

SERIES THREE | CYBERSECURITY



CRITICAL INFRASTRUCTURE SECURITY INSTITUTE

TEXAS TECH
National Security

- Developing cutting-edge solutions to emerging cyber-physical security challenges.
- Advancing the protection, resilience, and security posture of our nation's critical infrastructure systems.
- Real-time cyber-physical simulations, integrating hardware and controllers for cyberattack testing.
- High-resolution 3-D models to detect counterfeit electronics & unauthorized backdoors without physical disassembly.



CRITICAL INFRASTRUCTURE

RESEARCH AREAS



Defense



Telecom



Agriculture



Energy



Water



Healthcare



VULNERABILITY MITIGATION

Cyber-Physical Systems

Understanding, prioritizing and reducing risk in systems where a software flaw becomes a physical consequence.

Aaron Teinert
Cybersecurity Research Engineer

Agenda

Welcome and objectives	Why this topic, and what you will leave with
Part 1 - Foundations	What a cyber-physical system is, and the vocabulary
Part 2 - Why CPS security is different	Safety, uptime and 25-year lifecycles
Part 3 - The vulnerability landscape	Where weaknesses live and how they are exploited
Part 4 - Lessons from real incidents	Five attacks and what they have in common
Part 5 - The mitigation model	Seven layers of practical defense
Part 6 - Making it stick	Standards, governance and a 90-day roadmap
Key takeaways	The five things worth remembering
Questions and discussion	Open floor

OBJECTIVES

What you will walk away with

1

A shared language

You will be able to explain what a cyber-physical system is, how it differs from ordinary IT, and why that difference changes every security decision that follows.

2

A way to see risk

You will be able to locate where vulnerabilities live; in design, in code, in configuration, in the supply chain and in people. Reason about which ones matter.

3

A practical model

You will leave with a seven-layer mitigation model, a prioritization method for when you cannot fix everything, and a 90-day starting roadmap.

One idea to hold onto: in a cyber-physical system, the goal of security is not to protect data. It is to make sure the physical process keeps doing the right thing safely, even when part of the digital layer is compromised.

01

PART 01

Foundations

Before we can mitigate anything, we need a shared picture of what we are defending.

IN THIS SECTION


- What a cyber-physical system is
- Where they show up across sectors
- The layered anatomy of a control system
- The vocabulary you will hear all session



DEFINITION

What is a cyber-physical system?

A cyber-physical system is any system in which computation and networking are tightly coupled to a physical process. The software senses the physical world, makes a decision based on that data, and acts on it, continuously.



SCOPE

You are surrounded by them

Cyber-physical systems are not a niche industrial concern. They are the substrate of modern daily life.

**Energy**
Substations, generation, pipelines, wind and solar inverters, smart meters

**Water**
Treatment plants, chemical dosing, pumping stations, reservoir level control

**Healthcare**
Infusion pumps, ventilators, imaging, pacemakers, hospital building systems

**Transport**
Vehicle ECUs, rail signaling, air traffic, traffic lights, port cranes

**Manufacturing**
Robotic cells, process control, safety systems, warehouse automation

**Built environment**
HVAC, lifts, access control, fire suppression, data center cooling

The common thread: in every one of these, an availability failure or an integrity failure has a physical cost. A blackout, contamination, misdiagnosed patient, collision, damaged plant caused damage that no amount of data recovery can undo.

STRUCTURE

The anatomy of a control system

Most industrial CPS follow a layered reference model (the Purdue model). Attacks travel downward; consequences travel upward.

Level 5-4	Enterprise IT	Email, ERP, internet access, the corporate domain. Where most intrusions begin.
Level 3.5	Industrial DMZ	The buffer zone. Jump hosts, data historians, patch servers. The intended chokepoint.
Level 3	Site operations	Production scheduling, historians, engineering workstations, domain services for OT.
Level 2	Supervisory control	HMIs and SCADA servers. The screens an operator watches and acts through.
Level 1	Local control	PLCs, RTUs, DCS controllers, safety instrumented systems. The logic that runs the process.
Level 0	The physical process	Sensors and actuators. Valves, motors, breakers. Where consequence is real.

GLOSSARY

The Vocabulary

OT

Operational Technology - hardware and software that monitors or controls physical equipment. The counterpart to IT.

SCADA

Supervisory Control and Data Acquisition - the software layer that collects data from and sends commands to remote sites.

HMI

Human-Machine Interface - the operator screen. If an attacker owns this, they own what the operator believes.

SIS

Safety Instrumented System - independent system that trips the process to a safe state. The last line of defence.

ICS

Industrial Control System - the umbrella term for the control equipment inside an OT environment.

PLC / RTU

Programmable Logic Controller / Remote Terminal Unit - small ruggedised computers that execute the control logic.

DCS

Distributed Control System - integrated control for a continuous process such as a refinery or chemical plant.

IIoT

Industrial Internet of Things - newer connected sensors and gateways, often cloud-linked, often outside OT change control.

02

PART 02

Why CPS security is different

The same vulnerability, in a different context, demands a different answer.

IN THIS SECTION

- The priority inversion: safety and uptime first
- Seven structural constraints you cannot wish away
- When cyber risk becomes physical risk



The priority inversion

Traditional IT	Operational Technology
1. Confidentiality Keep data secret. A breach is a disclosure event. 2. Integrity Keep data accurate. 3. Availability Keep services up. Downtime is expensive but recoverable. <i>Typical response to a threat: isolate, patch, restore from backup. Taking a system offline is an acceptable safety move.</i>	1. Safety Nobody gets hurt. This outranks everything, including security. 2. Availability The process keeps running. Stopping a process can itself be dangerous. 3. Integrity Commands and readings are genuine. A false reading is as bad as none. <i>Confidentiality usually ranks last and taking the system offline may be the most dangerous option available.</i>

Practical consequence: an IT team that responds to an OT incident with the IT playbook of disconnect and patch immediately can cause more damage than the attacker. Response plans must be written jointly, in advance.

STRUCTURAL REALITY


Seven constraints you cannot wish away

1 Very long lifecycles Plant equipment is specified for 15-30 years. Systems still in service today were designed before the modern internet threat model existed.	2 You cannot simply patch Patching may need an outage window, vendor approval and safety recertification. Many devices have no patch at all.	3 Protocols with no security Modbus, DNP3, Profinet and others were designed for trusted serial links. Most have no authentication or encryption by default.
4 Fragile to scanning Standard IT vulnerability scanners have knocked over PLCs. Active tooling must be used with extreme care, or replaced by passive methods.	5 Weak or shared credentials Shared operator logins, hardcoded vendor accounts and default passwords remain common, and are hard to change safely.	6 Deep vendor dependency The integrator often holds the only working knowledge of the system, and often holds standing remote access to it.
7 Split ownership Security owns the policy, engineering owns the asset, procurement owns the contract. Nobody owns the risk end to end.		

CONSEQUENCE

When cyber risk becomes physical risk

In IT, the worst case is usually measured in records, downtime and fines. In CPS, the ladder keeps going.



Part 02 — Why CPS security is different

12

03

PART 03

The vulnerability landscape

Where the weaknesses live, and how they get turned into infiltration.

IN THIS SECTION

- What the disclosure data tells us
- Six classes of CPS vulnerability
- Insecure by design is not a CVE
- How intrusions unfold



THE LANDSCAPE IN NUMBERS

What the disclosure data tells us

12,174

vulnerabilities across 3,637 CISA ICS advisories since 2010, affecting 2,783 products from 689 vendors

8.0+

average CVSS severity of ICS advisories in 2024 and 2025. Up from 6.44 in 2010

22%

of OT/ICS vulnerabilities disclosed in 2025 had a matching CISA advisory, down from 58% in 2024

26%

of public advisories shipped with no patch and no mitigation advice at all

What is getting worse

- More disclosures, more severe, and more of them landing on the devices closest to the process
- Controllers, PLCs, RTUs and IEDs are now the most frequently named asset category
- A growing share of flaws never gets a central advisory at all


What that means for you

- A CVE-driven, patch-everything strategy will never converge - the queue grows faster than you can clear it
- You cannot rely on one feed. Vendor and national CERT sources both matter
- Prioritization by consequence, not by CVSS, is the only workable approach

Sources: Forescout Vedere Labs / ICS Advisory Project, ICS Cybersecurity in 2026 (data to 31 Jan 2026); Bitsight, 2026 Global State of ICS/OT Exposure.

TREND


The queue is growing faster than the fix rate



Source: Forescout Vedere Labs / ICS Advisory Project, ICS Cybersecurity in 2026. Intermediate years omitted for clarity; the trend is a steady rise.

TAXONOMY


Six places vulnerabilities live



Note the balance: only class 2 is fixed by patching. The other five are fixed by design, configuration, contracts and governance.

THE CONCEPT THAT CHANGES YOUR METRICS

Insecure by design is not a CVE




Why this matters

- A large share of real OT risk sits in the right-hand column, where your vulnerability scanner reports nothing and your remediation dashboard shows green
- Research programs such as OT:ICEFALL documented dozens of these design-level weaknesses across many vendors at once - not bugs, but assumptions
- This is why "we have no critical unpatched vulnerabilities" is a weak statement about an OT environment, and why architecture and segmentation carry so much of the load

TREND

The attack surface keeps expanding

Every one of these changes delivers real operational value. Each also erodes the isolation the original design depended on.




The air gap, if it ever really existed at your site, has almost certainly been bridged - usually by a project that was approved, documented and delivering value. Assume connectivity and design for it; do not assume isolation and be surprised.

ATTACK PATH

How intrusions unfold

Attacks on physical processes are rarely a single dramatic exploit. They are a patient sequence, and every step is an opportunity to detect and stop it.



STAGE 1 - IT intrusion (standard tradecraft, standard detections apply)

STAGE 2 - OT-specific attack (needs engineering knowledge and time)

Defender advantage: stages 4 and 5 are slow. An adversary who reaches your OT network still has substantial work to do before they can cause a physical effect. This means detection inside OT, even late, still buys you the time that matters most.

04

PART 04

Lessons from real incidents

IN THIS SECTION

- A fifteen-year timeline
- Ukraine 2015 - the first grid blackout
- Triton 2017 - attacking the safety system
- Colonial 2021 - IT ransomware, OT consequence
- Unitronics 2023 - low sophistication, real impact



TIMELINE

Fifteen years of escalation

2010	2015-16	2017	2021	2023-25
Stuxnet	Ukraine grid	Triton / Trisis	Colonial Pipeline	Unitronics / CyberAv3ngers
Purpose-built malware physically damages uranium enrichment centrifuges. Proof that code can break machinery.	Two winter blackouts. The first switched breakers by hand through hijacked operator sessions; the second used automated OT malware.	Malware written to manipulate a safety instrumented system at a petrochemical plant. A safety trip exposed it.	Ransomware on the IT side leads to a precautionary shutdown of fuel distribution across the US east coast.	Internet-exposed PLCs at water utilities compromised through default credentials; campaign later evolves to custom OT malware.

The trajectory: from a single state-built weapon requiring extraordinary resources, to reusable OT malware, to campaigns that need nothing more than a default password and an internet connection.

Ukraine grid, December 2015

What happened

- Months earlier: spear-phishing with a malicious document gave attackers a foothold in the corporate IT networks of three regional distribution companies
- They harvested credentials and quietly mapped the route into the control network
- On the day of the attack, they used legitimate remote access to open circuit breakers at substations by hand, through the operators own systems
- In parallel they wiped workstations, corrupted serial-to-Ethernet converters so breakers could not be reclosed remotely, and flooded the call center to delay reports
- Roughly 225,000 customers lost power. Crews restored supply by driving to substations and switching manually

Why it matters

- No exotic exploit was needed. The attackers used the utilities own remote access and their own HMI software
- The IT network was the entry point; the OT network was the objective
- The attack destroyed the ability to operate remotely and did more lasting damage than the outage itself
- Manual fallback saved them. Ukrainian operators still had the training and the physical means to run the grid without automation

Mitigation lesson: control and monitor remote access into OT as though it were the crown jewels, and keep a rehearsed manual mode of operation. The single most protective factor in this incident was that people knew how to run the system without the computers.

Triton / Trisis, 2017

What happened

- Attackers reached the OT network of a petrochemical facility and spent a long period learning the environment
- They built malware targeting a specific safety instrumented system, the independent controller that trips the plant to a safe state when conditions go wrong
- The apparent goal was to reprogram the safety controller so it would not trip, removing the last automatic protection
- The intrusion was discovered because the safety system detected an inconsistency and shut the plant down - the attackers made a mistake

Why it matters

- This crossed a line. The target was not production or data - it was the protective layer that exists to prevent people being killed
- It demonstrated that adversaries will invest years to reach systems that only matter in the worst moments
- It was found by luck and by safety engineering, not by security monitoring
- The same tooling and approach can be adapted to other safety controllers

Mitigation lesson

Safety systems must be genuinely independent of the control network, not merely logically separated. Keep key switches in RUN rather than PROGRAM, alarm on any change of state, treat every safety controller reprogram as a security event requiring approval, and never accept an unexplained safety trip as a nuisance - investigate it as a possible attack.

CASE STUDY 3 OF 4

Colonial Pipeline, May 2021

What happened

- A compromised credential for a legacy VPN account without multi-factor authentication gave a ransomware group access to the corporate IT network
- Ransomware encrypted IT systems, including the business systems used for billing and scheduling shipments
- The operator shut down pipeline operations as a precaution. The OT systems themselves were not encrypted
- Fuel supply across much of the US east coast was disrupted for days, triggering panic buying and emergency measures

Why it matters

- You do not need to touch OT to stop OT. Losing the ability to bill, schedule or measure can halt a physical operation just as effectively
- The decision to shut down was rational, without confidence in the boundary, continuing to operate was the riskier choice
- One legacy account without MFA was enough. This was not a sophisticated intrusion
- It reframed industrial cyber risk as a national resilience issue and accelerated regulation

Mitigation lesson

Map the IT dependencies your physical operation actually needs, and decide in advance how long you can run without each one. Enforce MFA everywhere, retire dormant accounts, and make sure you have enough network visibility to answer "did it spread to OT?" in minutes rather than days - because that question determines whether you shut down.

CASE STUDY 4 OF 4

Water utilities, 2023 onwards

What happened

- A politically motivated group began compromising a widely used brand of small PLC used by water and wastewater utilities
- The method was almost trivial: the devices were directly reachable from the internet and were still using the vendor default password
- Affected utilities, including small municipal authorities, had to switch to manual operation while devices were replaced
- The campaign continued and matured, with later reporting describing custom Linux-based OT malware designed to blend into legitimate traffic

Why it matters

- Sophistication is optional. Two basic failures of internet exposure and a default credential were the entry point of the attack
- It hit small operators with no security staff and no budget, which describes a large share of critical infrastructure
- Choosing a common product model gave one campaign reach across many independent organizations at once
- Low-skill entry can become a durable foothold as the attacker invests over time

Mitigation lesson

The unglamorous basics still deliver the largest risk reduction per pound spent. Find out what of yours is reachable from the internet, remove it or put it behind controlled access, and change every default credential. If you do only one thing after this session, do that.

SYNTHESIS

What the incidents have in common

1	The way in was already known	Phishing, a stolen credential, a default password, an exposed appliance. The OT-specific skill came later, not at the front door.
2	IT was the route, OT was the target	The corporate network is the on-ramp in the overwhelming majority of cases. IT security is OT security.
3	Legitimate tools, legitimate access	Attackers overwhelmingly preferred the operators own remote access and engineering software over custom exploits.
4	Detection was late or accidental	Found by a safety trip, an operator noticing something odd, or the attacker announcing themselves. Rarely by a security control.
5	Recovery depended on people	Where organizations coped, it was because staff could operate manually and had rehearsed the fallback.

Every one of these is a mitigation you can plan for. None of them requires you to predict which vulnerability an attacker will use.

05

PART 05

The mitigation model

Seven layers, in the order that a real program should build them.

IN THIS SECTION


- Know what you have
- Architect for containment
- Control access
- Manage vulnerabilities realistically
- Detect and monitor
- Respond and recover
- Build and buy securely



OVERVIEW


Seven layers of defense

No single control prevents a determined adversary. The objective is depth so that any one failure is contained, visible and survivable.



FOUNDATION

Layer 1 - Know what you have



Layer 2 - Architect for containment

Segmentation is the control that keeps working when everything else fails and the one that makes un-patchable devices survivable.



Zones and conduits

Group assets that share a security level and a consequence profile into a zone. Every path between zones is a named conduit with an explicit, documented, enforced ruleset. This is the core idea of IEC 62443.



An industrial DMZ

No direct traffic between enterprise and control networks. Everything terminates in a buffer zone - historians replicate outward, jump hosts broker inward. If you build one thing, build this.



Segment within OT

A flat control network means one compromised HMI reaches every controller. Separate by process area, by criticality, and always isolate safety systems.



Prefer pulling to pushing

Data should be pulled outward from OT rather than pushed inward. Where the requirement is genuinely one-way, a data diode enforces it in hardware.



Control the boundary hard

Deny by default. Allow named protocols between named hosts. Review the ruleset on a schedule. Firewall rules decay through project exceptions that nobody removes.



Design for failure

Assume a zone will fail. The test of your architecture is what an attacker can still reach from inside it.

Segmentation is how a 20-year-old unpatchable controller becomes an acceptable risk rather than an open door.

Layer 3 - Control access

Remote access, the priority

- No direct inbound paths into OT. Broker every session through a controlled jump host in the DMZ
- Multi-factor authentication on every route in, without exception for legacy or convenience
- Vendor access enabled on request and time-limited, not standing and permanent
- Record sessions into critical zones; the recording is both deterrent and evidence

Accounts and privilege

- Eliminate shared operator logins where safety allows; where it does not, compensate with physical and session controls
- Change every default and hardcoded credential you can find, and record the ones you cannot
- Separate OT identity from the corporate domain so an enterprise compromise does not grant control access
- Review dormant accounts quarterly. The Colonial entry point was one of these

Physical and local access

- Controller key switches in RUN, not PROGRAM, with position monitored and alarmed
- Locked cabinets and controlled access to panel rooms. Physical access usually defeats every logical control
- Govern USB media and transient devices: block by default, scan at a kiosk where genuinely needed
- Contractor laptops treated as untrusted until inspected

The tension to manage honestly

Every control here adds friction to someone who is trying to keep a plant running, sometimes at three in the morning during a fault. If your access design has no fast, safe emergency path, operators will build their own through a shared password on a sticky note, or a permanently open tunnel. Design the emergency path deliberately, with break-glass accounts that are monitored and alarmed, or you will get an undesigned one.

Layer 4 - Manage vulnerabilities realistically

The first question is not "how severe is it?" but "can it be reached, and what happens if it is?"

Q1	Is the asset reachable?	From where, by whom, over what path? An unreachable flaw on a well-segmented device is a low priority regardless of its CVSS score.
Q2	What is the physical consequence?	If exploited, does the process stop, misbehave, damage equipment or endanger people? Consequence, not severity, drives urgency.
Q3	Is it exploitable here?	Is the vulnerable feature enabled in your configuration? Is there known exploitation in the wild? Many advisories do not apply to your build.
Q4	Can it be patched safely?	Is a fix available, vendor-approved and deployable in a window? Roughly a quarter of advisories arrive with no fix at all.
Q5	If not, what compensates?	Restrict the path, disable the feature, monitor for the behavior, add manual verification. Document the decision and its owner.
The output is not a patch queue. It is a risk register in which every accepted vulnerability has a named owner, a documented compensating control and a review date.		

When you cannot patch

This is the normal case, not the exception. A practical catalogue, roughly in order of effectiveness.

1	Restrict reachability	Tighten the conduit so only named hosts can reach the device on only the required protocol and port. Often reduces risk more than the patch would.	2	Disable the feature	If the flaw is in a web interface, a discovery service or a legacy protocol you do not use, turn it off. Free, immediate and permanent.
3	Harden the path	Put a protocol-aware gateway, proxy or firewall in front of the device so it never sees untrusted traffic directly.	3	Monitor for the behavior	If you cannot prevent it, detect it. Alert on the specific traffic or command sequence the exploit requires.
5	Add procedural checks	Independent verification of critical changes, two-person rules for logic downloads, physical confirmation of set points.	5	Increase physical protection	Key switch position, locked cabinets, restricted panel access. The flaw becomes unreachable without physical presence.
7	Plan the replacement	Where the device cannot be secured, put it in the capital plan with a date. An accepted risk with no end date is a deferred decision.	7	Record the decision	Owner, rationale, control applied, review date. This is what makes an accepted risk defensible to a regulator or a board.

ASSUME PREVENTION FAILS

Layer 5 - Detect and monitor

What to do

- Passive network monitoring on OT segments. Listen on a mirror port, never scan
- Baseline normal, then alert on deviation: new devices, new connections, unexpected protocols
- Process-aware detection: a set point outside operating range, a logic download outside a change window, a safety controller leaving RUN
- Send OT alerts somewhere they will be acted on, with someone competent to interpret them
- Correlate with IT telemetry because stage 1 of the attack happens there

The uncomfortable baseline

Fewer than one in ten OT networks worldwide have network monitoring in place. In roughly a third of recent OT incident response engagements, the trigger for the investigation was a human operator feeling that something seemed wrong.

Operator intuition is a real and valuable detection capability. It is not a strategy, and it does not scale.

The advantage OT gives you

OT networks are far more predictable than IT networks. The same devices talk to the same devices, using the same protocols, in repeating patterns, often for years. Anomaly detection that would drown in false positives on a corporate network is genuinely workable here. This is the one place where the constraints of OT work in our favor.

Remember the kill chain: Stages 4 and 5 are slow. Detection inside OT, even months after the initial breach, still lands well before the physical consequence. This is the deadline that actually matters.

RESILIENCE IS THE REAL OBJECTIVE

Layer 6 - Respond and recover

1 A joint IT and OT response plan

Written together, with engineering holding the authority on any decision that affects the physical process. Isolation is an engineering call, not a security call.

2 Decision criteria agreed in advance

Under what conditions do we shut down? Who authorizes it? Deciding this at three in the morning, mid-incident, produces bad outcomes.

3 Backups (not just files)

Controller logic, device configuration, HMI projects, network device configs, safety parameters. Offline, tested by restoring, and version-controlled.

4 Known-good spares

For critical controllers with long lead times, a preconfigured spare on the shelf turns a months-long outage into a shift-long one.

5 Rehearsed manual operation

Can the process be run safely without the control system? Do current staff know how? When was it last practiced? This saved Ukraine.

6 Exercise it

A tabletop with operations, engineering, security, communications and executives at least annually. The first run always finds something broken.

Ask the question that matters: if the control system could not be trusted tomorrow morning, could you run safely without it - and who knows how?

STOP IMPORTING THE PROBLEM

Layer 7 - Build and buy securely

Most of the vulnerabilities you will spend the next decade managing are being introduced by procurement decisions taken this year.

Put it in the contract

- Security requirements in the specification, not the negotiation. Reference IEC 62443-4-1 and 4-2 capability levels
- A software bill of materials for every delivered system, kept current
- Committed vulnerability disclosure and patch support for the asset lifetime, not the warranty period
- No hardcoded or shared credentials; individual accounts and configurable authentication
- Defined remote support model, brokered and time-bound, agreed before installation

Manage the supply chain

- Know which vendors and integrators hold access to which systems, and review it
- Verify update integrity through signed firmware, verified sources, no updates over unauthenticated channels
- Assess integrator security practices; their laptop becomes your network
- Track component-level risk through the SBOM so you can answer "are we affected?" in hours
- Watch for monoculture: the same product across every site multiplies one flaw

The leverage point


A control system bought today will still be running in 2045. Every security property it lacks becomes twenty years of compensating controls, monitoring and accepted risk. The cheapest security work you will ever do is a requirement written into a specification before the purchase order is raised - and the most expensive is retrofitting the same property afterwards.

Practical first step: add a short, standard security annex to your OT procurement templates, and require the SBOM. It costs nothing and it changes what vendors offer you.

PRIORITIZATION

You cannot fix everything

Start from consequence, not from the vulnerability list. Work backwards from the handful of physical outcomes you genuinely cannot accept.



06

PART 06

Making it stick

Programs fail on ownership and follow-through far more often than on technology.

IN THIS SECTION

- Standards and regulation
- Governance, roles and metrics
- A 90-day starting roadmap



STANDARDS AND REGULATION

You do not have to invent this

IEC 62443

The primary international standard for industrial automation and control system security. Defines zones and conduits, security levels, and requirements for asset owners, integrators and product suppliers.

Use for: architecture, procurement, maturity targets

NIST SP 800-53 & 800-82

Practical US guidance on securing operational technology. Readable, comprehensive and a good starting reference for building a programmed.

Use for: programmed design and control selection

NIST CSF 2.0

A general risk-management framework. Includes govern, identify, protect, detect, respond, recover. Useful common language between IT, OT and the board.

Use for: governance and executive reporting

Sector regulation

NIS2 in the EU, TSA security directives for US pipelines and rail, NERC CIP for the bulk electric system, and medical device requirements from regulators.

Use for: knowing what is mandatory for you

Pick one as your backbone rather than mixing four. For most industrial asset owners that is IEC 62443, with the NIST framework as the language you use to report progress upward.

Ownership, roles and measurement

Who does what

- A single named accountable owner for OT cyber risk, senior enough to arbitrate between production and security
- Engineering and operations own the assets and hold the veto on anything affecting the physical process
- Security owns policy, monitoring, threat intelligence and incident coordination
- Procurement owns the contractual requirements from Layer 7
- A standing joint forum where all four meet monthly, with a decision log

Metrics worth reporting

- Percentage of OT assets in a maintained inventory, with consequence ratings
- Number of unmanaged paths between enterprise and control networks - target zero
- Percentage of remote access sessions brokered, authenticated with MFA and recorded
- Percentage of critical OT segments with monitoring coverage
- Backups of controller logic tested by restore in the last 12 months
- Date of the last joint IT and OT incident exercise

Metrics to be careful with

Patch compliance percentage and count of open critical CVEs both look rigorous and both mislead in OT. They reward activity on the vulnerabilities you can reach rather than risk reduction on the assets that matter, and they cannot represent the insecure-by-design weaknesses that carry much of the real exposure. Report them if you must, but never as the headline.

The most common failure mode of an OT security program is not a missing control. It is that nobody with authority owns the outcome, so every difficult decision is deferred until an incident makes it for them.

WHERE TO BEGIN

A 90-day starting roadmap

For an organization starting from a low base. All of it is achievable without major capital expenditure.

Days 1-30

Establish the facts

- Name the accountable owner and stand up the joint IT/OT forum
- Run the two-hour unacceptable-outcomes workshop from the prioritization slide
- Find everything of yours that is reachable from the internet, and remove or broker it
- Inventory and change every default credential you can safely change

Days 31-60

Close the obvious gaps

- Begin a passive asset inventory at your most critical site
- Map and document every path between enterprise and control networks
- Enforce MFA on all remote access into OT; disable dormant accounts
- Verify you hold tested offline backups of critical controller logic

Days 61-90

Build the durable capability

- Design the target zone-and-conduit architecture, starting with the industrial DMZ
- Deploy passive monitoring on one critical segment as a proof of value
- Run a joint IT and OT tabletop exercise and record what it breaks
- Add the security annex and SBOM requirement to OT procurement templates

Nothing on this page needs a business case. Everything on it makes the next, larger business case far easier to write.

KEY TAKEAWAYS

Five things worth remembering

- | | | |
|---|---|--|
| 1 | The consequence is physical, so the objective is different | You are not protecting data. You are making sure the process keeps doing the right thing safely, even when part of the digital layer is compromised. |
| 2 | Patching is not the strategy, and never will be | The disclosure queue grows faster than anyone can clear it, a quarter of advisories ship with no fix, and much of the real risk is insecure design with no CVE at all. |
| 3 | Attacks come in through the ordinary door | Phishing, stolen credentials, exposed remote access. The OT-specific skill is applied later. That means IT security is OT security. |
| 4 | Architecture and prioritization are what scale | Segmentation makes an unpatchable device survivable. Consequence-driven prioritization gives you a finite, defensible list of what to protect first. |
| 5 | Resilience is the honest goal | Assume something gets through. Detect it inside OT where the adversary is still slow, and rehearse operating without the system you no longer trust. |

Questions?

Thank you.

SERIES FOUR | COMING SPRING 2027



DEPTS.TTU.EDU/CISI
CISI@TTU.EDU

APPENDIX A | SLIDES 21-25

Sources for the incident case studies

Stuxnet, 2010 (slide 21)

Symantec Security Response, W32.Stuxnet Dossier - Falliere, O Murchu and Chien, v1.4, February 2011
https://archive.org/details/w32_stuxnet_dossier

Ukraine grid, December 2015 (slides 21-22)

E-ISAC and SANS ICS, Analysis of the Cyber Attack on the Ukrainian Power Grid: Defense Use Case, 18 March 2016
<https://nsarchive.gwu.edu/sites/default/files/documents/3891751/SANS-and-Electricity-Information-Sharing-and.pdf>
Underlying US government report: DHS ICS-CERT Alert IIR-ALERT-H-16-056-01, February 2016, cited throughout the above.

Ukraine grid, December 2016 (slide 21, referenced on 22)

ESET, Industroyer: Biggest threat to industrial control systems since Stuxnet, 12 June 2017
<https://www.welivesecurity.com/2017/06/12/industroyer-biggest-threat-industrial-control-systems-since-stuxnet/>
Dragos, CRASHOVERRIDE: Analyzing the Malware that Attacks Power Grids
<https://www.dragos.com/resources/whitepaper/crashoverride-analyzing-the-malware-that-attacks-power-grids/>
CISA alert: <https://www.cisa.gov/news-events/alerts/2017/06/12/crashoverride-malware>

Triton / Trisis, 2017 (slides 21, 23)

Dragos threat group profile, XENOTIME - the activity group behind TRISIS
<https://www.dragos.com/threat/xenotime>
MITRE for the US government, Cyber Risk to Mission Case Study: Triton (public release 22-3170)
<https://apps.dtic.mil/sti/trecms/pdf/AD1183008.pdf>
Originally disclosed jointly by Mandiant (TRITON) and Dragos (TRISIS) in December 2017; CISA calls it HatMan.

Colonial Pipeline, May 2021 (slides 21, 24)

CISA and FBI Joint Cybersecurity Advisory AA21-131A, DarkSide Ransomware, 11 May 2021
<https://www.cisa.gov/news-events/cybersecurity-advisories/aa21-131a>
Legacy VPN without MFA - CEO testimony to the Senate Homeland Security Committee, 8 June 2021
<https://www.cybersecuritydive.com/news/colonial-joseph-blount-ransomware-legacy-vpn/601523/>

Water utilities, 2023 onwards (slides 21, 25)

FBI, CISA, NSA, EPA, INCD, CCCS and NCSC, Joint Advisory AA23-335A, IRGC-Affiliated Cyber Actors Exploit PLCs
<https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-335a>
Continuation of the campaign against other PLC vendors: CISA Advisory AA26-097A, April 2026
<https://www.cisa.gov/news-events/cybersecurity-advisories/aa26-097a>

Sources for the disclosure data

Vulnerability counts, severity trend and advisory coverage

Forescout Vedere Labs with the ICS Advisory Project, ICS Cybersecurity in 2026: Vulnerabilities and the Path Forward - data from March 2010 to 31 January 2026. Primary source for the 12,174 / 3,637 / 2,783 / 689 figures, the 2011-to-2025 comparison, the CVSS trend and the 22% advisory-coverage figure.
<https://www.forescout.com/blog/ics-cybersecurity-in-2026-vulnerabilities-and-the-path-forward/>
Independent coverage of the same report, useful if the original moves: Infosecurity Magazine and Industrial Cyber
<https://www.infosecurity-magazine.com/news/industrial-control-system-vulns/>
<https://industrialcyber.co/threats-attacks/forescout-flags-spike-in-high-severity-ot-ics-flaws-exposing-visibility-gaps-that-leave-critical-infrastructure-at-risk/>
Primary feed the report is built on - CISA ICS advisories: <https://www.cisa.gov/news-events/ics-advisories>

Monitoring coverage and detection triggers

Bitsight, 2026 Global State of ICS/OT Exposure. Source for the under-10% monitoring coverage figure, the 30% of incident-response cases triggered by an operator, and the 26% of advisories shipping with no mitigation (slides 14 and 34).
<https://www.bitsight.com/blog/2026-global-state-of-ics-ot-exposure>

Attack paths, ransomware and adversary behavior

Dragos, OT Cybersecurity Year in Review. Source for the share of OT incidents originating in IT, ransomware-driven shutdowns, and adversaries operating inside OT at Stage 2 (slides 7, 19 and 26).
<https://www.dragos.com/ot-cybersecurity-year-in-review>

Note: all three are vendor research reports, each drawing on its own telemetry and customer base, so treat them as directionally reliable rather than precise.