Towards Zero Trust For Critical Infrastructure:
Rethinking The Industrial Demilitarized Zone

E. B. Boumhaout[1] , **A. S. Danielsen**[2] , **O. B. E. Pedersen**[3] , **M. Shahid**[4]
Massachusetts Institute of Technology
[1]elbachir@mit.edu, [2]akselsd@mit.edu, [3]olebjorn@mit.edu, [4]mshahid@mit.edu

Abstract

Industrial control systems (ICS) that power critical infrastructure such as electric grids and metro systems have become high profile targets for nation-state attackers. They are becoming more susceptible to cyber-physical attack with many ICS deployments becoming digitized and connected to the internet. We propose a system design approach to apply Zero Trust architecture to Industrial Control Systems (ICS). Zero Trust is a model that proposes shifting from perimeter defence approach to an access-based defense which is location-agnostic and more reliable in today's data ecosystem that no longer follows a hierarchical flow. We identify challenges to adopting Zero Trust in ICS architecture, as well as carve a path to solve them.

1 Introduction

From traffic lights and metro systems, to oil refineries and nuclear power plants, industrial control systems manage some of the most critical infrastructure services. They sustain various functions that are vital to our everyday life and are enabled by Industrial Control Systems. Today, many critical infrastructures are connected to the internet, either directly or indirectly, where digital and physical systems are converging. Thus, the security of critical infrastructure has to be of the highest standards and must be a priority to guard against the continuous threats of increasingly sophisticated malware and malicious attacks.

1.1 Industrial Control Systems

Industrial control system is a term used to encompass various automation systems and its devices involved in the control, security and operations of industrial systems. These include:

- Programmable Logic Controllers (PLC)
- Human Machine Interface (HMI)
- Intelligent Electronic Devices (IEDs)
- Supervisory Control and Data Acquisition (SCADA) systems
- Distributed Control Systems (DCS)
- Safety Instrumented Systems (SIS)
- Sensors

Each of these systems is crucial to support critical infrastructures. They constitute a blend between Operational Technology (OT) and Informational Technology (IT) that enable the monitoring and control of ICSes.

1.2 Zero Trust

The term Zero Trust was first introduced by Forrester Research in 2010 as a security model/paradigm for software defined networks that seeks to eliminate the concept of trusted/untrusted (i.e. internal/external) networks [Forrester Research, 2013]. The paradigm has three key concepts:

- Ensure that all resources are accessed securely regardless of location
- Adopt a least privilege strategy and strictly enforce access control
- Inspect and log all traffic

Zero Trust Networks challenge many of the assumptions of traditional perimeter defense networks, providing a new approach to security. It enforces adaptive controls, and continuously verifies trust. This approach can help prevent unauthorized access, contain breaches and reduce the risk of an attacker's lateral movement.

1.3 Threat Model

For this project, we are assuming that the ICS must be secure against Advanced Persistent Threats (APT). This means a skilled and resourceful group that are potentially payed salaries for their work and may also be backed by a government. The attackers may be motivated by a potential financial gain, or simply wish to destabilize or otherwise sabotage an ICS. We are assuming they can launch their attack over an extended period of time, in order to remain undetected and slowly gain more access and privileges in the system.

2 Related Work

To the best of our knowledge, there are no publicly available publications proposing or implementing a Zero Trust architecture on industrial control systems. However, there has been extensive effort to improve and standardize ICS

security. In 2014 the National Institute of Standards and Technology (NIST) published a Guide to Industrial Control Systems Security [Stouffer *et al.*, 2011] with the purpose of providing guidance for securing ICS, including its subsystems (SCADA, DCS, PLCs etc.). Specifically, the publication provides guidelines for designing a ICS security architecture. In general, there's extensive focus on classical boundary protection through the use of firewalls. However, they also have a great focus on central Zero Trust concepts such as network segmentation and segregation, authentication and authorization, and monitoring and auditing.

Aside from the general guidelines provided by NIST, there have been proposed several ICS architectures in literature. An interesting publication from Idaho National Laboratory proposes a hierarchical architecture that extends on the hierarchical viewpoint of the Purdue Reference Model, and integrates it for control systems in smart critical infrastructures [Zhu *et al.*, 2011]. Their proposal provides stricter segregation of layers in a vertical fashion. However, the proposed architecture provides no horizontal protection between individual units at the same horizontal layer. Therefore, it is still vulnerable to the type of security risks that the Zero Trust model aims to diminish.

Even though there has been no effort to implement a full scale Zero Trust architecture on an ICS, individual Zero Trust compliant components have been proposed for cyber-physical systems (CPS) [Chiluvuri *et al.*, 2015]. Specifically, a monitoring component coined a "trust-worthy autonomic interface guardian architecture" (TAIGA) is proposed and implemented on a System on Chip (SoC). The component is designed to separate trust between the physical control system and the production controller. Since the production controller often is connected to the enterprise network, such a component could be critical in a Zero Trust architecture.

3 Current State of ICS Security

3.1 Purdue Enterprise Reference Architecture

In this project, we adopt the Purdue model for ICS by ISA-99 as a concept architecture which is considered an industry best practice for industrial control system network segmentation. The model relies on the concept of zones to segment enterprise and ICS network into logical parts composed of systems that perform similar functions or have similar requirements. In our analysis and proposal, we focus on Level 3, the Industrial Demilitarized Zone and Level 4 highlighted in figure 1 as they constitute the epicenter of critical data flows and potential exploits.

Level 4 - The Enterprise Security Zone

Level 4 in the Purdue model is the first layer in the enterprise security zone, and is also the first level with internet access. The level provides services for all the enterprise levels above which may use summarized data from the lower levels [Ackerman, 2017]. The layer performs the following major functions: capacity planning, plant production scheduling, materials requirements planning, and all manufacturing related

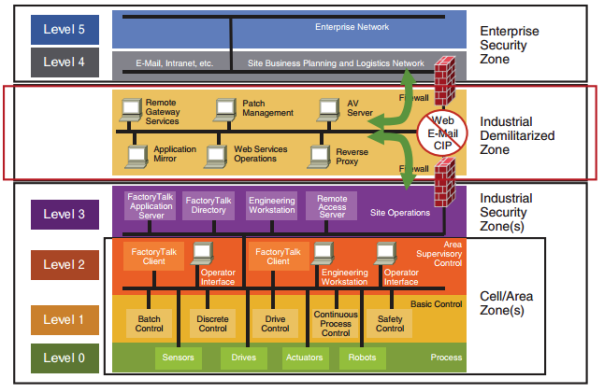

Figure 1: Purdue Model Architecture for ICS [Ackerman, 2017]

purchasing [Albert W. Jones, 2000]. Other typical services provided at this layer may include [Ackerman, 2017]:

- Email
- Enterprise applications, i.e. SAP and Oracle
- Remote desktop services to lower levels

The stakeholders at this level are mainly employees in operational management groups, as they will have access to the services and the data provided from the lower levels. The operational management groups are responsible for the major functions involving site business planning and logistics.

Level 3 - The Industrial Security Zone

Level 3 in the Purdue model is the upper level of the industrial security zone, and is responsible for all the communication with level 4. Some of the major activities performed at level 3 are [Albert W. Jones, 2000]:

- Resource allocation and control
- Data collection and acquisition
- Performance analysis
- Process management
- Quality management

Several computer services and systems are used to facilitate these activities, some of the typical ones are [Ackerman, 2017]:

- File servers
- OS/application patching service
- Engineering workstations with remote access service
- Inter-VLAN routing and traffic inspection
- Network and domain services, such as Active Directory (AD), Dynamic Host Configuration Protocol (DHCP), Dynamic Naming Services (DNS), Windows Internet Naming Service (WINS), Network Time Protocol (NTP) etc.

Level 3 of Purdue model is the level within the industrial security zone that sees most interaction. This is the area where operators in central control rooms log into shared computer systems, and observe and interact with plant wide systems

and applications [Ackerman, 2017]. The stakeholders at this level are the operators that physically, or through remote services, work at the industrial site.

Industrial Demilitarized Zone (IDMZ)

An IDMZ is a physical or logical subnet that separates an internal local area network (LAN) from other untrusted networks – usually the public internet. The IDMZ is an information sharing layer that is comprised of:

- Boundary or "edge" security appliances like firewall(s) that can inspect traffic as it enters and exits each security zone (enterprise and industrial).
- Appliances and servers that replicate services like web proxies, data proxies, file transfer proxies, application and operating system patch proxies, and application proxies [cisco, 2017].

An IDMZ is sometimes referred to a perimeter network that exposes an organization's trusted external services and data to an untrusted network.

IDMZ primarily holds Machine to Machine Interaction and checks traffic going from Enterprise Zone to Industrial zone, hence it has no direct stakeholders. However, stakeholders at level 3 and level 4 could be considered its indirect stakeholders.

4 Zero Trust Architecture

Zero Trust model proposes a shift from perimeter defence approach to an access-based defense. Instead of dividing the internal and external traffic as safe and unsafe, respectively, Zero Trust assumes the network is always hostile, and each interaction needs to be verified. One pillar of Zero Trust is the implementation of role-based access model where the concept of least-privilege is adopted to restrict access to users who legitimately need access to a resource.

Another pillar of Zero Trust is Micro-Segmentation. Micro-Segmentation is a security technique that enables fine-grained security policies to be assigned to data center applications, down to the workload level as well as devices. This means that security policies can be synchronized with a virtual network, virtual machine, operating system or other virtual security targets.

The concept of adaptable trust based on constant monitoring is important for Zero Trust as well. Micro-segmentation allows for much more granular data that can be used by the Trust Inferer to make policy-based authorization decisions for each transaction request. Trust Inferer is discussed in more detail in the next section.

Zero Trust is not a technology in and of itself but a shift in the design approach for cyber-security. It has to be customized to suit the need of the network to ensure security while not impeding availability, integrity and confidentiality.

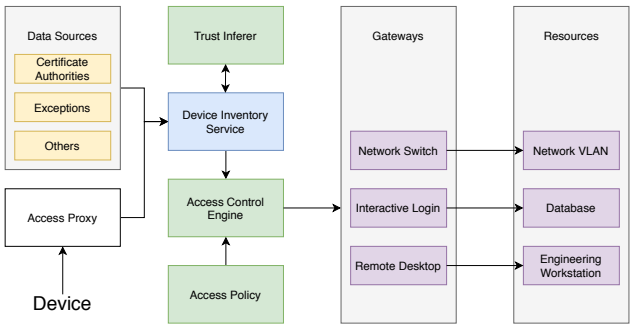

Figure 2: High-level Zero Trust architecture data-flow, modified from [Osborn *et al.*, 2016].

4.1 Zero Trust for ICS

Based on Google's corporate Zero Trust architecture [Osborn *et al.*, 2016], we propose a new high-level Zero Trust architecture intended for industrial control systems. The new architecture is fundamentally different from the current ICS security models in that it segregates every device and resource, as opposed to segregation of levels in a hierarchical fashion.

The high-level Zero Trust workflow model is illustrated in figure 2 2. The figure shows the data-flow of a device trying to access an ICS resource. The device connects to the system through an *Access Proxy*, the single entry point. Further, the device is routed to the heart of the architecture, the *Device Inventory Service*.

Device Inventory Service

The device inventory service is responsible for managing all devices in the system. It continuously collects, processes, and publishes changes about the state of known devices [Osborn *et al.*, 2016]. The continuous monitoring is imperative, as an adversary may gain control of a device after a connection is established. In order to evaluate the state of a given device, the device inventory service has to process a variety of data sources. The processed aggregate data is sent to the *Trust Inferer*, which indicates whether the given device state is trustworthy or not, by assigning a trust score. The trust score is important since each resource will have an associated minimum score required for access.

Data Sources

The possible data sources used in a Zero Trust architecture are many, however, they can be categorized in two main flavors: observed and prescribed [Osborn *et al.*, 2016]. Observed data is programmatically generated and may include items such as OS version and patch level, last security scan date and result, and a list of any apps installed on the device. For an ICS, the most important observed data sources, besides the general ones, include:

- Real-time process monitoring data. Using real-time data from process monitoring, irregular system behaviour could potentially be correlated with device activity.
- Modification time-stamps from control units. Suspicious device behaviour can be detected by comparing

control unit modification time-stamps with regular activity.

- Safety instrumented systems (SIS) data. SIS are dedicated safety monitoring systems. They are there to safely and gracefully shut down the monitored system or bring that system to a predefined safe state in case of a hardware malfunction [Ackerman, 2017]. Suspicious activity should be detected before the SIS is effected. However, as a last resort, SIS state data can be used to detect irregular activity.

Prescribed data, on the other hand, are manually maintained by IT Operations, and typically include a variety of assignments [Osborn *et al.*, 2016]. For an ICS architecture, the most important prescribed data sources include:

- Exception list. The device inventory services considers preexisting exceptions from the exception list, and can allow for overrides to the general access policy [Osborn *et al.*, 2016]. A well defined exceptions list is especially important for industrial control systems since they rely on legacy systems with specialized operating systems and protocols. Embedded systems are also likely to be included in the exceptions list, as the have limited resources to comply with the general policy. A more comprehensive description of how the architecture treats special considerations is found in section 4.2.
- User and group privilege assignments. Privilege assignments will vary greatly between different industrial control systems, however we would recommend that the highest privilege assignment would be monitoring resources. And that the control resources would be limited to only physical access. This would greatly reduce potential damage from an adversary.
- Device owner assignment. Assigning a device owners is crucial, since it aids the trust inferer in evaluating suspicious activity.

Trust Inferer

Once the data sources have been processed by the device inventory service, they're sent to the Trust Inferer for a trust evaluation. Typical requirements for a high trust level include device encryption, latest OS patches, and that the device execute all management and configuration agents [Osborn *et al.*, 2016]. The exact details of the trust evaluation will likely vary greatly between different industrial control systems. However, from a high level perspective, there are several common practises that should be in place. In general, the requirements should be stricter than for a pure enterprise architecture, i.e. the trust inferer should all ways require device owner assignment. Industrial control systems should also try to leverage the great amount of monitoring data that it has available. By utilizing real-time monitoring data from sensors and controllers, the trust inferer can correlate infrastructure behaviour with device activity. These types of evaluations can detect insider attacks from actors that are usually considered as trustworthy in other hierarchical security models. Further, it is of general importance that the trust evaluation is tailored to the specific ICS subsystems. The vast variety of ICS subsystems, as listed in section 1.1, makes this difficult. To limit the implementation complexity we would recommend that only the critical ICS subsystems are connected to a network. Once the critical systems are implemented and tested in the trust inferer, other systems can be gradually implemented. This would limit the data-flow between the enterprise and industrial zone, but would be crucial in implementing a effective Zero Trust architecture. Section 4.2 will discuss systems that requires special considerations in more detail.

Anomaly Detection

Since trust is a malleable state in the Zero Trust architecture, anomaly detection and user behavior analytics play a key role in the trust engine. Anomalies may remain unknown or undetectable for extended periods of time given the high complexity of the systems involved in critical infrastructure and its various data channels from hardware sensors data, to environmental conditions and valve states. With that being said, the visibility that the Zero Trust architecture provides over the network enables a large amount of data available for inference and advanced threat detection. In addition, this contributes to qualifying risk from an IT network and security perspective.

Access Control Engine

The access control engine is the policy enforcing authority, it gives a binary decision whether or not to allow a given device access to a resource. The access control engine uses the trust evaluation and the data sources forwarded by device inventory service, together with *Access Policy*, to determine whether or not to allow access. Each resource has an associated minimum trust score and other predicates that must be satisfied for successful authorization, these are defined by the access policy.

Access Policy

The access policy defines programmatically the requirements for authorization to access a resource. The access policy defines these rules based on the resources, trust scores, user/group assignments and device owner assignments. In general, the access policy will have to tailored to the needs of a specific ICS. However, the access requirements should always be stricter for lower industrial levels.

4.2 Special considerations for ICS

For most critical infrastructures, ICS are unique in two major ways. First, many ICS tend to be behind critical services such as the electricity grid, gas supply, metro system, etc. Hence, they are high profile targets for nation state attackers. This feature is not shared by many non-ICS networks whose interest is mainly to protect proprietary information and data.

The second unique feature is the necessary presence of OT (Operational Technology) - the infrastructure that does direct monitoring and/or control of industrial equipment, assets, processes and events. While many enterprises might have OT, ICS networks main functionality depends on control of OT. A change in a programmable logic circuit (PLC) by a malicious party could halt a manufacturing plant or turn an air traffic control tower offline. Heavy reliance on such critical infrastructure means that Zero Trust for ICS must focus on protecting these OT devices without affecting

the reliability of the machines and hardware to produce their output.

Legacy Systems

The first hindrance of Zero Trust comes into play when Legacy Systems are considered. Most ICS networks control hardware that is too old to be updated with the latest security controls or systems that cannot handle security controls because of limited device resources. BeyondCorp's Zero Trust implementation puts its legacy systems into an 'Exceptions List' which allows it to override the general access policy. This is low-risk to a company like Google whose focus is their IT systems and their large user-base. However, for ICS the external industrial process is the primary focus, hence generic 'Exception Lists' undermine the security of the ICS network itself.

To cater to this, we want to apply micro-perimeter around such legacy systems as well. However, due to lack of any access management in these systems, we would need to expand the microperimeter to encompass the closest possible Human Machine Interface (HMI) or any similar conduit used to communicate with the legacy system. Allowing such a modified microperimeter will provide a gateway to interact with the legacy system without bypassing the Access Control Policy of Zero Trust.

One such realization of this communication conduit could be the use of API façade design pattern. By leveraging APIs [Feuer, 2017], ICS enterprises can introduce an abstraction layer in front of the legacy infrastructure, making access to such legacy systems subjected to strict authentication mechanisms via the Trust Inferer.

Cloud Services and Third-party vendors

For Zero-trust to fully cover all the possible communication routes, it needs to be integrated with all vendor products and services interacting with the ICS network. This includes all the services migrated to the cloud as well.

We can apply the same Zero Trust concept to the cloud by driving access through a secure gateway for least-privileged access. This is especially necessary since many companies have realized it is more cost-effective to host an application on the cloud, rather than their own data centre.

However, clouds are not used only as Software-as-a-Service (SaaS), but public clouds for data storage and hybrid clouds are used as well. Moreover, it is not necessary for the ICS network to be interacting with one cloud provider but it could be any combination of Microsoft Azure, Google Cloud Platform, Amazon Web Services, or another cloud provider. Hence, it is necessary to know who is accessing their applications and data, what devices are being used to access them, and how data is being used or shared. To monitor and maintain the Zero Trust environment, it is important to map the different types of traffic, continuously inspect them, so that any changes in usage can be adjusted with the Access Control Policy [PaloAlto, 2017].

Industrial IoT

The proliferation of Industrial Internet of Things (IIoT) can be considered a major motivator to move to Zero Trust, since it upends the assumption the Purdue model is based on i.e. the hierarchical nature of data flow. The Purdue model assumes that data flows from level 5 to level 4 and moves incrementally downwards and vice versa. However, the introduction of IIoT devices as sensors to provide more granular data means that there are now possible data flows from level 0 or 1 to level 4, bypassing the middle levels. Some IIoTs send data directly to the cloud too, to be aggregated and analysed. Hence, IIoT stand to fragment the Purdue network, especially if IIoT devices communicate directly to the IT layers.

To accommodate IIoTs without increasing possible vulnerabilities of the system, it is necessary to gain detailed insight into every IIoT device on the network, including its business context and potential for risk. Since IIoTs are coupled with OT devices, we want a zoning solution that is centralized and does not require moving around of bulky industrial gear or re-engineering existing systems. This is provided by Zero Trust through micro-segmentation and defining Software Defined Network, where the control plane is centralised and ensures granular authorization and coordination.

4.3 Example User flow

To illustrate the functionality of our Zero Trust Architecture, we consider this user this user flow example where an engineer's goal is to adjust a setpoint for a PLC at some industrial plant:

1. The request is sent to the access proxy with her credentials, 2FA token and device certificate.
2. The Access Control Engine uses the inventory databases to verify that the credentials are correct and that the device is in an appropriate state.
3. It verifies that the user/device pair is allowed access to the specific resource according to the Access policy and that the request receive a sufficient trust score.
4. The request is finally forwarded to the correct gateway and on its way to the PLC.

4.4 Do we need to restructure Purdue?

Purdue was introduced as an operational model to segment the network into logical components, resulting in an enterprise zone, a Manufacturing/Industrial Zone and the IDMZ which in addition to firewalls provided a point of redundancy by hosting many replication services. The Purdue model was not meant to be a cyber security model, and additions in IDMZ to provide cyber-security was an afterthought.

Historically, ICSs were physically isolated or air-gapped from the outside world, so it was difficult to make a business case for cyber-security given that may were isolated in the past. The focus was on keeping the process running, not cyber-security. The conflation of the*operational* Purdue

Model to a cyber-security model can be considered one of the primary reasons why Purdue is not a good enough defence. The other reason would be the changing attack space and advanced malicious actors. Today, ICS networks are highly targeted which raises the question if we need the hierarchical Purdue model if Zero Trust is deployed effectively?

The answer is both: no and yes.

No. There is no need for the Purdue model as Zero Trust can replace it and provide better protection. Zero Trust authenticates each interaction, whether across zones or within a zone, and the only differentiating factor is the score the Trust Inferer assigns to a particular interaction based on the factors discussed in 4.1. So, there is no need to have an 'inside' and 'outside' network divided by the IDMZ, as under Zero Trust assumption, all network traffic is considered hostile until proven trust-worthy. There are many reasons why the IDMZ can no longer serve its function:

- Provides only 1 level of segregation, hence there is little to no granularity which makes monitoring and logging of very little use.
- The IDMZ does not cater for data flow that is not hierarchical - IoTs, remote access, cloud services might bypass IDMZ and provide attackers a way into the network.
- Some of the components inside IDMZ have migrated to the cloud. Every SaaS app a company deploys and every server they host move externally-facing infrastructure out of the data center and into the cloud
- Provides a central point of failure.

On the other hand, yes, we can still use the Purdue model. But with the caveat that it be used only for logical division of assets for business purposes and streamlining the ICS network's functionality. The Purdue Model and Zero Trust can theoretically work together as long as there is a clear understanding of the purpose of each, and this requires adequate training of the engineers tasked within an ICS. A modern interpretation of the Purdue model would account for the reality of increasing inter-connectivity with Industry 4.0 underway. It would provide a platform of interacting parts with some sort of common data and communications capability that would add the concept of cloud services to capture connections across zones.

5 Analysis and Evaluation

Our analysis is based on an example theoretical attack and how Zero Trust defends the system in comparison to an actual historical attack of similar mechanisms. We show that Zero Trust in superior in not only preventing a malicious actor but also mitigating the resources an attacker has access to, even if she is able to get into the system.

5.1 Historical attacks in ICS

According to an analysis by SANS [Lee *et al.*, 2016] almost 250,000 customers in Ukraine lost their power in 2015 after a coordinated attack on Kyivoblenergo, a power distribution company. The attackers used traditional phishing methods to gain access to computers in the enterprise network and pivoted from there to harvest credentials, escalate privileges and gain company VPN access to the enterprise network. From here they could explore the network for paths into the OT network, and used "native software to deliver themselves into the environment for direct interaction with the ICS components". They achieved this using existing remote administration tools on the operator workstations. The same strategy was used for the 2014 attack on a German steel mill, and is the most common attack vector for ICS attacks [Lee *et al.*, 2014].

Having insufficiently secured remote access software in production is a security issue in itself, but it illustrates one of the core problems of the IDMZ. Because of the trusted inside of the network, the perimeters are the primary line of defence. Achieving strict and secure segregation of the enterprise and industrial networks is often difficult or impractical. And with Industry 4.0 and the increasing amount of OT systems like HMIs, PLCs and sensors that are connected to the network and cloud, maintaining the proper network perimeters are becoming infeasible.

5.2 Security in Zero Trust

At its core, our Zero Trust architecture moves from the idea of a secure network to securing every individual connection. This provides a platform for building secure systems.

Consider an attacker attempting to use the same attack vector as described earlier. The active asset management and device state database immediately makes the attack harder, as the system requires devices to stay up to date with security updates. Let's say the attacker uses some zero-day exploit and get access to a enterprise machine. Until the attacker requires the 2FA token for the user(s) using the device, the attackers opportunity to perform further active attacks is very limited. Even if the attacker acquires the ability to produce valid 2FA tokens, he will only get access to the resources the compromised user/device has access to. If the attacker then begins to probe the network or otherwise transmit irregular traffic (high level of access to resources that employee seldom uses, sending requests at odd hours, irregular modification timestamps on control units etc) it's trust score should decrease and potentially cause a lock-out or a flag to be raised.

5.3 Limitations and challenges

Industrial control systems endure large technical debt where redesigning and deploying new systems can be disruptive and costly. The ability to apply micro-segmentation on legacy systems is one of the major limitations of Zero Trust in ICS. Even if we have modified micro-perimeters as mentioned in section 4.2: Legacy Systems, many of these networks cannot afford to put their systems offline to make the necessary changes. The old protocols in place make the task more challenging with many of these protocols being proprietary and vendor specific, particularly for automation and programmable logic controllers.

Moreover, Peer-to-Peer technologies and mesh networks technology work completely perpendicular to Zero Trust models, as they assume shared access. Role-based access and micro-perimeter controls may not be possible in this environment. If an ICS network heavily relies on mesh technologies, adopting Zero Trust might require a revamp of their architecture. However, with all these challenges, it is more imperative that ICS undergo fundamental architecture and systems changes to put security at the core of their design.

6 Conclusion

The existing perimeter defence architecture adopted in the Purdue model attempts to protect the OT network by separating it from the Internet connected IT network using the IDMZ. However, attackers has repeatedly demonstrated the capability of breaching this barrier, and with the increasing connectivity in IIoT and industry 4.0, the need for rethinking cybersecurity in ICS has become apparent. We believe implementing our architecture, which takes a more defensive and active approach to security with active asset management, granular access control, logging and threat detection, can help improve security in these critical systems. There are challenges and limitations that must be resolved, but with a digitized and automated society, cybersecurity is more important than ever and should be prioritized.

References

[Ackerman, 2017] Pascal Ackerman. *Industrial Cybersecurity: Efficiency Secure Critical Infrastructure Systems*. Packt, 2017.

[Albert W. Jones, 2000] Yuehwern Yih Albert W. Jones, Evan K. Wallace. *Monitoring and Controlling Operations*. National Institute of Standards and Technology, 2000.

[Chiluvuri *et al.*, 2015] N. Teja Chiluvuri, Omkar A. Harshe, Cameron D. Patterson, and William T. Baumann. Using heterogeneous computing to implement a trust isolated architecture for cyber-physical control systems. In *CPSS 2015 - Proceedings of the 1st ACM Workshop on Cyber-Physical System Security, Part of ASIACCS 2015*, 2015.

[cisco, 2017] cisco. Securely traversing iacs data across the industrial ... *Rockwell Automation*, May 2017.

[Feuer, 2017] David Feuer. Multicloud: Taming the rookery, Dec 2017.

[Forrester Research, 2013] Inc. Forrester Research. Developing a framework to improve critical infrastructure cybersecurity. April 2013.

[Lee *et al.*, 2014] Robert Lee, Michael Assante, and Tim Conway. German steel mill cyber attack. *SANS ICS 2014*, 2014.

[Lee *et al.*, 2016] Robert Lee, Michael Assante, and Tim Conway. Analysis of the cyber attack on the ukrainian power grid. *E-ISAC*, 2016.

[Osborn *et al.*, 2016] Barclay Osborn, Justin McWilliams, Betsy Beyer, and Max Saltonstall. Beyondcorp: Design to deployment at google. *;login:*, 41:28–34, 2016.

[PaloAlto, 2017] PaloAlto. What is zero trust for the cloud?, 2017.

[Stouffer *et al.*, 2011] Keith Stouffer, Joe Falco, and Karen Scarfone. GUIDE to industrial control systems (ICS) security. In *The Stuxnet Computer Worm and Industrial Control System Security*. 2011.

[Zhu *et al.*, 2011] Quanyan Zhu, Craig Rieger, and Tamer Başar. A hierarchical security architecture for cyber-physical systems. In *Proceedings - ISRCS 2011: 4th International Symposium on Resilient Control Systems*, 2011.