

ZTA may also provide an advantage by increasing the predictability of traffic found on the network. This is because when a user or machine on the network accesses a resource, it must follow a specific procedure to gain access to its desired network resource. For example, a network implementing an access policy similar to [7] would expect to see the first packet of the network traffic include some key used to identify the subject. Should a network analyst (either human or automated) observe network traffic that is not consistent with this procedure, there can be a somewhat high degree of confidence that the traffic originates from an attempt to gain unauthorized access. This information is then able to be processed for ongoing threat intelligence and shared within and beyond the organization, as appropriate.

[8] scenario 4 details how the flexibility gained with a ZTA implementation could offer advantages to an NPP. NPPs operate according to standards and guidance from a multitude of organizations including standards bodies like the Institute of Electrical and Electronics Engineers (IEEE), International Electrochemical Commission (IEC), and International Atomic Energy Agency (IAEA), and national

regulators like the US Nuclear Regulatory Commission (NRC). They maintain product supply chains originating from many different countries as well. A NPP which implements ZTA will have an easier time maintaining compliance and compatibility by instituting an automated audit system for continuous monitoring. Core principle 6, Alignment and Automation, is applied to the plant to enable real-time log capture, storage, and analysis, meaning that automated analysis rules can provide Operators with immediate notification of standards compliance. Upon updates to standards or guidance, the rules are updated, and all non-compliant assets or procedures are detected for remediation.

ZTA must also provide an asset-centric security approach to the network, as stated in core principle 8 [8]. This means security practices are tailored to the assets within the network rather than an approach that is application-centric, for example. This reduces the complexity of the network and enables an easier exchange of data between interfaces. Approaches such as format preserving encryption and tokenization reduce the complexity of interfaces that must be maintained between different types of assets. ZTA also helps to secure high-value systems by adapting to changes in the environment quickly and autonomously. ZTA specified secured zones, policy-driven access control, and context-specific data security. This means that policies can be adaptive and provide data security depending on situational risk at any time to systems without the need to network-wide changes or updates.

Challenges. There are many challenges that must be overcome to implement ZTA in an NPP; most of these challenges arise from the inherent lack of cybersecurity features in most ICS and SCADA systems and components. For example, generally the control system for the NPP will be implemented over a serial communication protocol. Common serial communication protocols like DNP3 or modbus do not implement encryption, authentication, or other security measures by default [9]. ZTA requires data to be protected at all points during its lifecycle, so there must be some hurdles overcome to enable encrypted communications within the plant's operating environment. There are some protocols which implement modern security standards, such as Modbus/TCP support for Transport Layer Security (TLS). An overview of the process for Modbus/TCP security is depicted in Figure 2 [10].

A critical component of the security for Modbus/TCP and TLS in general is labels 1 and 2 in Figure 2 where the client and server verify each other's identity using X.509 certificates. This requires a robust Public Key Infrastructure (PKI) be developed and maintained for the devices, which is currently not in practice. A PKI internal to the plant could be operated, but the storage of keys must also be secure. BeyondCorp can relieve itself of many challenges associated with this process by requiring that all assets on the internal network include a TPM that is managed by the organization [11]. This is infeasible using current ICS devices because manufacturers either do not include any TPM or TPM-like component or do not allow for operators to access these components through any provided programming software.

The challenge of maintaining secure private keys and certificates for TLS encryption gives rise to the larger issue of identity maintenance in general for ICS devices. In Enterprise environments, there has been decades of research and development invested in the creation of protocols and architectures for identity management of both users and their devices, such as Active Directory. Active Directory uses Kerberos to issue and keep track of access requests and "tickets" that are served to users based on successful requests for access to assets [12]. There is no such architecture in place for ICS, so some identity and access management system must either be developed and tailored to ICS devices or manufacturers of ICS devices must make major changes to both included hardware and device firmware to make them compatible with existing solutions in the IT space. It may also prove to be challenging to maintain required real-time constraints on ICS devices when implementing modern encryption and authentication protocols.

ZTA implementations must also maintain robust databases for users, not just digital assets. This poses a unique problem with ICS devices because there is not generally any user space on the device. The concept of separate address spaces for user and kernel space is implemented in Operating Systems that are used in IT environments. This logical separation ensures that a typical user on a device is not able to access fragile components of the system's kernel or directly interface with the system's hardware. The kernel implements many security measures to provide this assurance. Real-Time Operating Systems (RTOSs) operate exclusively in the kernel address space or do not protect kernel address space from processes running in user space. This means that any user who gains access to an ICS device running an RTOS can masquerade requests as coming from the digital asset rather than the user. Additionally, kernel space

processes that are used for reporting accesses by users or other important security information that is required to be logged and analyzed constantly by ZTA can be subverted relatively easily.

Significant challenges for ZTA stem from the needs of the operational environment present in NPPs. The communications in OT are far more important to the safety of a facility than communications in an IT network. A PLC reporting a safety value, or a stop command, must reach its destination to ensure safe operation. It is also vital in OT systems that information is timely, and communications cannot be encumbered by encryption or other security controls that delay safety systems or stresses the limited computational resources on controllers, especially in situations where performance of these PLCs could be degraded. The security of the information communicated is far less important than ensuring the information's integrity and validity, which presents a very different paradigm from IT security.

For continued safe operations of a process control system, some controllers may need to be allowed to communicate and operate even if de-authenticated. For many systems in an NPP the loss of authentication on a PLC and subsequent communications halt could blind operators and reliant control systems from critical process data. A ZTA implementation will need flexibility to allow some contingencies for scenarios where a device critical to operations can be untrusted but allowed to operate on the network. This is a unique requirement of OT for security, and therefore some method of retaining operational safety while ensuring security must be implemented.

A method of remote authentication of devices may need to be developed. Should a device fail authentication from network faults or environmentally induced noise, it may need to be reauthenticated remotely. Demanding a full stop of production or a reactor shut down for a device to be re-authenticated would be an undue burden on operators and reduce the likelihood of adoption and retention of ZTA security practices. This could be aggravated in emergency conditions where areas of the control system may be inaccessible to workers, but control over the system must be guaranteed. This presents a major issue in ensuring the security of a ZTA network, potentially introducing a backdoor to later exploit. But operation outside of normal parameters and expectations must be considered to ensure that potential implementations of security measures do not hinder or limit emergency responses or compromise system safety.

Implementation of ZTA would not contradict the cybersecurity regulations according to CFR 73.54 [13] and would likely improve cybersecurity posture of plants. However, a partial or complete implementation of ZTA would likely result in systemic changes to a plant's network structure, cybersecurity and incident response plans, and communication flows. This means that a new Cyber Security Plan must be developed in accordance with NEI 08-09 [14].

Many nations derive NPP cybersecurity regulations and guidance from IAEA documents on the subject. For example, Section 5.2 (Cybersecurity Measures) of Canadian Nuclear Safety Commission DIS-21-03 [15] references three IAEA documents: NSS 17-T [16], NSS 23-G [17], and NSS 33-T [18].

IAEA NSS 17-T [16] recommends establishing logical and physical boundaries for information flows based on associated risk levels of information. Boundaries increase as risk decreases. ZTA implementations are not conducive to logical boundaries of data, as the purpose is for a deperimeterized network. So, while a ZTA implementation can provide rigorous data protections that allow for data to safely leave local network segments, guidance provided by IAEA may not allow for this flexibility. ZTA Tenet 2 states that access requests originating internally and externally must meet the same requirements, and that all communications must be secured regardless of network location [1].

IAEA NSS 23-G [17] provides guidance on establishing confidentiality, integrity, and authenticity of information and communications within an NPP. This guidance allows for information to be shared with outside organizations, such as appropriate state agencies, external states, or the public, when necessary. The guidance also provides the ability to change security measures of information according to regular audits and investigations. The Information Security Plan guidance in 23-G aligns heavily with ZTA principles. For example, the plan should provide for regular monitoring and review to ensure that procedures remain relevant and effective, which aligns with ZTA Tenet 5 [1]. Additionally, for sensitive information, information must be restricted to those who need access to perform their duties, have been granted the authority, and who have undergone a trustworthiness check commensurate with the classification level of the information [17]. ZTA provides a framework for which this trustworthiness

check can be implemented as well as allowing for the requirements of access to be dynamically updated according to new information, and subsequently enforced by the PEP autonomously.

IAEA NSS 33-T's objective is "to provide guidance for the protection of I&C systems at nuclear facilities on computer security against malicious acts that could prevent such systems from performing their safety and security related functions" [18]. A key concept to this guidance is a risk informed approach to cybersecurity. This is in line with other IAEA guidance on cybersecurity, following closely in line with information classification levels, and security measures that line up with those measures. A ZTA implementation is well suited for compliance with NSS 33-T, as a strong relationship between risk management and computer security teams can be made effective and efficient with dynamic policy updates, continuous monitoring, and revocation or adjustments of access privileges based on real time threat intelligence and changing risk factors / consequences. NSS 33-T prescribes a strict access control policy and a minimal number of access points. This is implemented in ZTA with a centralized PDP and PEP.

6. CONCLUSIONS

ZTA has the potential to provide a more secure posture for OT systems soon and allow modes of operation not previously possible that would provide major economic benefits to new designs. Many of the new generation of reactors are designed around the concepts of inherent and passive safety. This major shift in design principles moves away from active safety systems that must intervene to return the plant to a safe condition. Safety is central to the design so that no operator or control system intervention is required to ensure that the reactor is within a safe operational envelope. This could allow for a larger number of OT devices that are classified at lower safety or security levels, i.e., the regulatory requirements on the strict bounds of communications could be lessened. This should reduce cost of advanced reactor designs and may make ZTA far more relevant to these plants.

For the current fleet of reactors, implementing ZTA would come at a high cost of equipment replacement. Because these reactors have some inherent cybersecurity from their diversity and redundancy, there would be little benefit to implementing ZTA for the excessive cost of a full system replacement. Staged deployment as systems are upgraded in their natural replacement cycles could be a reasonable path if the equipment cost differential was not excessive. ZTA in the current fleet would be better suited to IT systems and improving the cybersecurity posture of administrative systems and systems in higher levels of the nuclear control system hierarchy as outlined by the IAEA.

For reactors that are yet to be constructed the implementation costs are significantly lower. With designs that intend to have centralized control rooms and streamlined control systems it will be necessary to consider the importance of secure communications. Designers will need to evaluate the potential for malicious operation and if the control systems have the capability to put any part of the plant in dangerous operational modes. Since the cost of implementation would lower than the current fleet, and some protection via diversity is lost, if the technology is available for ZTA implementation it may provide significant improvement to the cybersecurity posture of future reactors.

If remote operation is viable, it would require some method of highly secure communication architecture like ZTA. The cost to implement ZTA for remote operations would be less concerning than the critical need of cybersecurity in a remote operation application. The cost of R&D would be the critical factor for enabling this mode of operation.

Overall, the most significant technical hurdle to ZTA implementation is dynamic session-based system-to-system authentication. Incorporating this kind of authentication using techniques typically used today via remote device databases and certificate authorities creates additional layers of complexity in control systems. This additional complexity imposes new single-point and common-cause failure risks and creates unpredictable increases in latency. These new failure risks and latency increases could potentially be eliminated with new engineering approaches, but this is a currently unexplored area of research.

ZTA approaches can increase system security overall but require careful and thoughtful application to be cost and functionally effective.

References

[1] S. Rose, O. Borchert, S. Mitchell and S. Connelly, "NIST Special Publication 800-207: Zero Trust Architecture," 2020.

[2] J. Forum, ""Identity" Commandments".

[3] IAEA, "NSS 27-G: Physical Protection of Nuclear Material and Nuclear Facilities," IAEA, Vienna, 2018.

[4] Deloitte, "Achieving a Zero Trust Architecture in an Industrial Environment with Multiple Facilites," Deloitte, 2021.

[5] "BeyondCorp," Google, [Online]. Available: https://cloud.google.com/beyondcorp.

[6] P. A. Grassi, J. P. Richer, S. K. Squire, J. L. Fenton, E. M. Nadeau, N. B. Lefkovitz, J. M. Danker, Y.-Y. Choong, K. K. Greene and M. F. Theofanos, "NIST Special Publication 800-63c: Digital Identity Guidelines," 2017.

[7] C. DeCusatis, P. Liengtiraphan, A. Sager and M. Pinelli, "Implementing Zero Trust Cloud Networks with Transport Access Control and First Packet Authentication," 2016.

[8] T. Ghosh, N. Kumar, S. M. Sakuru, P. Shirazi, M. Simos, A. Valani, A. Carrato, S. Whitlock, J. Hietala, J. Linford and A. Szakal, "Zero Trust Core Principles," The Open Group, 2021.

[9] S. East, J. Butts, M. Papa and S. Shenoi, "A Taxonomy of Attacks on the DNP3 Protocol," 2009.

[10] Modbus, "MODBUS/TCP Security," 2018.

[11] R. Ward and B. Beyer, "BeyondCorp: A New Approach to Enterprise Security," Google, 2014.

[12] Microsoft, "Kerberos Authentication Overview," [Online]. Available: https://learn.microsoft.com/en-us/windows-server/security/kerberos/kerberos-authentication-overview.

[13] U.S. NRC, "Protection of digital computer and communication systems and networks.," [Online]. Available: https://www.nrc.gov/reading-rm/doc-collections/cfr/part073/part073-0054.html.

[14] Nuclear Energy Institute, "NEI 08-09 [Rev. 6]," 2010.

[15] Canadian Nuclear Safety Commission, "DIS-21-03, Cyber Security and the Protection of Digital Information".

[16] "IAEA Nuclear Security Series No. 17-T," International Atomic Energy Agency, Vienna, 2021.

[17] IAEA, "NSS 23-G: Security of Nuclear Information," 2015.

[18] IAEA, "NSS 33-T: Computer Security of Instrumentation and Control Systems at Nuclear Facilities," 2018.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.