NextGen Scientific Publication
American International Journal of Computer Science and Technology

Volume 4 Issue 5, Pg. No. 52-60, AIJCST-V4I5P106, 2022
https://doi.org/10.63282/3117-5481/AIJCST-V4I5P106

Original Article

IT/OT Convergence: A Zero Trust Architecture for Energy Sector

***Vilas Shewale**

Independent Cyber security Researcher, USA.

Abstract:

The barrier separating IT from OT has narrowed over the past 10+ years. The rate at which it is happened recently has increased. It is no longer enough for utilities and pipelines / refiners to digitally transmit log data, and they should expect cloud backed analytics, to let engineers adjust parameters using remote control rooms and virtual digital twins of physical assets, which allow them to manage facilities across vast geographies. All that information, the ability to use digital technologies, provides tangible business value to these operations, which results in correspondingly wider exposure of the systems running them. Today's paper examines IT/OT convergence as a business requirement that represents an ongoing security challenge. It discusses and proposes an adaptable zero trust reference architecture comprising five layers-identity, device, network, application and data-that energy sector environments may deploy within existing frameworks to better mitigate cybersecurity threats. We rely on case examples drawn from actual cyberattacks and disclosures over 2020 through mid 2022 to inform this architecture, which also reflects recommendations made by organizations such as NIST, CISA and IEC. Than offering a vendor neutral checklist or prescriptive list of recommendations, the goal of this discussion is to provide a strategic framework within which architects can make decisions about OT/IT convergence as it relates to operational security.

Keywords:

IT/OT Convergence, Zero Trust, Energy Sector, Industrial Control Systems, Micro Segmentation, Critical Infrastructure.

Article History:

Received: 04.08.2022

Revised: 19.08.2022

Accepted: 04.09.2022

Published: 17.09.2022

1. Introduction

IT/OT integration was little more than buzzwords on a PowerPoint presentation barely ten years ago. Today, it is part and parcel of nearly all utility companies and energy operators in North America. Process control centers use both business intelligence dashboards as well as SCADA displays. Refinery Historians dump raw sensor data into cloud-based data analytics platforms. Wind farm SCADA information can stream to the respective manufacturer's diagnostic cloud for proactive maintenance efforts. Resulting improvements in both transparency and operational effectiveness in North America would be hard to imagine 20 years ago. But it comes at the cost of the transparency and insulation that was created between disparate industrial control system platforms, both built for different threat realities.

This paper posits that the most effective security tactic when dealing with industrial control system integration is not to slow or reverse the integration. This phase is well behind the sector and is irreversible. In contrast, this paper asserts that security personnel must build a reference architecture with defensive capabilities in mind, viewing industrial control system integration as an unchangeable constant and carry out Zero Trust principles that take advantage of integration to improve security. Industries are often classified according to how their critical infrastructure falls within broader national framework policies, both motivating nation-state attacks while increasing federal regulatory scrutiny [1]. The energy industry is one of these classified industries. Furthermore, the

sector suffers from a significant long tail of legacy systems that prevent all systems from being regularly patched. Finally, several recent attacks have drawn much-needed attention to the cost of not addressing these issues: in April 2022, the INCONTROLLER attack toolkit was published [2], the Industroyer2 attack against the Ukrainian electric power sector was deployed around that time [3] and pipeline and utility company ransomware threats began before the pandemic and have shown no signs of slowing down.

Section 2 will explain what this means practically and why it is unstoppable. Section 3 reviews the current threat landscape through the first half of 2022, including OT-aware malware systems. Section 4 outlines NIST SP 800-207's Zero Trust model [4]. Section 5 introduces a five-pillar reference framework that illustrates how it could map to the energy sector's needs. Section 6 outlines a framework for implementation and how to execute it, with energy sector applications. Finally, Section 7 covers the limitations inherent in this approach.

A word on scope: this paper discusses architecture and does not get into the details of which specific security vendors to deploy, as such market offerings are vendor neutral and prone to change and the intent is for these frameworks to be used regardless of which vendor suite is employed. Operators reading this should expect to use these guidelines and principles irrespective of their security stock position.

2. The Convergence Phenomenon: What and Why

IT/OT convergence entails the fading divide between systems to manage business and systems to manage an industrial operation. For many years, a widely recognized scheme designated these as separate domains, linked by a slender Industrial DMZ [5]. Although this model was the norm, it is diminishingly relevant as modern operations blur the boundaries, consider Fig. 1, which portrays the reality faced by energy providers.

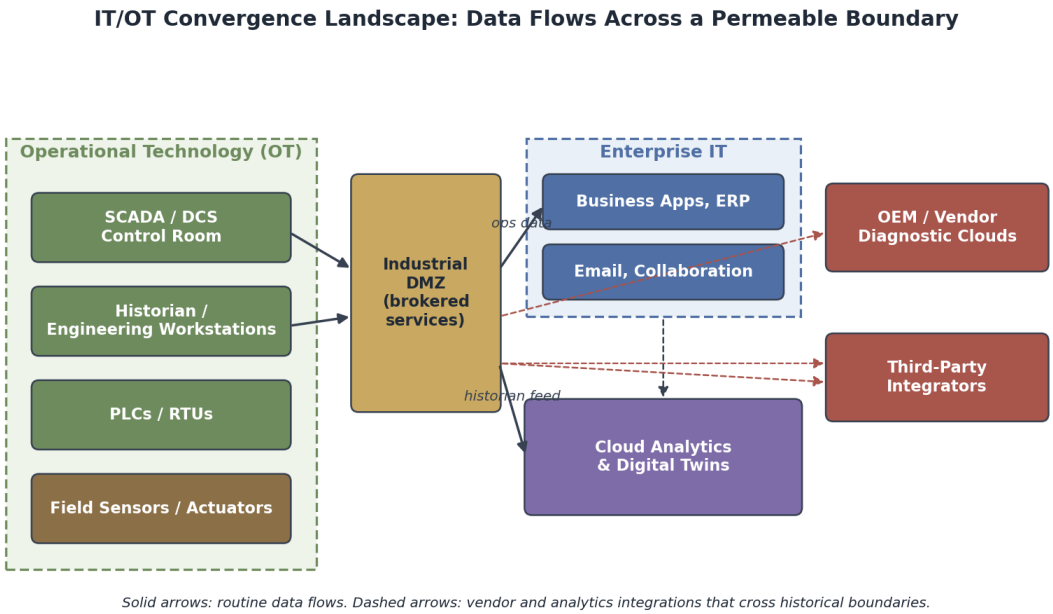


Figure 1. IT/OT Convergence Landscape. Solid Arrows Show Routine Data Flows; Dashed Arrows Show Vendor and Analytics Integrations That Increasingly Cross Historical Boundaries

Cloud Analytics, Digital Twins. Vendors developed fully functional platforms capable of reading historian data and using them to create predictive maintenance solutions, anomaly detection and improve the process. With the cost of a cloud-based solution being much less than an on-premise installation, operators have decided to send their historians outside the control network to the vendor's cloud analytics solutions through brokers near or in the IT/OT DMZ. This is old technology now that sends data outside of the plant.

Industrial IOT sensors. Vendors have brought to market affordable and reliable wireless sensors capable of easily instrumenting assets which have remained historically "dark". A pump with a gauge on it 10 years ago now comes with temperature, pressure, vibration and current sensors all feeding data to the local IoT Gateway. These gateways are usually on the OT network, but the data has to egress the plant to be useful. The gateway operating system usually resembles a general purpose OS than a SCADA or ICS device.

Remote operations centers. Plants and facilities have consolidated their control rooms out of local plant environments to larger regional operations centers. A single operator used to operate a hundred separate pumping stations out in the field, he now operates three different regions plus the national disaster backup from 3 regional centers plus the backup. The WAN infrastructure used for this operational network often runs on the same physical and vendor managed infrastructure as corporate network data.

Workforce realities, the fourth force, play a crucial role. The pandemic hastened the adoption of remote engineering, extending even to OT tasks previously regarded as exclusively hands-on. While some degree of this trend has reversed, a significant proportion of remote engineering access has become permanent [6].

Vendors have their own support models, for example, original equipment manufacturers (OEMs) must be able to remotely communicate with the connected hardware in their deployed customer base to diagnose faults, provide software/firmware updates and handle warranty claims. These support contracts often assume the existence of an IT/OT network boundary connectivity for such communications, details often driven by procurement agents less concerned with the security architecture.

Anyone trying to maintain an air gap soon discovers that it is impossible to do in situations where the stakes are too high or the deadline too urgent, e. g. If the security guard inadvertently leaves their personal USB stick plugged into a PC and leaves it overnight for malicious actors or, if there is a critical vendor bug which needs hot fixing before the window closes during an emergency maintenance period. Reality requires us to plan on the assumption that the boundary is permeable, but ensure that permeable pathways are deliberately configured, monitored and governed by strict policy.

3. The Threat Landscape Shaped By Convergence

Convergence just increases the damage of existing problems by increasing the attack surface after an attack. A hospital phishing scammer getting access to a medical office desktop could, in an unsegmented network, access machines controlling the ventilation or the robots used in surgery. Alternatively, compromised Amazon account where healthcare providers kept their EMR might mean multiple sites got hit in one simultaneous go, either because the account had admin permissions or the same login was used for multiple locations. This came up repeatedly during early-2020 through mid-2022 as there was a surge in attacks.

3.1. OT-Aware Malware Frameworks

In April 2022, the Cybersecurity and Infrastructure Security Agency (CISA) of the U. S. Department of Energy (DOE), the FBI and the National Security Agency (NSA) issued a joint advisory detailing advanced ICS malware, known as INCONTROLLER or PIPEDREAM. This toolkit attacks Programmable Logic Controllers (PLCs) and certain engineering workstations found throughout U. S. Critical infrastructure sectors [2]. Mandiant's analysts reviewed samples of the malware toolkit and discovered several useful functions such as a scanner to probe OPC Unified Architecture (OPC-UA) servers, multiple components designed to interact with controllers at their native protocols and modules that can take down PLCs, trigger their safety logic systems to fail, thereby stopping plant operation.

Coincidentally, ESET researchers joined with Mandiant analysts to present on an Industroyer malware variant that was deployed against an energy utility in Ukraine. The malware is a descendant of that used in the 2016 blackout, although this time, it was tailored to assault a portion of the IEC-104 industrial control protocol that operates across Ukraine (Industroyer was specifically engineered to operate with one of the IEC-101 versions of the standard). In this case, the attack was disrupted, with no large-scale outages resulting. Nonetheless, it made plain that nation-state adversaries continue to invest in developing offensive capabilities tailored for disruptive purposes on power grid equipment.

There are two critical conclusions from these events. The first conclusion is the increasing pervasiveness of OT cybersecurity awareness, with several nations-states displaying significant understanding of ICS technologies and flexible, sharable tools. The second conclusion is a sharpening of the attackers' ability to conduct tailored assaults against targeted industrial control systems. Whereas

Industroyer included more generic functionality, the modern, custom toolkits examined in the advisories discussed here were specialized to exploit known issues on specific PLCs and devices that were believed to be in the attack's primary intended target.

3.2. Ransomware and Supply Chain

In 2021 and early 2022, ransomware attacks were responsible for most major disruptive incidents in industrial organizations according to Dragos' Year in Review [7]. As per the 2022 Verizon Data Breach Investigations Report, human behaviors such as phishing, stolen credentials and errors contribute to approximately 82% of all data breaches [8]. IBM's 2022 Cost of a Data Breach Report showed a new record global average of $4.35 million per breach, with incidents in critical infrastructure costing more than $5.5 million [9].

The second major vector concerns third parties and the software supply chain, so-called third-party risks. When vulnerability CVE-2021-44228-also known as Log4Shell-was made public in December 2021, it also affected industrial companies because Apache Log4j is a key library within so many commercial software products and components, some of which are often used on systems shared across IT and operational environments, as Log4j is also frequently utilized in manufacturing applications. As this log4shell exploitation has progressed throughout 2022, the full extent of its influence continues to become clearer, even as organizations attempt to resolve the many outstanding issues.

3.3 Geopolitical Pressure

Despite that it actually was the day Putin rolled his forces into Ukraine-February 24, 2022, the United States Cybersecurity and Infrastructure Security Agency launched the "Shields Up" program to remind United States Critical infrastructure owners and operators about rising threats. CISA immediately started providing technical advice and alerts, among other useful things, out to its entire constituency. Things did not get better for summer. Russia's geopolitical and security pressures continue and Russia even encouraged non-energy sector operators of critical infrastructure in countries outside of combat with "an elevated risk profile" until it believes things "cool down".[10]

4. Zero Trust: First Principles

The concept of a Zero Trust architecture originated with John Kindervag, then a security analyst at Forrester, in 2010 [11]. The Zero Trust security principle was officially published by the National Institute of Standards and Technology (NIST) in their SP 800-207 guidelines in August 2020[4]. U. S. Executive Order 14028 will require its adoption by the U. S. Federal government by mid-2021 [12], to boost national cybersecurity. Energy companies with a federal contract or operation that are federally regulated are likely within scope for the EO and Executive Order for supply chain reasons, even if they are not named directly in the EO. NIST 800-207 does not respect an implicit assumption of security due to having traversed a network firewall to enter the internal network perimeter. In general, NIST 800-207 has the following tenets:

- All data sources and services are viewed as resources and protected, regardless of location.
- All access to resources is provisioned for a limited duration (per session), based on dynamically policy enforcement and assessed rigidly.
- Security and governance policy measures shall be applied to access for all resources.
- Default communications are expected to be encrypted.
- Identity, device characteristics and location, along with the identity of the device/resource making an attempt, may determine security policy enforcement for accessing the target resource. Zero Trust is therefore more accurately identified as a strategy than a single product.

For OT environments, we suggest that a slightly different approach needs to be considered since 1990s controllers do not readily accept device certificates, a closed control loop cannot accommodate latency inherent in fetching policies remotely over the network, a network function would inherently decrease network ports exposed (increase vulnerability) if introduced at a higher OSI model layer than that expected of an OT network, where ports must remain minimal and protocols specific. Instead, a fortified infrastructure designed for OT needs to be implemented with Compensating Controls where specific tenets do not fit into the standard format. Below is a proposed reference architecture featuring 5 key pillars and specifying what each pillar means for the OT energy sector:

5. A Five-Pillar Reference Architecture

Figure 2 depicts the 5 primary security areas (pillars), with the policy layer managing related security concerns: identities, devices, networks, applications/workloads and data. Information flows from each pillar to the policy engine, which compares the data against relevant corporate policies. The engine then triggers enforcement of specified policies by activating security controls spread across the entire system, highlighting each pillar in an electric utility.

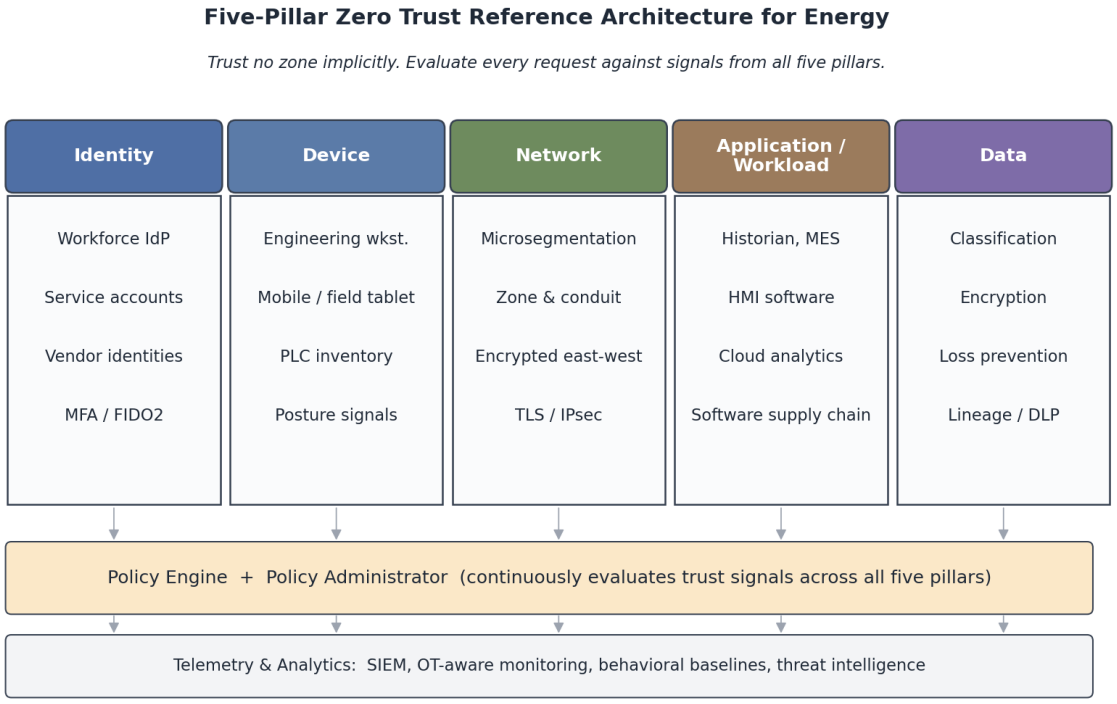


Figure 2. Five-Pillar Zero Trust Reference Architecture. Identity, Device, Network, Application/Workload, And Data Each Contribute Signals To A Central Policy Engine. Telemetry From Across the Environment Informs Continuous Re-Evaluation

5.1. Identity

To get there, identity first represents the easiest, cheapest, first win. We need a unified identity platform to centrally manage all our staff, service accounts and vendors (including where to enable MFA protection for credentials into all privileged domains). A trust posture means that all employees who access sensitive IT assets need to be MFA-enabled, engineers and any privileged accounts must use hardware MFA tokens, we also offer sponsored guest and contractor accounts (with defined lifetimes) who are managed through our Identity and Access management system than their vendor's identities directly.

Another place trust comes into play is in service accounts. In OT (operational technology), those accounts typically hold the passwords machines use to talk to each other, they are known to exist for decades, some accounts often come with excessive privileges and we rarely (if ever) have password expiration policy in place for them.

5.2. Device

The problem is, device covers up all things holding identities in this world: engineering workstations, field mobiles, HMIs, historians, the PLC list itself. IT-tier devices have similar posture attributes familiar to us: patch level, EDR status, encryption status, known good state. The difference that Zero Trust is intended to make is that posture is checked the moment that device is attempting to make its access, not a monthly compliance run. OT-tier devices, on the other hand, cannot even take an agent. So the response to that is basically: one: an accurate inventory of all OT devices and their known issues, including all versions and firmware levels and

two: a strong posture-based decision-making process enforced by the network infrastructure surrounding that OT device. The OT device itself does not have to validate it. The gateway infrastructure is what is responsible for validating the posture.

5.3. Network

Micro segmentation is the core of the network pillar and the IEC 62443 zone and conduit model in OT outlines how to create microsegments: Create areas that are zones with a similar level of security, show traffic flowing from one zone to another through conduits and set the conduit-specific rules of how traffic is permitted in and out of each conduit. The zone boundaries that are created through policy enforcement as well as the zone boundary around safety are explained in Fig. 3.

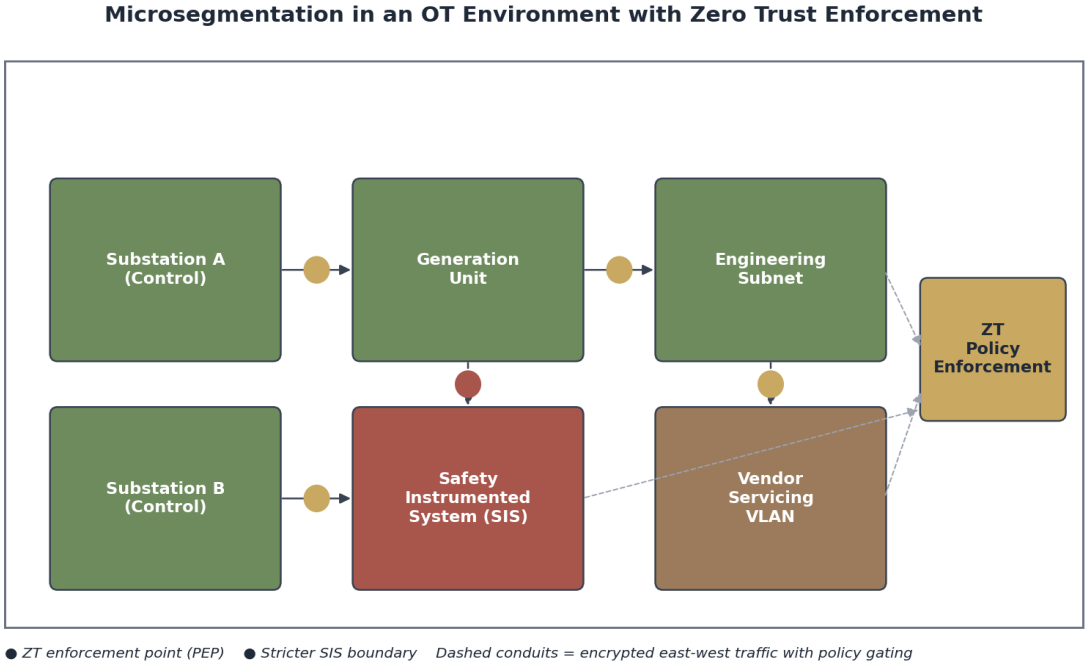


Figure 3. Micro segmentation in an OT environment with Zero Trust enforcement points. Safety-instrumented systems are isolated behind stricter conduit policy; vendor servicing traffic is routed through a low-trust VLAN

East/West traffic is now a must encrypt than optional. The network is not a trust boundary and if a threat actor compromised outside of the perimeter and has reach into the OT systems, encrypting communication between them would prevent them from capturing authentication credentials and process secrets.

5.4. Application and Workload

We start from historian to cloud analytics to SCADA. Everything the operations guy touched from 7am till 10pm should now fall into this. Zero Trust has expectations on MFA into those apps, authorization at a per-resource basis for every user (user x machine) and transparency into that software supply chain. We will hear the term Software Bill of Materials in the OS and IT world, but there will need to be equivalents there. We need vendors to show that we can audit where code came from and make informed risk-based decisions around the patches we use.

Zero Trust needs to be able to extend into cloud and edge environments, which makes this part of the domain "workload protection". My cloud historian that mirrors my on-prem historian in my company cloud should be treated as an identified Zero Trust asset and be hardened and monitored the same way its on-prem brother would be.

5.5. Data

Data usually does not make it into many OT zero trust conversations, partly because IT and OT have different views of it-the former see it as regulated information and the former view it mostly as telemetry. Data from a process control system at a refinery or gas pipeline, though, does have direct commercial consequences and has regulatory implications. When process data must leave OT networks to cloud-hosted applications or services for analytics, it must be properly classified and, if possible, encrypted and tracked for both security purposes and to ensure compliance with rules for industries such as nuclear power generation and oil and gas.

5.6. Policy Engine and Telemetry

Those three pillars are tightly coupled. The security engine relies on threat intelligence feeds from each, the telemetry system provides SIEM/OT with data for historical analysis and to feed back into behavioral baseline calculation. Generally speaking, OT monitoring platforms and SIEMs in energy sector facilities tend to be distinct platforms, though ideally they would sit atop an identical analytical core. True detection requires data correlation across environments, otherwise, how do you reliably detect malicious PLC code, commands outside of PLC operating parameters, logins from a new, anomalous vendor or configuration changes outside an approved window? You cannot detect them without both IT and OT environments contributing the necessary contextual data.

6. Implementation Roadmap

Reference architectures are nice pictures, but actually building a Zero Trust architecture is harder. An organization operating from a typically converged OT/IT network is facing three distinct phases of work.

The first year will focus on visibility and identity. This means you need to inventory all assets, including OT equipment that has not historically been part of IT asset management. Additionally, you must onboard vendor and service accounts into the same identity provider that your workforce uses. You must put multifactor authentication on all of your remote access paths. Get a vendor, ideally an OT-aware one, in place to do network monitoring if you do not already have one. These are not the most glamorous tasks, but they lay the critical foundation needed for anything else.

Once you can clearly see and manage all of your assets and users, over the next 18 to 30 months, you will tackle segmentation and enforcement. Use standards such as IEC 62443 to define zones and conduits across the network. Establish enforcement points at conduits so traffic between zones is evaluated. Look for opportunities to micro-segment parts of the network-beginning with the highest-risk areas such as safety systems. Finally, migrate your vendors' and contractors' remote access to a privileged access solution that your company can manage.

Over the subsequent 3-5 years, you will focus on the data and application pillars. Start to classify process and operational data to understand risks. Require software bill of materials from critical hardware suppliers. Inject policy engine considerations into the procurement and architecture review processes, so new applications and integrations align with Zero Trust principles instead of conflicting with them. Engage executives to reframe security discussions away from compliance posture and toward operational resilience. Two primary factors differentiate organizations that achieve Zero Trust in OT from those that stall out. The first is sequencing. Most organizations fail attempting to carry out all five pillars simultaneously. Zero Trust is not a "Big Bang" project, it requires an evolutionary approach, where foundational capabilities like identity and inventory need to mature before other elements can build on top of them. The second factor is people.

Implementing Zero Trust in an operational technology environment is not a purely technical project, but a profound change in the way engineers, operators and vendors work. Organizations that made the upfront investment in training their OT teams, forging collaborative relationships between IT and OT security personnel and engaging their OT teams throughout the design and implementation phases progressed much faster toward the second horizon than did those that viewed Zero Trust as solely an IT security initiative.

7. Limitations and Open Problems

Let me provide a couple of caveats to the proposed architecture that might not be obvious reading through, but become critical in real-world deployment. These are not "gotchas" that are missed, but architectural trade-offs being made.

First, it does not fully protect legacy assets. If a controller was commissioned in 2005, it does not currently possess a device certificate, nor will it ever install an EDR agent. It is rare to convince an operator to rip out a 15-year-old controller purely for security reasons when it works. The architecture serves to "wrap" legacy assets, than upgrade them to modern security capabilities. Zero Trust aims to make such devices irrelevant (or isolated).

Secondly, real-time control needs are strict with their latency budget. It simply is not possible to slot a policy evaluation in the middle of a deterministic control loop. Trying to shoehorn Zero Trust enforcement into a closed-loop control path is nearly impossible and operators almost universally yank those features out after the first incident of a process being destabilized by the policy evaluation. Zero Trust makes sense to apply at the engineering/management planes, but not at the control path.

Safety overrides all other security concerns. You do not, cannot and should not compromise a Safety Integrity System (SIS) or other safety instrumented systems to enforce security. This architectural choice means separating SIS and safety systems from systems subject to OT zero trust more aggressively. The zero-trust enforcement logic will be applied to systems that interact with SIS, not directly to SIS controllers.

There is also a significant disparity in the maturity of ICS software vendors. Some vendors (e. g. , some from the Siemens TIA portal) have built modern authentication capabilities, complete software bill of materials reporting and runtime protections within their software. Other vendors-those tied to older PLC/DCS lines-do not. To get your infrastructure to "good" (based on what the architecture is trying to accomplish) requires designing to the weakest link (vendors who have not put the effort into securing their software) while encouraging vendors to improve their software through purchasing incentives.

Finally, the threats against OT will not sit still. The OT-awareness described in Sec. 3 represents tooling we did not know existed in 2018 and felt cutting-edge. In 2021, it has become of the baseline. Designing systems based on the 2018 threat environment will have predictable and disappointing results in 2025.

8. Conclusion

IT/OT convergence does not come to a project deadline or to a closure milestone. It is just how industrial work operates now and it is not changing. Out of all industrial sectors, energy bears a huge amount of exposure owing to the sensitivity of the sector's infrastructure, aging installed base and persistent attention from motivated cyber threat groups. Security programs have to assume these conditions hold and accommodate themselves. Zero Trust comes to handy to the degree it works regardless of an IT or OT network setup or a particular set of manufacturers. A company or organization might begin to assess identity, devices, networks, apps, data, all within a framework that relies on continually verifying and testing these assets and then implementing policy, security procedures or systems of checks to make them work together.

Zero Trust provides useful design blueprints that organizations customize according to their needs and priorities, compliance regulations, existing infrastructure, workforce profile or region. What should not be thrown out is the concept of trust as something that is earned for an individual user, device, app or function and is never an inherent trust just because a device is on an OT or IT network or from a trusted brand. At the end of the 2020s, businesses or organizations with good security practices will be those who designed their OT security from scratch for converged IT/OT environments than those who struggled to adapt existing IT solutions.

The current circumstances brought about by a global outbreak created this demand for improvement. The news reports in late February 2023 about the discovery of sophisticated OT offensive operations toolchains eliminated any excuse to defer security improvements. While the task is large, the required course of action is clear and cannot be ignored.

References

[1] U.S. Department of Homeland Security, “National Infrastructure Protection Plan: Energy Sector-Specific Plan,” most recent revision.
[2] U.S. Cybersecurity and Infrastructure Security Agency, Department of Energy, Federal Bureau of Investigation, and National Security Agency, “APT Cyber Tools Targeting ICS/SCADA Devices,” Joint Advisory AA22-103A, April 13, 2022.
[3] ESET Research, “Industroyer2: Industroyer Reloaded,” white paper, April 2022; Mandiant, “Industroyer.V2: Old Malware Learns New Tricks,” April 2022.
[4] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, “Zero Trust Architecture,” NIST Special Publication 800-207, August 2020.
[5] T. J. Williams, “The Purdue Enterprise Reference Architecture,” Computers in Industry, vol. 24, no. 2–3, pp. 141–158, 1994.

[6] World Economic Forum, "Global Cybersecurity Outlook 2022," January 2022.
[7] Dragos, Inc., "Year in Review 2021: ICS/OT Cybersecurity," February 2022.
[8] Verizon, "2022 Data Breach Investigations Report," May 2022.
[9] IBM Security and Ponemon Institute, "Cost of a Data Breach Report 2022," July 2022.
[10] U.S. Cybersecurity and Infrastructure Security Agency, "Shields Up" guidance and advisory series, February 2022 onward.
[11] J. Kindervag, "No More Chewy Centers: Introducing the Zero Trust Model of Information Security," Forrester Research, September 2010.
[12] The White House, "Executive Order 14028: Improving the Nation's Cybersecurity," May 12, 2021.
[13] International Electrotechnical Commission, "IEC 62443: Security for Industrial Automation and Control Systems," multi-part series, 2013–2020.
[14] K. Stouffer, V. Pillitteri, S. Lightman, M. Abrams, and A. Hahn, "Guide to Industrial Control Systems (ICS) Security," NIST Special Publication 800-82 Revision 2, May 2015.
[15] National Institute of Standards and Technology, "Framework for Improving Critical Infrastructure Cybersecurity," Version 1.1, April 2018.
[16] European Union Agency for Cybersecurity (ENISA), "Threat Landscape 2021," October 2021.
[17] U.S. Department of Homeland Security, Transportation Security Administration, "Security Directive Pipeline-2021-01 and Pipeline-2021-02 (revised)," 2021–2022.
[18] U.S. Cybersecurity and Infrastructure Security Agency and Federal Bureau of Investigation, "DarkSide Ransomware: Best Practices for Preventing Business Disruption from Ransomware Attacks," Joint Advisory AA21-131A, May 2021.