Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

Securing OT, IoT, and ICS Networks Using Zero Trust: A Structured and Operationally Safe Adoption Framework

Dr. Aniket Satish Deshpande
Post-Doctoral Researcher, Sunrise University, Alwar, Rajasthan, India

ARTICLE INFO

Received: 01 Nov 2024

Revised: 15 Dec 2024

Accepted: 26 Dec 2024

ABSTRACT

Industrial control networks have slowly become more connected over the years, mostly because it has made day-to-day operations easier. Remote diagnostics, centralized monitoring, and data sharing were added for practical reasons, not because anyone planned a major architectural change. But those small steps have created openings that did not exist before, and many systems still operate as if they were isolated. The problem is no longer just the equipment or the protocols—it is the assumption that everything inside the network can be trusted by default. Zero Trust is useful here, but it has to be introduced in a way that does not interfere with how plants run or how safety is maintained. This paper lays out a gradual approach that starts with understanding what is already happening in the system, then shaping communication boundaries, and only later tightening access based on identity and context, so that security improves without disrupting operations.

Keywords: Zero Trust, Operational Technology Security, ICS Cybersecurity, IIoT, MITRE ATT&CK for ICS, ISA/IEC 62443, Network Segmentation, Industrial Risk Mitigation

INTRODUCTION

Industrial systems run real things in the world. Pumps, breakers, valves, turbines, conveyor lines. When they work, they tend to work quietly and predictably. People trust that. And for many years, the networks that controlled these systems were kept separate on purpose. Not because someone wrote a strategy document about segmentation, but because the simplest way to avoid trouble was to avoid unnecessary connectivity in the first place (NIST, 2023).

That has changed. Slowly at first, and then very quickly. Maintenance teams needed remote access. Vendors wanted telemetry. Management wanted dashboards. Optimization systems are needed to process data from different sites. All of this was reasonable at the time it was introduced. No one made a single, dramatic change; the system just got more connected around the edges (McLaughlin & Skowyra, 2022). And when that happens, the old assumptions about “trusted internal networks” stop holding without anyone noticing the exact moment it happened.

The recent incidents people think about — Colonial Pipeline, Oldsmar — were not sophisticated attacks on PLC instruction sets. They were footholds taken somewhere adjacent to operations. A remote access tool was left open. A shared credential. A monitoring workstation connected to two networks at the same time. It was the pathways that created the problem, not the devices themselves (CISA, 2023). The risk did not arrive because OT changed. The risk arose because the context around OT changed.

Zero Trust gets mentioned a lot here, often as if it were a product or a box to install. But the core idea is simpler: no access is assumed safe just because it exists (Rose et al., 2020). The difficulty is that ICS and OT environments have very real constraints. You cannot patch a PLC the same way you patch a web server. You cannot restart a process historian during active production. And you cannot introduce a control that alters timing on the network without risking process stability (ISA, 2021). So whatever security model is used has to fit the rhythm of the plant, not the other way around.

The point of applying Zero Trust in OT is not to lock everything down. It is to make sure that every connection, every action, and every change in these environments is deliberate, justified, and visible. The work is less about blocking and more about understanding what already happens, and then shaping it in a way that reduces the chance of surprises. Because surprises are the real enemy in industrial systems, once something unexpected happens in a control environment, you are already late.

Journal of Information Systems Engineering and Management
20--, Vol(Iss)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

This paper takes that problem seriously and lays out a staged approach to introducing Zero Trust in industrial networks — starting with visibility, then building toward identity-based enforcement, and finally into adaptive monitoring that respects operational stability. The aim is not to redesign the system from scratch. The aim is to let the system keep running while reducing the amount of silent trust inside it.

BACKGROUND AND STANDARDS LANDSCAPE

Industrial networks did not initially connect to anything else. For a long time, they were built to solve one problem only: make the equipment run safely and predictably. The documentation, design decisions, and even vendor training were all based on the idea that the control network was its own world. No one expected it to share data with ERP systems or cloud dashboards. So the older standards, especially the early automation layering models, were written with that assumption in place (NIST, 2023).

When people refer to the Purdue Model, they often focus on the neat stack of layers, but the more important part is the expectation that each layer has boundaries that should not be crossed casually. In reality, those boundaries have been weakening for years. A historian server pulls data from Level 1 and sends summaries to Level 4. A condition monitoring device forwards analytics to a vendor portal. A contractor's laptop gets plugged into the plant network for configuration. None of these decisions is wrong on its own. They happened because they made sense at the time. But taken together, they changed the shape of trust in the system (McLaughlin & Skowyra, 2022).

Standards evolved to catch up. NIST SP 800-82 tried to describe how segmentation and monitoring should work once the network became a shared environment. ISA/IEC 62443 added something different — it acknowledged that industrial systems are long-lived, patched slowly, and full of components made by different manufacturers. So instead of assuming one organization controls everything, 62443 splits responsibility across asset owners, system integrators, and product vendors (ISA, 2021). This matters because a security strategy that ignores operational maintenance realities will fail in practice.

Zero Trust, as described in NIST SP 800-207, is not about blocking everything. It's about asking a simple question every time something interacts with the system: *Do we know what this is, and is it doing something that makes sense right now?* (Rose et al., 2020). That principle aligns well with industrial control, but the IT-style implementation does not. In IT, you can patch or reboot systems on demand. In OT, you often cannot, because these systems run physical processes that may not have a safe stopping point. So the timing of security controls matters just as much as the controls themselves.

CISA's ICS advisories add one more piece: most real incidents start with credential misuse, shared access paths, or remote access left open longer than intended (CISA, 2023). It is not usually an attacker writing ladder logic directly. It is someone getting into a place they were not expected to be in, because the system assumed they belonged there.

So across all of these frameworks, three practical ideas repeat:

- Know what you have. Not just asset labels, but what talks to what, how often, and why.
- Do not assume internal equals trusted. Most incidents begin "from the inside," but only because the inside was left wide open.
- Security must not break the process. Controls have to fit around the operational tempo, not disrupt it.

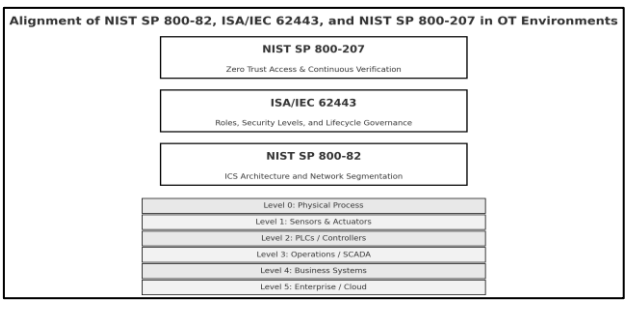

Figure 1: How NIST SP 800-82, ISA/IEC 62443, and NIST SP 800-207 align in OT environments

Journal of Information Systems Engineering and Management
20--, Vol(Iss)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

2.1 Existing OT/ICS Security Platforms

Many organizations begin improving security in OT environments by bringing in a visibility platform. This is mostly because the first problem is almost always the same: no one has a complete, up-to-date picture of what is connected, what it is talking to, and why. Over the last several years, several vendors have developed tools to help address this, each with a slightly different focus.

Some platforms, like Claroty, are built around understanding device roles and protocol behavior. They listen to traffic and try to determine not just who is talking, but what the communication means. This works well in environments with equipment from multiple vendors, especially when the documentation has gaps. Claroty tends to be used early in security programs, when the priority is building an accurate map of the system as it actually operates.

Other platforms, such as Nozomi Networks, focus more on recognizing when behavior changes. Instead of asking "What devices do we have?" the emphasis is on "How do these devices usually behave?" This is useful in plants where the baseline process is stable. Small shifts in timing or message patterns can stand out clearly once they are being watched consistently.

Scadafence approaches the issue from a different angle: it tries to connect OT visibility with the monitoring practices and SOC workflows already present in the enterprise environment. This is helpful when OT and IT teams need to collaborate but use different monitoring tools. The benefit is less about technology and more about shared understanding.

Darktrace's OT module is more behavior-driven. It builds its own sense of "normal" and raises alerts when something does not fit that pattern. It is sometimes used later in the maturity cycle, once the basic environment is known well enough that deviations are easier to interpret.

Then there are platforms like Armis, which are useful at the very beginning, especially in facilities where equipment has accumulated over decades. Armis focuses on identifying devices that people may not even be aware are still present. This can include small wireless gateways, vendor appliances, or older workstations that were never decommissioned.

Large network vendors such as Cisco and Palo Alto now offer OT-focused modules that tie discovery and policy into existing infrastructure. These tend to be adopted in organizations where the network team already has a strong role in managing access boundaries, and where enforcing segmentation through switches and firewalls is more practical than adding new appliances.

Across these platforms, one pattern is clear. They are most effective when they support the maturity stage the plant is currently in. They do not replace the staged Zero Trust approach. They fill in the awareness and control needed at different points in the progression.

- When the environment is not well understood → visibility-first tools make sense.
- When the environment is known but risk is uneven → tools that help define influence boundaries are useful.
- When roles and responsibilities are understood → identity-based controls matter more than new monitoring.
- When the system has settled into a predictable rhythm → behavior-focused analytics become valuable.

Vendor	Primary Strength	Typical Use (Zero Trust Stage)
Claroty	Deep OT asset & protocol visibility	Stage 1 & 4
Nozomi Networks	Behavior & anomaly detection	Stage 4
Scadafence	IT/OT SOC workflow integration	Stage 1 & 2
Darktrace (OT)	Behavioral ML deviation detection	Stage 4
Armis	Unmanaged/unknown asset discovery	Stage 1
Cisco Cyber Vision / Palo Alto IoT	Network-enforced segmentation support	Stage 3

Figure 2. Vendor comparison overview of selected OT/ICS security platforms

Journal of Information Systems Engineering and Management
20--, Vol(Iss)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

The platforms are different, but the problems they address are consistent:

Industrial environments change slowly, but they accumulate complexity quickly. Tools help reveal that complexity so it can be managed deliberately rather than by assumption.

INDUSTRIAL ZERO TRUST ADOPTION MODEL (OVERVIEW)

Zero Trust gets discussed a lot in industrial environments, but it often feels abstract until it is tied to the way a plant actually works day to day. The priority in these systems is simple: the equipment must run, and it must run safely. So any security approach has to respect that reality first before anything else. When you look at where risks show up in most facilities, the issue is rarely that a PLC or RTU is inherently weak. The issue is that there are assumptions baked into the network about which machines should be talking to which others, and those assumptions are often never checked once the system is in operation.

The practical way to apply Zero Trust, then, is not to start by blocking things. It is to start by understanding what is already happening. Many plants have a network diagram that represents what was intended at commissioning, but the actual communication patterns usually look different. Equipment gets added, vendor laptops connect temporarily, a SCADA server gets moved to a virtual instance, and an IIoT sensor gets plugged in because someone needed trend visibility quickly. None of these decisions is harmful by itself. But they create pathways that no one is really tracking. That is what Stage 1 is about: just getting a clear picture.

Once you can see the network for what it is, the next question is: What matters most if it goes wrong? Not every device or data flow is equally critical. A historian going offline is inconvenient. A compressor controller changing mode at the wrong time is not. The risk discussion needs to be grounded in process consequences, not vulnerability scores taken out of context. And if a device cannot be patched or updated because it is tied to ongoing production, then the approach has to shift from "fix the device" to "limit the ways it can be misused." This is where virtual hardening comes in — restricting programming commands to specific workstations, requiring two people to approve control logic changes, and setting certain controllers to read-only during normal operation. These are small controls, but they reduce a significant amount of silent risk.

Only after that foundation is in place does segmentation make sense. And not the kind that simply separates "OT" from "IT." That division is already blurred in most places anyway. Segmentation in a Zero Trust sense is more like: this part of the process can talk to these specific systems, for these specific reasons, under these conditions. It is purpose-based, not subnet-based. It aligns with how operations staff already think about the plant: as functional areas that interact only in defined ways.

The final piece is noticing when something happens that doesn't fit the usual rhythm of the environment. Operators and control engineers already know this intuitively. They know when a set of messages feels out of place or when a device is being accessed at a strange time. Zero Trust just formalizes that intuition into something the system can use — an alert, a slowdown, a request for confirmation — before a mistake or intrusion becomes a physical consequence.

So the progression is not "trust nothing" in the sense of fear or restriction. It is trust what you can explain, monitor the parts that drift, and tighten control only where that drift carries real consequences.

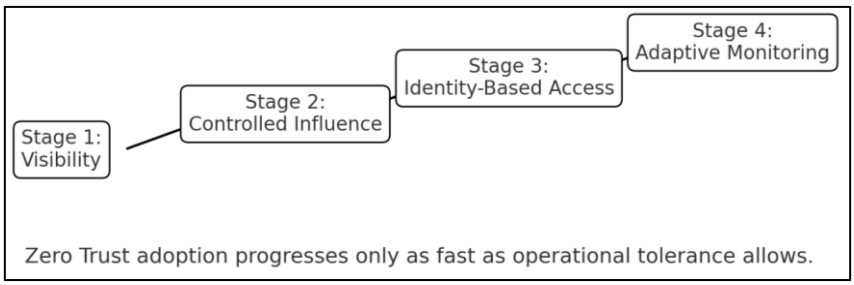

Figure 3: Zero Trust adoption as gradual alignment with operational reality.

Journal of Information Systems Engineering and Management
20--, Vol(Iss)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

ASSET AND COMMUNICATION VISIBILITY — STAGE 1

Most industrial networks do not look the way people assume they do. The documentation may show defined control zones, clear network boundaries, and a clean system hierarchy, but the environment changes faster than drawings are updated. A contractor may connect a programming laptop for diagnostics. A network switch may be replaced with something that has extra ports. A historian may be moved to a virtual server without considering routing changes. Little things, added up over the years, can shift the system into a shape no one has fully mapped.

Stage 1 is about looking at what is really there, not what the plant believes is there. And this is not a one-time exercise. It's a continuous effort, because the system will continue changing as part of normal work.

The visibility effort covers a few things:

- Devices: Not just PLCs, HMIs, and SCADA hosts, but also sensors, gateways, vendor diagnostic tools, wireless bridges, safety PLCs, building controls that share the network, and "temporary" equipment that ends up staying forever.
- Connections: Who talks to whom, and how often. Regular cyclic traffic is different from occasional bursts or configuration commands. That difference matters.
- Purpose: A historian reading data every second is normal. A historian issuing write commands to a controller is not. The goal is to understand the intent of communication, not just the flow of packets.

In a water treatment facility, for example, the PLC controlling chemical dosing will communicate with flow sensors and operator HMIs at regular intervals. If a new system suddenly begins querying that PLC at a faster or irregular rate, the question is not "Is the traffic allowed?" but "Why is it happening now?" In a manufacturing line, if an engineering workstation that is usually idle begins pushing logic downloads to robot cells during production hours, that is not routine activity. The point is to develop a sense of what makes sense in this environment.

This stage does not require blocking anything yet. That part comes later. What matters right now is building a trusted map — a picture of communication that makes sense when you look at it closely. Once that picture is clear, other issues start to reveal themselves on their own:

- Devices that no one remembers installing
- Systems communicating across process areas that should be isolated
- Remote access tools that were added temporarily and never removed
- Engineering laptops that still have cached passwords or trust entries

Most plants discover something unexpected during this step, not because the plant is negligent, but because industrial systems accumulate complexity naturally. This is why visibility is not a security product; it is a habit.

This stage ends when the organization can explain:

- What each device is supposed to do
- Who is expected to interact with it
- When those interactions are expected to occur

When those answers feel obvious and unforced, the foundation is strong enough to move to the next stage.

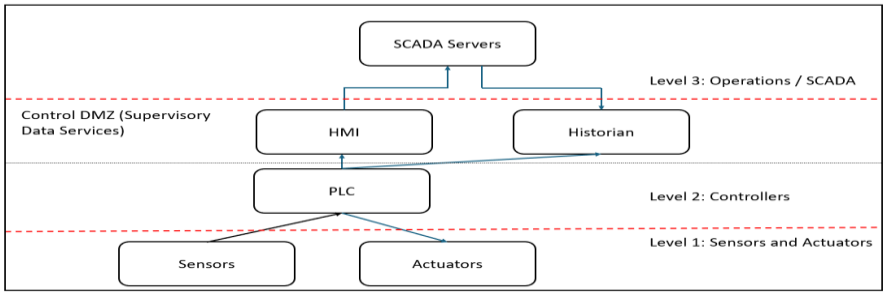

Figure 3: Example of communication baselining in a mixed control environment – Baseline Communication patterns are stable and form expected operational behavior.

Journal of Information Systems Engineering and Management
20--, Vol(Iss)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

RISK CONTEXT AND VIRTUAL HARDENING — STAGE 2

Once there is a clear picture of how the system currently behaves, the next question is not "How do we secure everything?" but "What could cause real harm if it changed unexpectedly?" Not all assets in an industrial environment carry the same weight. Some equipment can malfunction without serious consequences; others sit directly in the path of core process control or safety logic. Understanding this difference is what shapes meaningful security work in OT.

Traditional vulnerability scoring systems like CVSS tend to assume that a high-scoring vulnerability must be remediated quickly. In industrial environments, that is not always realistic. A controller that manages furnace temperature or material feed rates may not be able to be taken offline for patching simply because a scan flagged a vulnerability. The equipment may be operating continuously, with no safe shutdown window available for weeks or months. In regulated environments such as pharmaceuticals or food processing, even small configuration changes may require revalidation cycles that are lengthy and expensive (ISA, 2021). So the real question is not "What is the vulnerability score?" but "What is the operational consequence if something here goes wrong?"

This is where risk context matters more than risk classification. A historian server being out of date is inconvenient but rarely dangerous. A PLC controlling a high-pressure pump, on the other hand, plays a role in both process output and safety. A compromise there is not simply "a security issue" — it is a potential operational hazard. So the evaluation becomes situational, not generic. Different devices matter for different reasons.

Once the critical components are understood, the next move is not to replace them or force patches. Instead, the goal is to reduce how much influence other systems can have on them. This is what we call virtual hardening. It is a way of protecting a system by shaping the interactions around it, rather than altering the device itself.

Examples of virtual hardening include:

- Restricting configuration changes to only occur from specific engineering workstations.
- Forcing operator-confirmation steps before certain setpoint changes can be applied.
- Requiring a supervisor present (physically or digitally) during logic uploads or program edits.
- Setting SIS and protection PLCs to read-only mode under normal operating conditions.
- Time-bounding when maintenance access is allowed — such as only during scheduled maintenance windows, not around the clock.

These controls do not change how the equipment works. They change how the equipment can be influenced.

In a combined-cycle power plant, for example, the gas turbine control logic rarely changes during a run. There is no need for real-time write access. So, configuration writes can be disabled except during planned maintenance. In a manufacturing environment, robot cell programming should not occur during active production. So, even if the network technically allows it, the policy disallows it except during shift turnover or engineering rounds. The patterns are already intuitive to operators. Zero Trust simply turns those intuitions into enforceable conditions.

The key principle is this:

If a device cannot safely be changed, then limit the ways it can be influenced.

Risk is reduced without touching firmware, without rebooting anything, and without introducing downtime.

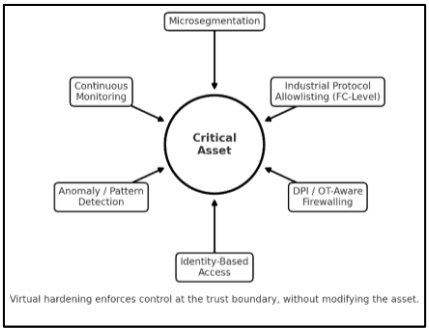

Figure 4: Virtual Hardening Controls Applied Around Critical Devices.

Journal of Information Systems Engineering and Management
20--, Vol(Iss)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

So Stage 2 is not about "fixing vulnerabilities." It is about defining what safe operation looks like and then putting guardrails in place so that the only actions allowed are those that align with that safe state. In industrial environments, security must work with the grain of operations, not against it.

Once the environment has visibility (Stage 1) and influence boundaries (Stage 2), the system finally has a foundation sturdy enough to support identity-based segmentation, which is where Zero Trust begins to take recognizable shape.

STAGE 3 — IDENTITY-BASED SEGMENTATION AND ACCESS CONTROL

Once there is a clear picture of how systems communicate and which components carry the most operational consequence, the next step is to shape who and what is allowed to interact with those systems. In many industrial environments, internal network placement is still treated as the main indicator of trust. If a device sits on the OT network, it is often assumed to be legitimate. If a user has access to a workstation on that network, it is assumed they belong there. These assumptions are convenient, but they are also where most real incidents begin (CISA, 2023).

Segmentation in a Zero Trust sense is not about drawing VLAN boundaries or firewall zones. Those controls help, but they still rely on location-based trust. Zero Trust segmentation asks a simpler and harder question:

What is the purpose of this interaction, and does it make sense in this context?

This shifts the focus from networks to identities:

- Human identities (operators, engineers, contractors, vendors)
- Machine identities (PLCs, HMIs, SCADA servers, historians)
- Service identities (applications, monitoring agents, analytics gateways)

Each identity is given permission to do only what aligns with its role in the process. Not because of where it is, but because of what it is supposed to do.

For human identities, this often means adding conditions to access rather than simply adding credentials:

- An engineer may adjust logic, but only during scheduled maintenance windows.
- An operator may change setpoints, but only within defined ranges, and only from approved stations.
- A vendor may connect remotely, but only when a supervisor is actively aware and the session is logged.

For machine identities, it means recognizing that:

- A historian reads data, not writes.
- A safety PLC controls trip logic, not normal operations.
- A sensor gateway forwards telemetry, not control commands.

Each device has a "shape" to its expected behavior. That shape can be expressed as policy, not just routing.

Now, how this looks in practice differs by industry, but the logic is the same:

- Manufacturing Example

On a production line, robot cells and conveyors follow predictable routines. During active production shifts, no logic uploads or axis reconfiguration commands should occur, even if the workstation has permission to send them. So instead of blocking programming access entirely, the factory limits it to specific hours and ties it to identified engineering roles. The rule is not "don't connect." The rule is "connect only when the plant is in a safe state to accept change."

- Utilities Example

In water treatment and power distribution, SCADA systems operate with steady polling rhythms. Operator actions, on the other hand, occur by shift. So identity-based control can map user permissions to time-bound roles. A remote operator can acknowledge alarms from home overnight, but cannot change control logic unless physically present in the control room. The context — location, time, purpose — becomes part of the authorization decision.

- Oil & Gas Example

Journal of Information Systems Engineering and Management
20--, Vol(Iss)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

In refinery environments and pipeline stations, certain controllers are tied directly to safety instrumented systems (SIS). These devices may technically support programming access, but the expected mode outside of shutdown periods is read-only. So instead of "blocking configuration," the system enforces two-person integrity and explicit override acknowledgment when changes are needed. Not to prevent action, but to ensure action is always intentional.

This Stage accomplishes identity-based segmentation, restoring clarity without reducing flexibility:

- People perform the tasks they are responsible for.
- Devices perform the functions they were designed for.
- Interactions that fall outside those expectations are visible instead of silent.

And the key outcome is practical:

- Security shifts from trusting where something is, to trusting what it is doing — and why.
- No system is locked down.
- No workflows are removed.
- Nothing interrupts production.

What changes is the default assumption: Internal traffic is no longer automatically safe just because it is "inside."

Identity Role	SCADA / HMI	PLC / Controllers	Historian / Data Store
Operator	Monitor Only	Control	No Access
Maintenance Engineer	Monitor + Logs	Full Control (Authorized)	Write Config / Retrieve Data
Vendor / OEM Support	Temporary View Access	Guided Control Sessions	Upload / Patch Only (Supervised)

Figure 5: Mapping identity to purpose-based access in mixed OT environments

CONTINUOUS MONITORING AND ADAPTIVE TRUST — STAGE 4

Once identities and access boundaries are shaped around purpose, the system becomes easier to "read." Patterns start to stand out. Most industrial facilities already operate this way informally. Experienced operators will say things like, "That pump doesn't usually cycle like that," or "I've never seen that workstation access that PLC before." They are noticing behavior that does not fit the usual rhythm of the plant. Stage 4 formalizes that skill into something the system can use.

In an OT network, the vast majority of communication is predictable. PLCs poll sensors, HMIs poll PLCs, historians read everything, SCADA coordinates supervisory control. The timing is steady. The message structure is repetitive. The peers rarely change. Which means abnormal activity does not hide well, once you are looking for it. The challenge is that many environments are not looking — not because they don't want to, but because the tools they have were never set up to monitor for meaning, only connectivity.

Continuous monitoring in a Zero Trust model is not about watching everything. It is about watching what changes, because change is where risk appears.

Some examples of meaningful change:

- A PLC that normally only receives read requests suddenly receives a write request.
- An engineering workstation sends a logic upload during active production.
- A historian begins communicating with devices outside their usual scope.
- A remote access session appears at a time when no one is scheduled to be off-site.
- A device communicates using a protocol it has never used before.

None of these events guarantees malicious intent. But each one is a signal. And industrial security improves dramatically when the environment pays attention to signals instead of reacting only to consequences.

The goal here is not to block every unexpected action. It is to slow the system down just enough to ask whether the action makes sense. Sometimes this means triggering a prompt to an operator. Sometimes it means requiring stronger authentication. Sometimes it simply means logging the action with clarity, so there is a record later.

Consider three real contexts:

Journal of Information Systems Engineering and Management
20--, Vol(Iss)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

- Manufacturing: A programming laptop connects to a robot cell. If this happens during shift change, it is expected. If it happens in the middle of a product run, the system should ask, "Who requested this, and why now?"
- Utilities: In a water treatment plant, setpoint adjustments usually occur in small, predictable increments. A sudden large change is not inherently malicious, but it always warrants acknowledgment.
- Oil & Gas: SIS controllers rarely enter "programming mode." If one does, the system should require explicit operator validation, because the consequences of incorrect logic can be severe.

This is adaptive trust in practice:

- The system raises the level of verification when the situation looks unusual.
- It does not assume every anomaly is an attack.
- It simply treats every anomaly as something worth noticing.

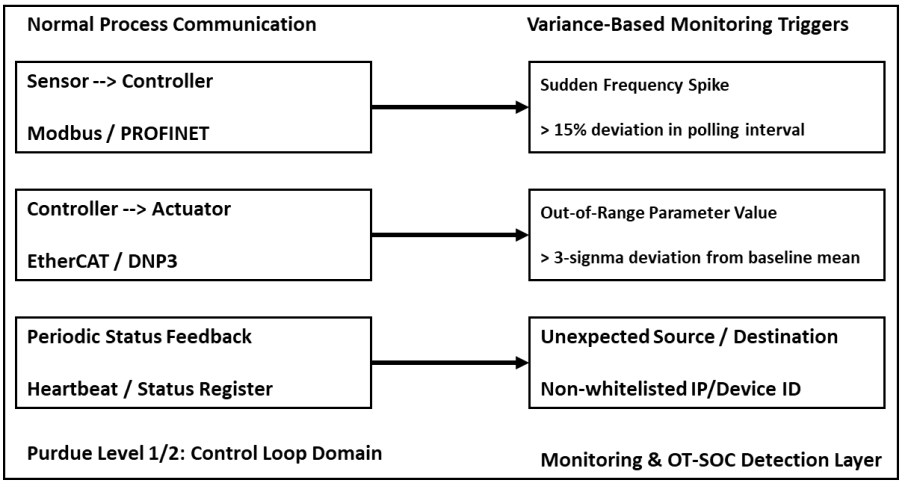

Figure 6: Normal process communication vs. variance-based monitoring triggers

What makes this stage sustainable is that it reinforces how operators already think. They notice when something feels off. Stage 4 gives the environment the ability to notice along with them — and to do it consistently, even when people are overloaded, tired, or distracted by alarms.

This stage closes the loop that began in Stage 1.

You began by learning the system's reality. Now the system has learned to notice when reality shifts. Security becomes awareness, not friction.

ARCHITECTURE MAPPING AND IMPLEMENTATION GUIDANCE

The traditional way of describing industrial network structure is through the Purdue Model, which presents the environment in layers — field devices at the bottom, control equipment above them, supervisory systems higher still, and business or enterprise systems at the top. The model remains useful because it reflects the functional separation between real-time control and higher-level decision systems (NIST, 2023). What has changed is not the logic of the hierarchy, but the way information now moves across it.

Data flows upward more frequently: historians feed dashboards, maintenance systems pull equipment health metrics, and planning applications request performance data. Access also flows downward more than it once did, particularly where remote support or centralized oversight is involved. In many plants, this happens quietly, through incremental additions rather than architectural redesign. Over time, these flows create pathways of trust. Some were deliberate. Some were accidental. Stage 4 made those pathways visible. Stage 5 now asks how to reshape them without interrupting normal operations.

The shift to Zero Trust does not require discarding the Purdue Model. Instead, Zero Trust overlays its principle — trust must be justified, not assumed — across the existing layers. The functional boundaries remain. What changes is that communication across those boundaries becomes purpose-bound. A SCADA server may read values from field

Journal of Information Systems Engineering and Management
20--, Vol(Iss)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

controllers, but not modify their configuration unless specific conditions are met. A historian may query process data, but it does not need write access to anything. A vendor access tunnel may be available, but only while a designated operator or supervisor is aware and present.

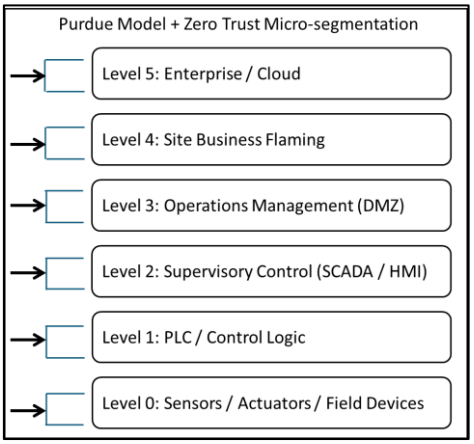

Figure 6: Purdue Model with Zero Trust Overlay – Function Boundaries, Not Location-Based Trust.

This approach means the plant does not need to restructure networks. It redefines how the system understands intent. The network can stay routable. Communication can stay real-time. What changes is the assumption about why a device is communicating, and what it is allowed to do during that communication.

Because most industrial environments contain a mixture of equipment from different vendors, the technology that enforces this segmentation rarely comes from a single platform (ISA, 2021). Some solutions provide deep protocol awareness and communication baselining. Others focus on identity and access management. Others support remote access governance or network micro-segmentation. None of them resolves the full picture on their own.

What ties these components together is not a product, but an approach:

- Define purpose before applying control: Every access rule is tied to what the device or user is actually responsible for.
- Introduce changes gradually: Policies take effect in observation mode first, enforcement later.
- Align enforcement to operational windows: Security adjustments occur when the process is stable, not during active production or peak demand.

These principles apply across different industries, even though the equipment and constraints vary.

- In manufacturing, logic changes are timed around planned production stops.
- In utilities, operator control authority is tied to shift responsibility and physical presence.
- In oil and gas, SIS configurations require multi-party validation because the safety consequences are high.

The implementation strategy remains the same: treat identity and purpose as the basis of trust, not network position.

As the environment transitions into this model, the system becomes easier to reason about. When an unusual action occurs — a change in setpoint behavior, a new communication peer, a configuration command sent at an unexpected time — the context needed to interpret that action is already in place. This is what allows adaptive trust, introduced in Stage 4, to function in a meaningful way.

Zero Trust, applied in this manner, does not tighten the system. It clarifies it. The goal is not to create barriers. The goal is to ensure that every action taking place in the control environment is something the organization can explain, defend, and repeat intentionally.

Journal of Information Systems Engineering and Management
20--, Vol(Iss)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

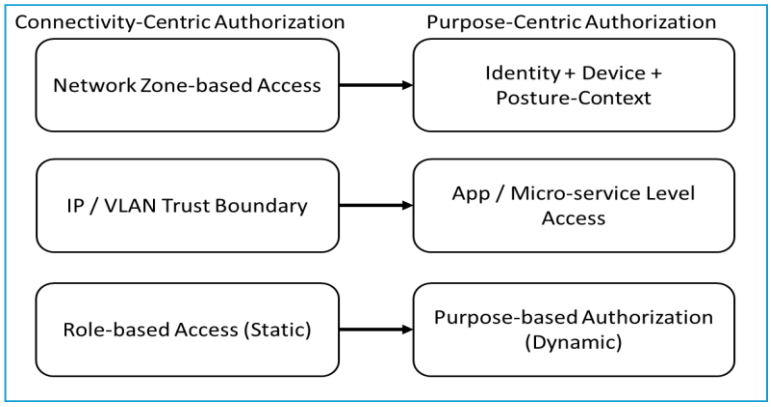

Figure 7: Architecture Transition Path: From Connectivity-Centric to Purpose-Centric Authorization.

CONCLUSION

When industrial systems become more connected, the assumptions that once kept them safe begin to fade. The equipment may still look the same, but the ways in which people and systems interact with it have changed. Remote access, centralized monitoring, and data sharing across networks have brought convenience, but they have also created openings that are not always obvious. These openings are where risk tends to enter.

The approach outlined in this paper does not depend on redesigning networks or replacing control equipment. It works by making the environment easier to understand and harder to misuse. The first step is gaining visibility into what devices are present and how they communicate in practice, not just in documentation. From there, attention shifts to the systems where unexpected changes would matter most, and the goal becomes limiting how those systems can be influenced. Identity-based segmentation replaces network-based assumptions, so access reflects purpose instead of location. Continuous monitoring ties it all together by noticing when something occurs outside the usual rhythm of operations.

The intent is not to restrict work, but to make sure every action can be explained. Industrial systems function best when they are predictable, observable, and deliberate. Zero Trust, applied in this gradual way, supports those qualities rather than disrupting them. The result is an environment with fewer hidden dependencies, fewer silent access paths, and fewer surprises. In systems that must run reliably and safely for long periods, that clarity is the strongest form of protection.

REFERENCES

1. CISA. (2023). *Industrial Control Systems Cybersecurity Year in Review*. Cybersecurity and Infrastructure Security Agency.
2. CISA. (2024). *Zero Trust Maturity Model 2.0*. Cybersecurity and Infrastructure Security Agency.
3. CERT-In. (2022). *Guidelines for Securing Industrial Control Systems and Critical Operational Technology*. Indian Computer Emergency Response Team.
4. CERT-In. (2023). *Critical Infrastructure Threat Advisory Bulletin*. Indian Computer Emergency Response Team.
5. DAE. (2021). *Cyber Security Framework for Critical Power and Nuclear Control Networks*. Government of India, Department of Atomic Energy.
6. ENISA. (2023). *Good Practices for Industrial IoT Security*. European Union Agency for Cybersecurity.
7. Gandotra, E., & Singh, M. (2024). Industrial IoT security challenges and trust models. *International Journal of Industrial Informatics*, 19(2), 88–104.

8. Gill, S., & Sharma, V. (2023). Secure convergence of IT and OT in manufacturing. *Journal of Cyber-Physical Systems Security*, 14(1), 22–37.
9. ISA. (2021). *ISA/IEC 62443 Series: Security for Industrial Automation and Control Systems*. International Society of Automation.
10. Khan, A., & Deshmukh, R. (2023). Industrial Zero Trust maturity adoption models. *IEEE Transactions on Industrial Informatics*, 19(8), 7712–7725.
11. McLaughlin, S., & Skowyra, R. (2022). Threat modeling in critical infrastructure control systems. *ACM Transactions on Cyber-Physical Security*, 5(3), 1–29.
12. MITRE. (2024). *ATT&CK for ICS Knowledge Base*. MITRE Corporation.
13. Moore, C., & White, R. (2023). Attack surface expansion in distributed industrial systems. *Journal of Operational Technology Defense*, 9(4), 55–73.
14. NCIIPC. (2023). *Critical Information Infrastructure Protection Policy Framework*. National Critical Information Infrastructure Protection Centre, Government of India.
15. NIST. (2023). *SP 800-82 Revision 3: Guide to Industrial Control Systems Security*. National Institute of Standards and Technology.
16. NIST. (2020). *SP 800-207: Zero Trust Architecture*. National Institute of Standards and Technology.
17. NIST. (2024). *SP 800-207A: Implementing Zero Trust in Enterprise Environments*. National Institute of Standards and Technology.
18. Nozomi Networks. (2024). *Industrial Threat Landscape Review*. Nozomi Networks Inc.
19. Parekh, N., & Jha, T. (2023). Securing process automation under real-time constraints. *IEEE Industrial Cybersecurity Review*, 12(2), 102–115.
20. Rao, P., & Srinivasan, R. (2022). Digital transformation and cyber resilience in Indian power grids. *Power Systems Journal of India*, 7(3), 44–59.
21. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero Trust Architecture*. National Institute of Standards and Technology.
22. SANS Institute. (2024). *ICS / OT Cybersecurity Survey Report*. SANS Institute.
23. Schneider Electric. (2023). *Secure ICS Lifecycle Modernization Whitepaper*. Schneider Electric Cybersecurity Group.
24. Sharma, P., & Kulkarni, R. (2023). Identity-driven segmentation in refinery automation networks. *Industrial Network Security Journal*, 11(1), 61–79.
25. Siemens. (2023). *Defense-in-Depth for Industrial Control Systems*. Siemens Industrial Cybersecurity Division.
26. Tenable OT Security. (2024). *Risk-Based Vulnerability Prioritization for Industrial Systems*. Tenable Holdings.
27. Tiwari, A., & Goel, S. (2022). Behavior-based anomaly detection in cyber-physical plants. *Journal of ICS Threat Intelligence*, 6(4), 34–49.
28. Wired. (2021). Colonial Pipeline Cyber Attack After-Action Analysis. *Wired Security Review*, 29(6), 12–18.
29. World Economic Forum. (2024). *Industrial Cybersecurity Readiness Benchmark*. WEF Center for Cybersecurity.
30. Vendor & Platform Capability Papers (Cited in Technology Section)
31. Claroty. (2024). *The State of Industrial Network Security*. Claroty Research.
32. Darktrace. (2023). *Self-Learning AI for OT Environments*. Darktrace Industrial AI Division.

Research Article

33. SCADAfence. (2023). *Operational Threat Detection in Distributed ICS Networks*. SCADAfence Research.
34. Cisco Systems. (2024). *Zero Trust for OT with Cyber Vision and ISE Integration*. Cisco Secure Industrial Strategy Division.
35. Forescout Technologies. (2024). *Unified Device Visibility for Enterprise and Operational Networks*. Forescout Research.