Industrial Cyber

Industrial Cybersecurity Critical Infrastructure Handbook 2027



DRAGOS DISPEL AYRON Macrium Software

Chapter 1 Protect the Value, **Not Just the Asset**

Chapter 2 Who Actually Owns **the Decision?**

Chapter 3 Threat Hunting Beyond **the OT Boundary**

Chapter 4 When AI Enters the **Control Loop**

Chapter 5 Remote Access: **Same Risk, Different Rulebooks**

Chapter 6 When Regulation Meets **Operational Reality**

Chapter 7 Backed Up, But **Can You Recover?**

Chapter 8 Keeping Essential Services **Running After Compromise**

Chapter 9 Making Residual **Risk Defensible**

Chapter 10 Secure by Design Is a **Shared Responsibility**

DRAGOS DISPEL VYRON Macrium Software



Intro

Critical infrastructure organizations cannot protect everything equally, and the services they deliver cannot simply stop. Every security decision therefore has to answer a basic question: what does it protect?

At Takepoint Research, our OT Cyber Value Protection Framework starts from a simple position. Industrial cybersecurity exists to protect value: safety, production stability, resilience, and integrity. The order in which decisions are made matters. Value comes first, then the consequence of losing it, and only then the capabilities needed to protect it.

This handbook draws on the sessions and discussions at Industrial Cyber Days: Critical Infrastructure 2026, where practitioners from power, water, oil and gas, manufacturing, and other sectors across APAC, EMEA, and the Americas shared what has worked and what has not. Despite different regions, regulations, and levels of maturity, they kept returning to the same starting point.

The ten chapters follow that logic. They begin with what matters and who owns the decision, move through visibility, AI, remote access, and supplier obligations, and close with recovery, continuity, defensible risk decisions, and shared responsibility for secure products. Each chapter stands on its own and ends with practical takeaways. If this handbook helps you make even one better-informed security decision, it has served its purpose.

Jonathon Gordon, Directing Analyst,
Takepoint Research

Research

Protect the Value, Not Just the Asset

Moving from threat-driven cybersecurity to protecting operational outcomes.



Introduction: Why fear stops working

Fear helped put OT cybersecurity on the agenda and, in many organizations, secured an initial budget. However, it does not sustain that investment. Continued funding depends on demonstrating what the money protects, which risks it addresses, and what improvements the organization can reasonably substantiate. Fear is not a strategy.

Threat-driven programs also tend to produce the wrong kind of control. A mindset built around keeping attackers out adds equipment and restrictions that operators experience as obstacles, and the people who run the process learn to work around them. Leaders who hear repeatedly that the danger is severe, without seeing which operations are at stake, have little basis for judging whether the next request is necessary.

The alternative starts from the other end: the services an organization delivers, the outcomes it cannot afford to lose, and the systems and people that keep those outcomes intact.

Systems, not assets

Asset-centered security pulls teams into silos. Each device becomes a separate problem with its own vulnerabilities and owners, and the conversation narrows to the item in front of the team. Viewing assets as a collection that works together to deliver value changes the questions. What is that value, what is it worth, and what would it take to disrupt it?

A Windows workstation illustrates the point. As an inventory entry, it is a commodity endpoint. In practice, it may be part of a distributed control system that the plant cannot run without, and its importance comes from that role.

Port operations show how far the dependencies reach. A container terminal relies on a terminal operating system, SCADA, cranes, application interfaces, and yard systems that are often managed separately but only deliver value together. What the terminal actually protects is the flow of cargo. If a crane stops, vessels are delayed, the yard congests, and ships queue at the gate, which can create regulatory problems of its own. The effects extend to other ports and vessel operators in the same ecosystem. Protecting the crane as an asset misses most of that picture.

Starting from what cannot be accepted

A consequence-led program reverses the usual order. Instead of beginning with an asset list and a vulnerability scan, it begins with the impacts the business cannot accept and works backward:

- Define the outcomes that must be protected, such as the flow of a service or the output of a plant.
- Identify the dependencies that sustain them, including assets, communications, and third parties.
- Establish who is responsible for each of those dependencies, so accountability doesn't stop at the plant boundary.
- Assess exposure against those consequences, including how long an outage can last before the impact becomes unacceptable, and design controls around them.

Duration belongs in the assessment. The case for a redundant system depends on what happens if the primary is unavailable for five or ten minutes, and that answer differs from one process to the next.

Consequences also differ between facilities. A generation outage can ripple into grid stability and downstream supply, while a refinery shutdown can bring production, safety, environmental, and supply-chain impacts at once. Even with similar technology, acceptable interruption periods and operational priorities rarely match.

Controls that fit the process

Consequence also determines whether a control is appropriate. Most people accept that a banking app locks them out after three failed password attempts. Few operations managers would accept the same behavior from an operator workstation during a shutdown, when someone under pressure mistypes a password. In OT, an identity control or a poorly timed patch can become a self-inflicted denial of service.

The analogy cuts the other way as well. Nobody would use a banking app over open public Wi-Fi without encryption, yet many industrial environments have historically run with security settings turned down to the minimum. The aim is the right setting for the process, avoiding both underinvestment and controls that disrupt the operation they are meant to protect. Well-designed controls make correct operation the easy path, so that people fall into good practice almost by default. Much of what is labeled security turns out to be good engineering and good system management.

The analogy cuts the other way as well. Few people would use a banking app over open public Wi-Fi without encryption, yet many industrial environments have historically run with security settings turned down to the minimum. The aim is the right setting for the process, avoiding underinvestment and controls that disrupt the operation they are meant to protect. Well-designed controls make correct operation the easy path. Much of what is labeled security turns out to be good engineering and good system management.

Frameworks help when they are applied as concepts. The Purdue Model provides a useful reference for understanding industrial architectures, although real environments rarely follow its layers precisely. IEC 62443 offers a more flexible approach through zones and conduits, but its application still requires an understanding of the specific environment. Distributed utilities, for example, must account for remote sites, wide-area communications, and external dependencies that differ from those of a centralized factory. Teams that copy a reference diagram, or accept an AI tool's average answer, miss what is specific to their own system.



A shared language built on consequence

Each function describes risk in its own terms. Security teams talk about threats and vulnerabilities, network teams about connectivity, engineering and safety teams about keeping equipment running, operations about keeping the business running, and management about financial, reputational, and statutory exposure. Teaching every group the others' vocabulary is less effective than giving them one common language, and consequence provides it: if this system is down, what stops, for how long, at what cost, and with what safety implications?

A common language still needs checking. 'OT' means different things to different people, and the same words carry different meanings in factory automation and utility SCADA. A conversation can run for half an hour before it becomes clear that two people define a key term differently. Stating which systems, networks, and environments a term covers, and asking simple questions early, prevents that. Plain explanations also carry weight when budgets are decided.

A common language only works if people establish it before they need it. In many organizations, the first time security, operations, and network teams sit together is at the incident table. Framing this as IT versus OT only deepens the divide, since both sides want the same thing: safe, secure, and reliable organizational operation.

Earning the next dollar

Investment cases that succeed are framed in the business's terms. Security will rarely increase production, but it can protect against the loss of production, and the potential loss can often be estimated. Risk assessments expressed in financial terms, with a phased



plan that addresses the most serious risks first, give management something it can evaluate. Many organizations look for the option that balances moderate risk against moderate cost, and a consequence-based case shows what that balance actually buys.

The case continues after approval. Leadership will ask what an investment has achieved, and programs need metrics that provide credible evidence of improvement, reported consistently over time.

Sponsorship at the top remains decisive. The story that works helps senior leaders see their responsibility for keeping the organization operating. For organizations with hazardous processes, the most accessible route is often to integrate cyber risk into established safety and risk management practices, while recognizing that deliberate cyberattacks introduce threats that conventional equipment reliability assessments may not fully address.

Establishing who is responsible for each dependency is where many programs stall. Security cuts across organizational divides, with OT often split between a CIO or CISO on one side and operations and engineering on the other. Once an organization knows what it must protect, the harder question is who holds the authority to decide how.

Key takeaways

- Fear can secure initial attention, but sustained investment depends on showing what is protected, which risks are addressed, and what improvement can be substantiated.
- Prioritize by the operational flows and outcomes that assets support, since an asset's importance comes from its role in the wider system.
- Begin with the consequences the business cannot accept, then work backward to dependencies, owners, exposure, and controls.
- Judge every control against the process it protects, avoiding both minimal security settings and measures that cause self-inflicted outages.
- Use consequence as the common language across security, engineering, operations, and management, and define key terms before relying on them.
- Present investment in financial and operational terms, phase it by risk, and use metrics as credible evidence of improvement rather than precise measures of eliminated risk.

Who Actually Owns the Decision?

Accountability when ownership, operational control, and cybersecurity responsibilities are divided.

Introduction: Divided control, single accountability

Knowing which consequences matter does not settle who acts on them. Critical infrastructure is rarely run by a single organization. Asset owners rely on operating companies, OEMs, system integrators, maintenance contractors, and service providers, each controlling a different piece of the environment. The party that detects a problem may not be able to isolate the equipment involved, and the party maintaining a system may not be the one that bears the consequences of taking it offline.

Regulation sharpens the problem without resolving it, as the entity stays accountable to the regulator even when the people with hands on the equipment sit in another organization. European regulatory frameworks, including the EU's NIS2 Directive and Critical Entities Resilience Directive, and the UK's NIS Regulations establish obligations for entities within their respective scopes. These include incident notification requirements, although the applicable triggers, procedures, and responsibilities differ. Outsourcing operational activities does not, by itself, remove those regulatory obligations, so an entity can remain accountable for events it has limited power to act on. The practical question is where decision rights sit, and whether they are agreed before an incident or improvised during one.

Ownership shifts across the lifecycle

The division between owner, builder, integrator, operator, and maintainer becomes visible as soon as someone proposes a security improvement. A maintenance contractor may agree that a site needs better endpoint protection, perimeter controls, or secure remote access, yet be unable to act because its contract covers maintenance for a fixed term and nothing more. The request returns to the owner, the risk analysis has to be reopened, and the change waits on a decision nobody planned for.

Construction adds another layer. A large facility may spend several years being built, with many companies connecting systems, aligning configurations, and establishing external connections. A zones and conduits architecture may exist in the design documents, but when the plant is handed to the team that will run it, someone has to verify that systems actually communicate as designed. If the party commissioning the plant is not the party operating it, that verification can fall between them.

Time widens the gap. Technology selected at the design stage may be obsolete by commissioning, and its replacement was never in the original budget. Most parties will agree the work is necessary. The open questions are who performs it, who funds it, and who repeats the risk analysis.



Shared assessment, named decision

Assigning ownership of a risk requires knowing what assets exist, what they depend on, and who is responsible for them. Without that knowledge, ownership is nominal. Effective governance defines the asset owner, the operator, and the executive sponsor for the business line in which each asset sits.

The security function typically acts as a liaison between the business and the plant, and its recommendations are only as good as the information it holds. The most reliable view of what could go wrong comes from the people who run the site: which process could fail, what a vendor connection exposes, which production lines would stop. Every facility carries out - dated systems and known vulnerabilities, but not every one produces operational impact, so assessment emerges as a collective exercise. Because not every cyber risk produces op - erational impact, assessment is a collective exercise. It should involve the asset owner, the operator of the essential service where that is a separate organization, operations and engineering, and the system integrator that built the control environment. Where continuity is at stake, management belongs in the decision.

Collective assessment still needs a named decision-maker. When cybersecurity and operations disagree about whether a plant should keep running, someone must have the final word, and that authority should be written down. Security teams identify the risk and make sure those accountable for the operation and the value it produces understand what they are accepting or managing, because the decision itself usually rests with them.

Contracts, funding, and the limits of transfer

Contracts can allocate operational duties and commercial liability, but they cannot transfer statutory obligations imposed directly on a regulated entity. Some organizations consult their contracts for the first time after a breach. More mature ones bring partners into incident response planning and program design, with legal input where needed, so that roles are understood on both sides in advance.



Older agreements are a particular weakness. An OEM contract signed years ago may say almost nothing about security. Treating contracts as living arrangements, reviewed against current requirements and regulation instead of being left until renewal, gives the owner a way to close that gap. Some larger organizations have gone further, helping smaller critical suppliers meet new obligations with practical and financial support, on the basis that a weak supplier is a weakness in their own operation.

Funding exposes the same fault lines. A risk raised by a service provider has to be reanalyzed and funded before anyone can address it. In unregulated settings, the owner decides whether to accept the risk. In heavily regulated sectors such as electricity distribution, investment may depend on tariff and subsidy decisions that require regulatory or government approval, and a weakness rooted in aging technology may affect thousands of sites at once. Water follows a similar pattern where public bodies own infrastructure that private companies operate and maintain. The interval between identifying a risk and funding the remedy can stretch over years, which has led some in the sector to argue that new cybersecurity obligations should come with provisions for funding them.

Regulation can also break deadlocks. Where every party in a supply chain agreed that security mattered but none wanted to carry the cost,



placing a clear obligation on one entity has given it reason to bring its partners along.

Testing who decides

Decision rights that exist only on paper rarely survive an incident. Crisis management needs defined thresholds for when an incident becomes a crisis, built on risk assessment, business continuity planning, and an escalation plan. Interdependence raises the stakes, since a weak point in one supplier can disrupt operations well beyond its own boundaries. Exercises are where these arrangements are proven.

Simulations show whether procedures work, whether technology performs as expected, and who actually owns each decision. An IT security operations center may detect and coordinate the response to an incident, but preventing operational harm also depends on plant personnel, engineering expertise, and clearly established authority to intervene. Exercises need to test coordination between field staff, regional teams responsible for widely distributed sites, central offices, the SOC, and subcontractors who may have security staff but none on site. Communication and leadership deserve as much attention as technical response, and simulation tools, including digital twins, can make the scenarios more realistic.

Organizations at any stage of maturity can make progress within a quarter:

- Confirm which regulatory requirements apply and what notifications they demand.
- Build a responsibility matrix stating who decides on risk prioritization, who advises on cyber risk, and who owns operational risk.
- Map key processes and systems to the assets they depend on and the owners of those assets.
- Align incident response and regulatory communication with that matrix, then test it.
- Extend the same clarity to suppliers through contract review and supply chain assurance.

Leadership commitment and a culture that treats cyber risk the way the industry learned to treat safety make these arrangements last. If a tabletop exercise cannot establish who makes the call, a real incident is unlikely to either.

Every one of these decisions depends on knowing what is actually in the environment. Unmanaged connections, shadow IT, and now unsanctioned AI tools mean that many organizations are assigning responsibility for systems they cannot fully see.

Key takeaways

- Map ownership, operation, integration, and maintenance roles for each site, because the party able to act is often not the party accountable for the outcome.
- Assign cybersecurity commissioning and handover responsibilities explicitly, so that design intent is verified before operations take over.
- Treat risk assessment as a shared exercise with operators, engineers, and integrators, and document who holds final authority when security and operations disagree.
- Review supplier contracts against current security and regulatory requirements instead of waiting for renewal or an incident.
- Build funding and approval cycles into remediation planning, particularly in regulated sectors where investment depends on external approval.
- Test decision rights through exercises that include field teams, regional and central functions, and subcontractors, not only the SOC.

Threat Hunting Beyond the OT Boundary

Why the assets outside traditional OT visibility are critical to defending the process.

Authored by Dragos

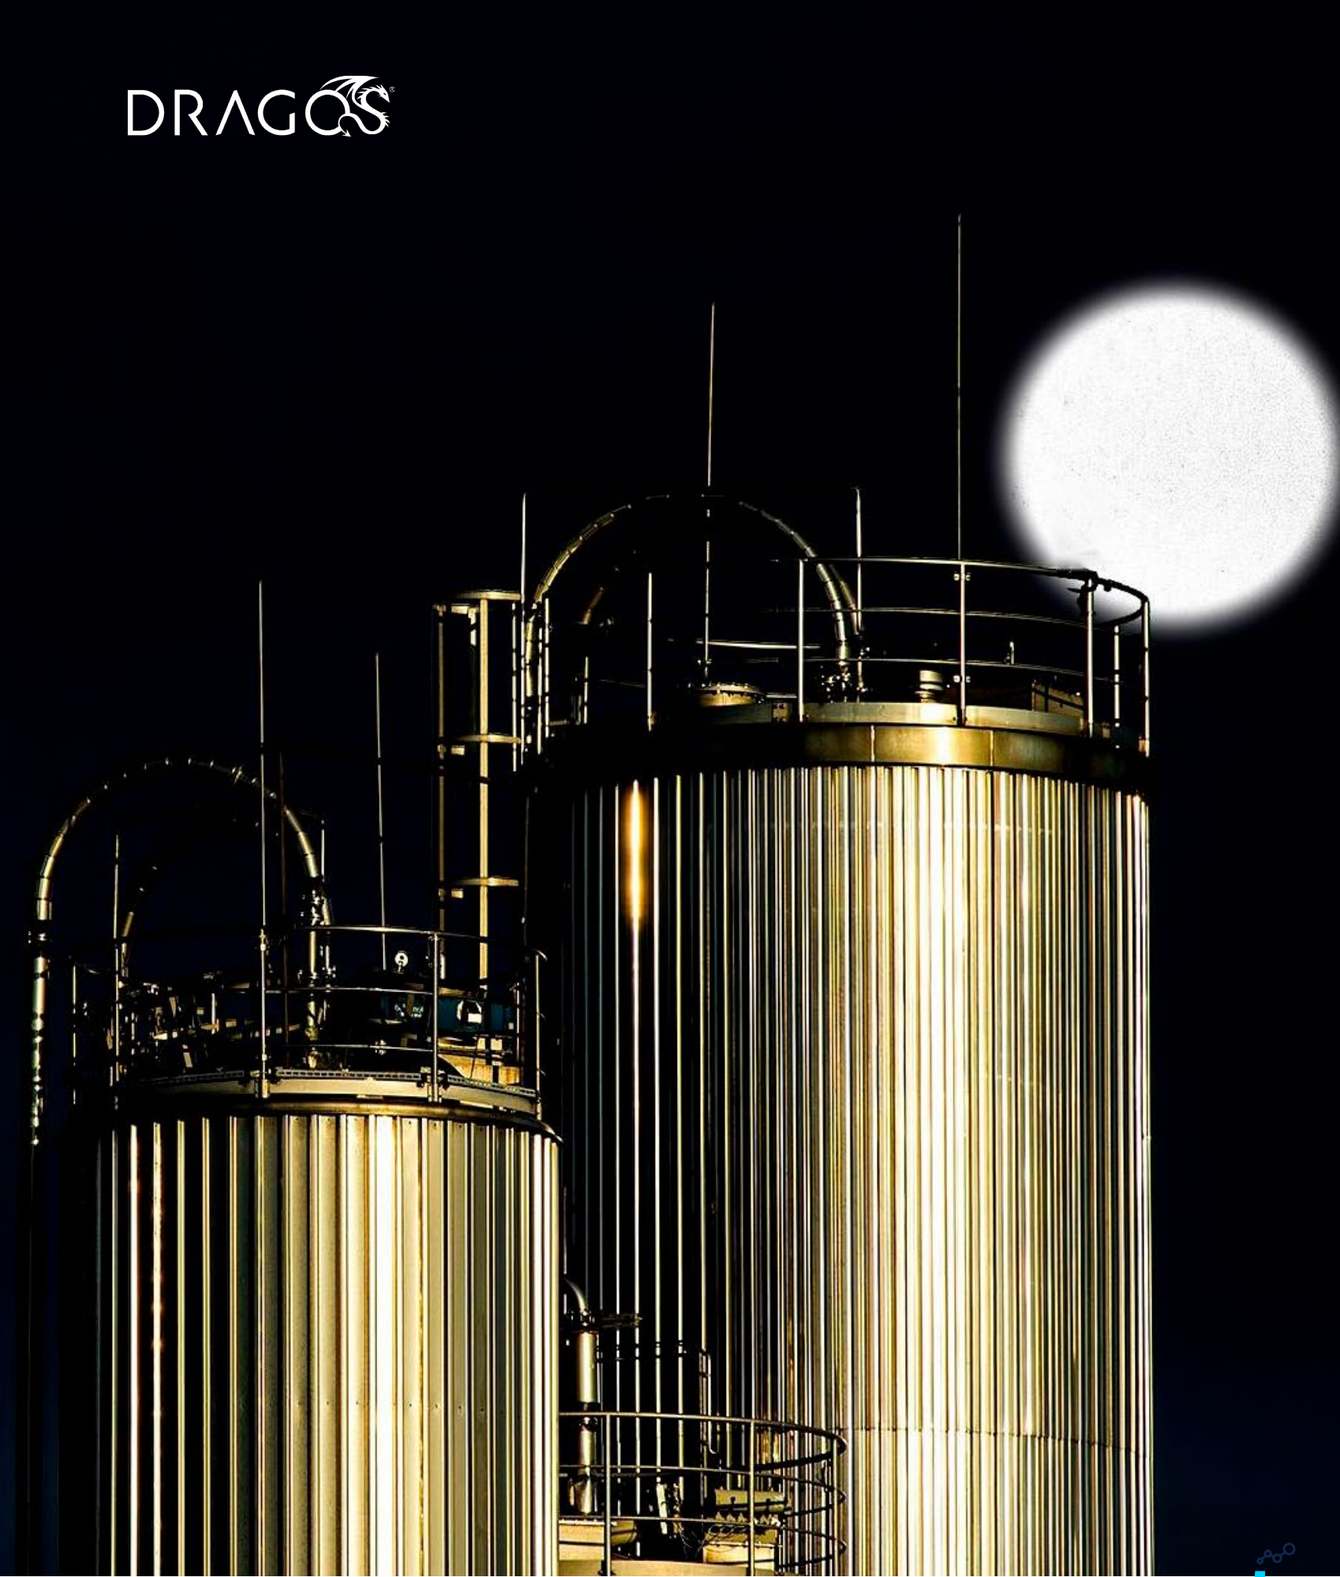

Introduction: The hunt that came back clean

Every industrial security leader is eventually asked: Are we more secure today than we were yesterday? And what can we do today to be more secure tomorrow? Most answers describe purchases or projects underway. A credible answer has two parts: evidence that something measurably improved, and a specific next step that will improve it further.

Consider a well-resourced team that stands up threat hunting. Its first hunts across controllers, HMIs, engineering workstations, and historians come back clean. Leadership is reassured. Yet the hunt never covered the building management system holding the temperature the process depends on, the vendor remote access path, the environmental sensors feeding the control system, or the cameras and edge devices sharing the same network. None of them were in the inventory, so none of them were hunted. The hunt did exactly what it was asked to do. The confidence was built on an incomplete inventory.

Defining the extended OT environment

Most programs cover the traditional OT core reasonably well: PLCs, RTUs, HMIs, safety systems, engineering workstations, and historians. The environment that sustains production is wider. In a manufacturing plant, a production order may flow from ERP to the line with no person in the loop. MES coordinates production, machine vision inspects every unit and feeds quality data back into the process, a vendor logs in remotely to tune robots, the building management system maintains environmental conditions, and an industrial label printer determines whether finished goods can ship. Any of these can stop production, not because a controller fails, but because the conditions the process needs no longer exist.

This extended OT, or xOT, environment can be defined with a simple test. Does the system influence a physical process, or does it drive a high-impact operational outcome? If so, it belongs in scope. Systems whose loss is damaging but does not halt operations, such as analytics platforms that only store history or corporate email, sit outside it.

How adversaries see the same environment

An asset inventory typically lists IP addresses and device types. Adversaries build something different: a model of relationships. A VPN gateway leads to an engineering workstation that can reach a PLC. ERP feeds MES, which feeds the line. A camera nobody thinks about can reach the network shares. The model shows them what normal looks like and where the least effort produces the greatest physical effect.



The threat group Dragos tracks as VOLTZITE, which overlaps with activity publicly reported as Volt Typhoon, has been observed collecting GIS data, OT network diagrams, and SCADA configurations, as well as investigating conditions that could trigger process shutdowns. That is closer to an operating manual than a device list, and it raises an uncomfortable question: does an adversary understand the environment better than its owner does?

Ransomware operators exploit the same gaps. Dragos has tracked 628 Akira ransomware incidents since the start of 2025. In one publicly reported case, attackers entered through a compromised VPN and attempted to run an encryptor on Windows hosts, where its endpoint detection and response (EDR) blocked it. They then found a Linux-based camera on the network that could not run an EDR agent and sat outside central management, and used it to encrypt network shares. The camera never touched the process; its position on the network was enough. The same applies to any connected device running default credentials, outdated firmware, or expired certificates, because fixing them manually across thousands of devices is slow, and adversaries exploit that gap.

Dragos's 2026 OT Cybersecurity Year in Review reports 119 ransomware groups targeting industrial organizations in 2025, up from 80 the year before, alongside an average OT dwell time of 42 days. A program that monitors only part of its environment detects only part of the activity within it, however capable its tools are.

A tooling problem or a sequencing problem

When capable teams fall short, the cause usually falls into one of two categories. The first is tool fit. IT security tools are pointed at environments they were never built for: EDR agents that cannot run on embedded devices, scanners that are unsafe to use against controllers, and patch cycles that cannot keep pace. No amount of process maturity fixes a tool that does not fit the environment.

The second is the sequence. Organizations advance detection, hunting, and response well ahead of the inventory those capabilities depend on. The capabilities at the top under-deliver because they can only see what the foundation provides. Telling a tooling gap from a sequencing gap shows where the next investment should go.

A progression in six levels

A maturity progression addresses the sequencing problem when tools built for xOT environments, not IT tools retrofitted into them, are in place. Each level answers one question and builds specific SANS ICS Critical Controls:

- Discover: What do we actually have? (SANS Controls 2, 3)
- Prioritize: What do we fix first? (SANS Control 5)
- Fix: How do we fix it at scale and prove it worked? (SANS Controls 2, 4, 5)
- Detect: Is this a real threat? (SANS Controls 3, 4)
- Hunt: What are we missing? (SANS Control 1)
- Respond: What do we do next? (SANS Control 1)

Each level depends on the one before it. Nothing can be prioritized before it is found, fixed at scale before it is prioritized, or detected where nobody is looking. The progression is also a loop. Response covers far more than incidents: an alert, an unscheduled configuration change, a shifted segmentation rule, or a newly disclosed vulnerability. Each response feeds back into discovery.

Every level produces evidence and a next move. Discovery shows fewer unknown devices than last week; an inventory that cannot list every device with an expired certificate or a password unchanged for 30 days lacks the detail it needs. Prioritization depends on recognizing that only a small share of OT vulnerabilities require immediate action, and on maintaining a “Now, Next, Never” list that shrinks over time. Where patching is not possible, an alternative such as a configuration change, network control, or detection should be documented rather than left to wait for a distant maintenance window.

Fixing is where many programs stall, because remediation is still done one device at a time. The strongest evidence of progress is validated change: rotated credentials, updated firmware, renewed certificates. The move: schedule the next remediation job across the fleet. The camera in the Akira case was a way in only because nobody was managing it, which makes it a Fix gap. Detection must cover remote access paths and native industrial protocols, because attacks can look like normal engineering work, and detections should map to the threat groups targeting the sector.

The move: close the next coverage gap, starting wherever discovery just found something new. In the 2024 FrostyGoop attack involving a Ukrainian municipal energy company, attackers manipulated controller-reported temperatures, and the heating system failed, leaving hundreds of apartment buildings without heat in the middle of winter. No controller was disabled. A hunt succeeds when its findings change a detection, a hardening standard, or an architectural decision. Everything a hunt can find is bounded by discovery: it searches only what the inventory includes, which is why hunting has to extend across the entire xOT environment. The move: plan the next hunt from what threat intelligence says is coming.

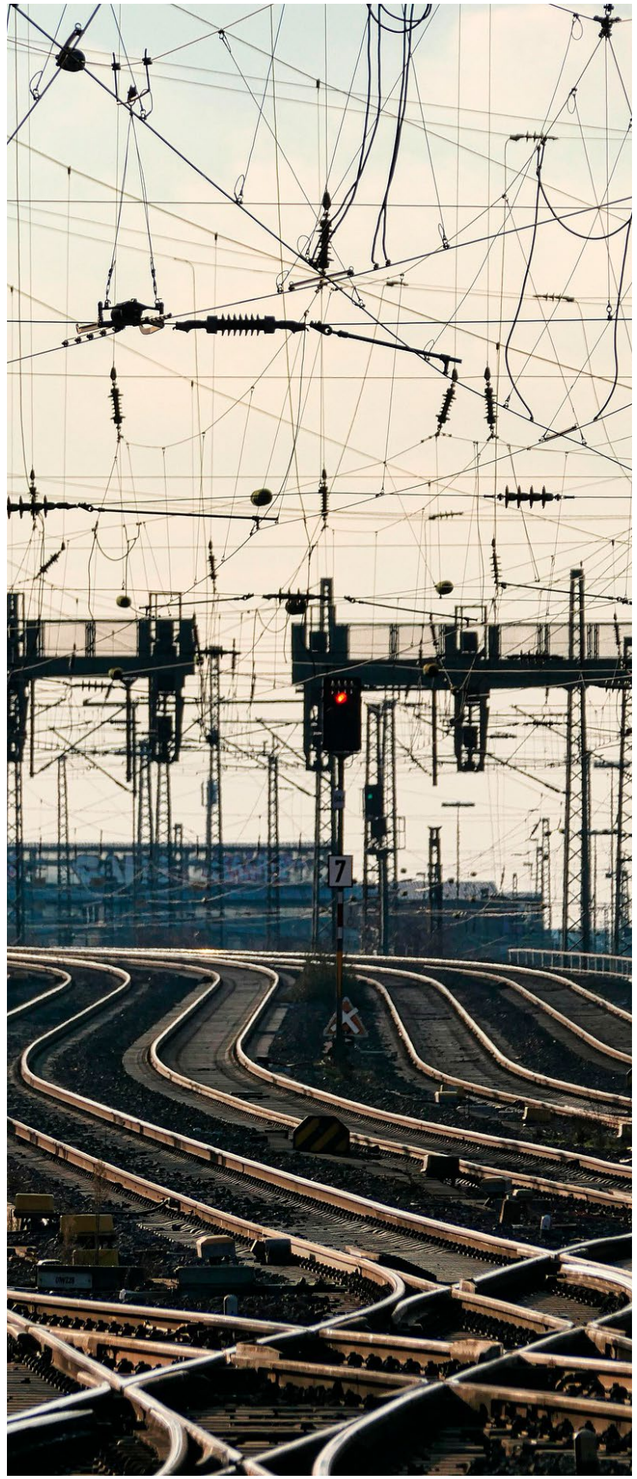

Response comes sixth in the progression, but the incident response plan does not come last. It should exist from day one. The sixth level turns that plan into an operating rhythm in which every event is closed and its lessons return to the inventory and the detections. The move: update the playbook, then run the next tabletop to test it. Playbooks must be built for xOT environments, because an IT playbook run against a live process can cause the outage the attacker wanted.

The progression maps to the SANS ICS Five Critical Controls. It does not replace them; it sequences them. Every control is built by at least two levels, and the incident response plan is in place from day one. The model adds scope, so that every system influencing the process is covered and sequenced, so that each investment rests on the one beneath it.

Reading maturity honestly

Organizations often judge maturity by their strongest capability. Read in order; maturity stops at the first level that is not genuinely established. If discovery falls short, the program is at level one, and yet its green detection dashboards rest on an inventory nobody trusts.

Three habits help teams answer the leadership question with evidence:

- Run the scope test. Sit down with operations, list every system whose failure would stop production, OT and IT alike, and compare that list with what is monitored today. The difference is the real scope.
- Find the first gap. Score each level honestly, and count any uncertain capability as absent until it is confirmed.
- Measure, then move. Track one proof point per level every week, and direct the next investment to the level where the program actually sits.

Do that, and the next time leadership asks whether the organization is more secure than it was yesterday, the answer is not a list of projects. It is a number that moved, and the move being made next.

The scope test asks whether a system influences a physical process. A growing number of systems now go further, shaping or making operational decisions themselves, and that raises a different set of questions about trust.

Key takeaways

- Answer “are we more secure today than yesterday?” with evidence of what improved and the next move that improves it tomorrow.
- Define scope by whether a system influences a physical process or drives a high-impact operational outcome, not by whether it sits on the traditional OT network.
- Treat a clean hunt as evidence only for the assets in scope, and test the inventory before trusting the results.
- Diagnose whether shortfalls stem from tool fit or from sequencing before deciding where to invest.
- Build capabilities in order, from discovery through prioritization, remediation, detection, hunting, and response, with response feeding back into discovery and each level building specific SANS ICS Critical Controls.
- Prioritize the small share of vulnerabilities that require immediate action, and document compensating measures where patching is not possible.
- Judge maturity by the first level that is not established, and report progress with weekly evidence and a next move rather than project lists.

When AI Enters the Control Loop

Autonomy, operational data trust, and the consequences of delegating decisions to machines.

Introduction: From advice to action

Artificial intelligence entered industrial operations as an advisor. It is now beginning to act, and its trajectory points toward the control loop. That progression raises three questions critical infrastructure operators cannot defer: how far to trust the decisions AI makes or informs, how far to trust the operational data behind them, and who answers for the outcome when something goes wrong.

The benefits are real. The risk lies less in the technology itself than in complacency, where people come to trust AI output in the same way organizations sometimes over-trust individual operators without defense in depth.

Where AI is already at work

Most current deployments support people rather than replace them. Water utilities apply predictive models to chemical dosing and energy use, two of the highest operating costs after personnel, where slow process response makes anticipation valuable. Predictive maintenance applies learning algorithms to vibration analysis and thermography to anticipate equipment failure. In wind generation, models track structural stress over time. Computer vision analyzes images of structures and equipment for corrosion, reducing the need to send inspectors into the field.

Security operations are another established use. Behavioral modeling of OT network traffic, treating machines as entities whose normal behavior can be learned, has been in use at some utilities for close to a decade. It flags devices communicating in ways they never have before and allows a small team to cover more ground. Security operations centers with SIEM and SOAR platforms increasingly integrate AI into triage and patch prioritization, typically starting at the DMZ, with the prospect of moving to lower levels.

Appetite varies by sector and function. Optimization on a manufacturing line carries different consequences from AI influencing a safety-instrumented system in a refinery or chemical plant, where the quality of the data used to train and tune a model would need to be beyond question. For now, most operators remain reluctant to let AI make process-level or control decisions. Whether a human sits in the loop, approving each action, or on the loop, supervising, is an open debate that depends heavily on context.

AI arrives through the back door

AI does not always arrive as a deliberate project. It can come bundled in a software update. In the nuclear sector, some of the first AI capabilities to appear inside protected

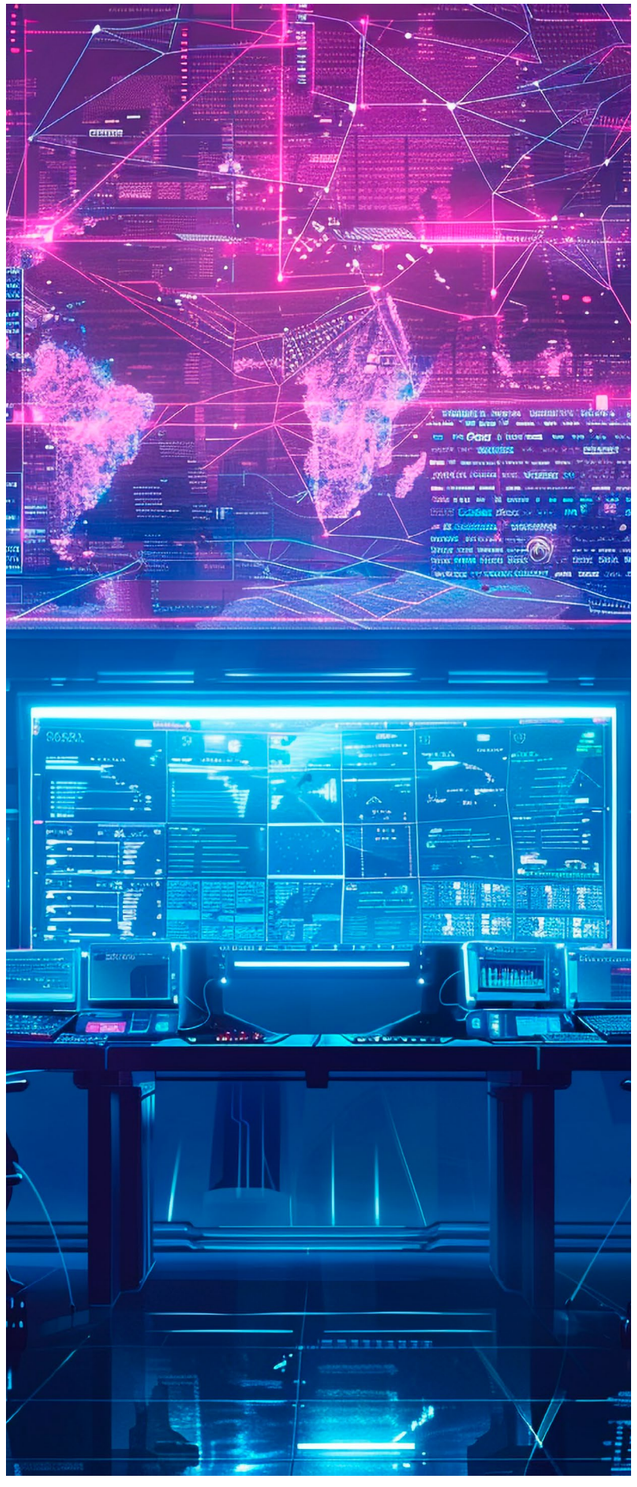

OT environments arrived through updates to existing security tools, and organizations did not always realize they were there, let alone approve them. Integrators and managed service providers can introduce AI agents to simplify their own work, leaving owners who have outsourced those functions unaware of what is now running in their environment.

The connectivity behind these features matters as much as the features themselves. An AI capability that depends on a cloud service may be communicating out of an environment its owner believes is isolated. A system connected to a cloud-based model and sitting in an OT DMZ also accumulates detailed knowledge of the environment, which an adversary could exploit to move from IT into OT where segmentation is weak.

Adversaries are already using AI. In a reported intrusion involving a Mexican water utility, attackers used commercially available AI assistants to support reconnaissance and analysis. The case illustrates how AI can accelerate established attack techniques without requiring fundamentally new methods. Consistent application of basic controls such as segmentation and keeping systems off the internet would likely have slowed or obstructed it.

Fundamentals first, then the new risks

Organizations with mature security practices start from a stronger position, regardless of size. Segmentation, access control, and management of change processes ensure that a new capability, including one delivered by an update, is tested and evaluated before it reaches production. The same processes should catch unexpected external connections and prevent anyone from establishing an API connection between a new application and an existing historian without review.



AI also introduces concerns that those processes were not designed to cover. How a model communicates, how inference is performed, how prompts and responses are logged, and how retrieval draws on other data sources all need to be traced. A trusted model can still be compromised through poisoned training data, insider action, or a malicious update. Protecting it means covering every pathway into OT (wired, wireless, physical, supply chain, and portable media) throughout the model's operational life.

Data governance becomes a frontline issue. Data uploaded to an external AI service may leave the organization's controlled environment and, depending on the provider, account settings, and contractual terms, may be retained or used for model training. A workforce that uses approved internal tools at work may still turn to public services from home. Clear rules on which models may be used for which data, supported by training, are part of the control set. Data integrity, long a lower priority in OT discussions, becomes central when that data informs automated decisions.

Smaller utilities do not need large AI budgets to benefit. Some are pairing open-source network monitoring with locally hosted models to produce useful dashboards at low cost. A dedicated on-premises AI subnet running local models avoids both cloud dependency and unpredictable usage costs. Local models may be less capable than frontier services and require effort to tune to a specific plant, but they keep data and decision support within the organization's boundary. Adoption in small utilities also tends to depend heavily on individual skill, which makes governance and peer support more important.

Accountability stays with people

When AI contributes to a bad decision, its involvement does not remove the responsibilities of the people and organizations operating the facility. In regulated sectors such as water and nuclear power, licensed operators retain important operational responsibilities for decisions made under their supervision. That has practical consequences. People need to stay in the loop, review AI-informed decisions, and confirm that the data informing the model is sound. The risk to guard against is rubber-stamping, where operators approve recommendations as a matter of routine, along with gradual erosion of the technical skills needed to challenge them.

Operators preparing for greater autonomy are approaching AI through the same human performance and defense-in-depth principles applied to people. No single human makes every critical decision unchecked; procedures, peer verification, and structured communication catch errors. Some failure modes remain unresolved. Hallucination, distinct from drift or data manipulation, resembles an operator arriving unfit for duty, a condition procedures and checkpoints are not designed to detect in a machine. Until such gaps are closed, full autonomy in safety-sensitive operations remains out of reach.

A practical working assumption is to treat any AI system as another entity in the environment, one that could act as an insider threat, and to govern its access, monitoring, and authority accordingly.

The paths by which AI reaches operations, through updates, integrators, and outside connections, are the same paths by which people and vendors reach them. Controlling those external pathways is the next part of the exposure question.

Key takeaways

- Distinguish AI that supports an operator's judgment from AI that exercises operational authority, and set governance for each.
- Assume AI can arrive through software updates and service providers, and use change management to detect and evaluate it before it reaches production.
- Verify where AI-enabled tools send data and whether they depend on cloud connectivity, particularly in environments believed to be isolated.
- Apply security fundamentals first, then extend controls to model integrity, training data, prompts, logging, and retrieval data flows.
- Keep licensed operators accountable and in the loop, with safeguards against rubber-stamping and skill erosion.
- Treat every AI system as an entity that could become an insider threat, and limit its access and authority accordingly.

Remote Access: Same Risk, Different Rulebooks

Defending the pathways into electric and water utility operations.

Authored by Dispel

Introduction: The 2 a.m. call

It is 2 a.m. A pump station or a substation is misbehaving, and the vendor or OEM that supports it needs access now. What happens next reveals more about an organization's control of its operations than any audit. The session that lets a specialist diagnose the fault may also reach the setpoints that govern chemical dosing in drinking water or the protection settings in a substation. Discomfort begins when an external connection crosses from the business network into the operational zones that control those outcomes.

Electric and water utilities face that moment under very different rules. In North America, applicable entities and assets within the bulk electric system have been subject to mandatory NERC CIP requirements for years, while the water sector is still working toward a common baseline. Both share the same weak point: the paths into control systems. Vendor sessions, integrator tunnels, remote access to HMIs, and shared credentials are how much of the work gets done, and they are also where much of the exposure sits. The question that matters is less who is more compliant than who actually controls who can reach OT.

Two ends of the spectrum

Remote access failures run in both directions. At one end are environments still running operating systems far beyond vendor support, default passwords, and certificate infrastructure so dated that modern systems no longer recognize it. At the other are large organizations that have outsourced access approval to a third party, where emergency onboarding can take weeks and removing access can take as long again, leaving firewall rules open in the meantime. Neither serves the 2 a.m. call.

Imperfect controls are still better than none. A workable arrangement may involve a VPN, a vendor calling in for a one-time code, and several jump hosts before reaching a controller. It is not elegant, but it is far safer than equipment left exposed on the internet. Better options now exist, and the standard to aim for has risen accordingly.

Why remote access remains a primary vector

Threat intelligence continues to show remote access as one of the main ways adversaries reach industrial environments. The exposure is broader than user-to-system connections. It includes system-to-system links and the appliances used to provide access, such as firewalls and VPN gateways, where newly disclosed vulnerabilities have been exploited within a day of publication. Adversaries also compromise third parties, using an approved OEM or service provider connection to reach the real target.

Secure remote access is one of the SANS ICS Five Critical Controls for that reason. Deploying a tool does not satisfy it. What matters is how the tool is operationalized, governed, and staffed. A tool is not a control framework.

"We have a VPN" often ends the remote access conversation when it should begin it. A VPN typically provides encryption in transit and some form of authentication, which addresses only a handful of controls. The rest depend on answers the VPN does not supply. Who validated that this person should have access? Who approved it? Is the session recorded, and is anyone reviewing the recording? Who responds if something looks wrong? Each answer needs to be set in policy, supported by procedures, and carried out by people with the time to do it.

The starting point is a framework, not a product. Options include the SANS ICS Five Critical Controls, guidance from the American Water Works Association, and NIST SP 1800-45,

DISPEL

Cybersecurity for the Water and Wastewater Sector practice guide that demonstrates how to secure remote access to operational technology. Organizations outside mandatory regimes, such as electric distribution utilities beyond NERC CIP's scope, have taken a similar route: adopting an established standard, adapting it to their own operations and third-party contracts, and managing risk against it. Debating the relative merits of two good frameworks is a sign of progress.

What each sector can learn from the other

NERC CIP concentrates on the electronic security perimeter and on who may enter it and how. Its compliance overhead shapes behavior. Granting third parties access inside the perimeter requires significant effort, so organizations tend to do it only when there is a compelling reason. That has benefits, and it also drives workarounds.

Practical improvements need not wait for full IT and OT separation. One effective early step is separating identities, so that a single compromised corporate credential no longer provides a direct path into operational systems. It is logical rather than a complete separation, but it removes one of the most direct routes an attacker can take.

The water sector has little appetite for adopting NERC CIP wholesale. Sector bodies are instead pursuing a similar partnership model between industry and government, aligned with NIST and avoiding unnecessary burden. Practitioners warn that compliance can breed complacency, and that the controls must work for the smallest utilities, where one person may run operations and IT. Guidance written in plain language, with few acronyms, matters as much as the controls themselves.

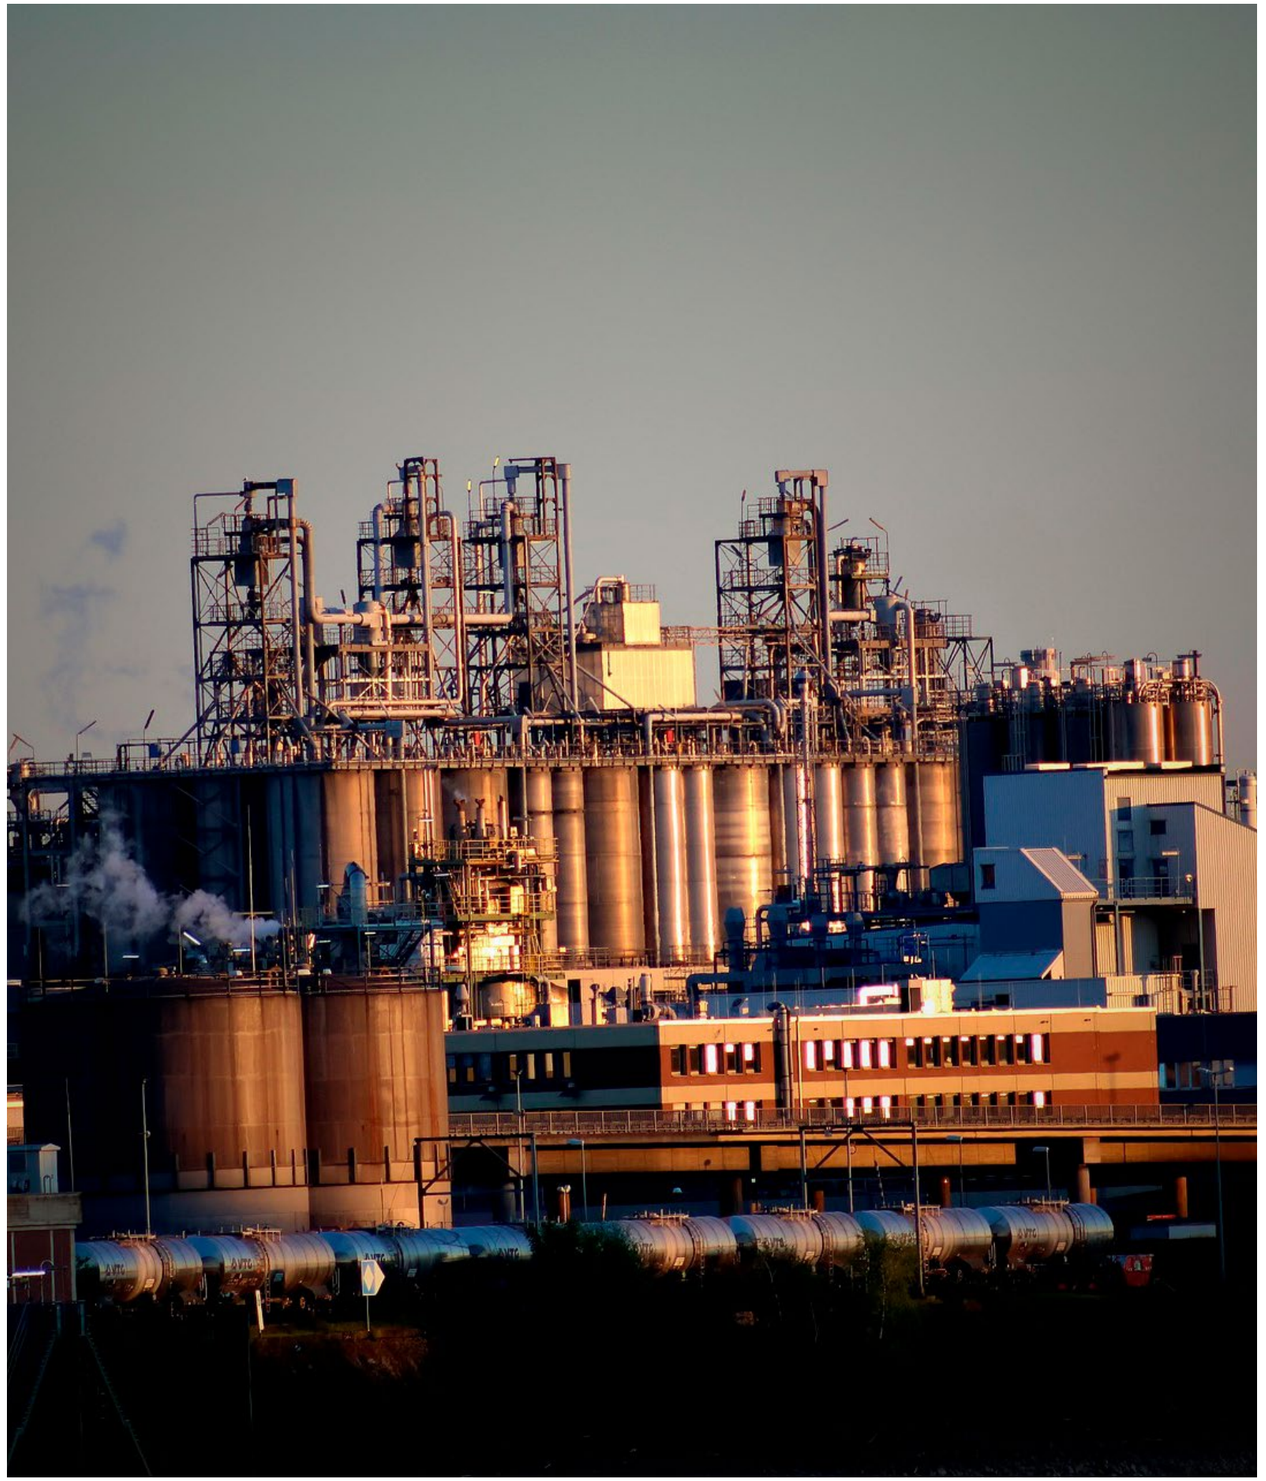

Approaches elsewhere vary widely. Some countries apply strict requirements to facilities above defined criticality thresholds, in some cases mandating data diodes and prohibiting remote access, including at desalination plants. Others have little regulation. Water is often left out because individual systems fall below thresholds set on economic grounds, despite being essential.

Making the 2 a.m. decision workable

The person approving an access request in the middle of the night is usually not a security specialist and will not wait for the SOC to analyze the request. Effective platforms give that person enough context to decide. One approach is a simple traffic-light rating: green requests can be approved, amber requests are left to discretion or local policy, and red requests are denied automatically or require additional steps to reduce the risk. Combining risk signals from different tools on the approver's behalf, while still passing raw logs to the SOC, moves the heavy lifting away from the operator.

Just-in-time access marks a substantial change from the always-on VPN. Instead of handing out copies of a front door key and never changing the locks, the asset owner approves each session before a vendor can enter and can isolate a session that behaves unexpectedly.

Knowing who has access today comes first. Point-in-time packet captures miss any session that is not active when they run, so continuous network monitoring is foundational. Physical walk-downs are equally important. Cellular modems in control cabinets and network cables run beside HMIs to bypass hardened boundaries are common discoveries, because people will do what they must to get the job done.



Can you disconnect in minutes?

A final test is whether an organization could cut all external access to OT quickly without stopping operations. Well-prepared utilities can keep treating water and running pump stations on their local control systems, with staff traveling to sites to carry out work that would otherwise be done remotely. Others prepare staged breakpoints in advance, with rules ready to disconnect remote access protocols between corporate and OT networks or to separate them entirely, depending on the threat.

The remote access platform itself must fail safely. It should sit at the boundary, not deep within critical networks, and if it fails, it must do no harm to operations.

Every one of these pathways ultimately terminates in a product. When the vulnerability lies in the access appliance or the equipment behind it, the question becomes what manufacturers owe the operators who deploy their products.

Key takeaways

- Judge remote access by who controls entry to OT, not by which compliance regime applies.
- Treat third-party, system-to-system, and access appliance pathways as primary attack vectors that require monitoring, not one-time configuration.
- Start from a control framework and map tools to it; a VPN alone addresses only a small part of the requirement.
- Separate operational identities from corporate credentials as an early, achievable step.
- Give approvers clear context and just-in-time control, so that urgent access decisions can be made quickly and safely.
- Establish who has access today through continuous monitoring and physical walk-downs, and prepare to disconnect external access without stopping operations.

When Regulation Meets Operational Reality

What the Cyber Resilience Act owes operators and what NIS2 asks in return.

Authored by Vayron

Introduction: A new exchange between vendor and operator

Every remote access path and control system runs on someone else's product. When that product carries a vulnerability, the operator depends on its manufacturer for warning, information, and a fix, which until recently rested largely on the vendor's goodwill. The EU Cyber Resilience Act (CRA) turns much of it into obligation.

The CRA matters beyond the EU. Vendors selling into the European market must meet its requirements, and few are likely to maintain separate product lines for other regions. That is a reading of the economics rather than a legal requirement, but if it holds, plants elsewhere will receive the same products and advisories. An organization that only buys and operates products takes on no CRA duties itself. Its obligations come from national law, including the transposition of NIS2 within the EU.

What the CRA gives operators

The CRA establishes five entitlements that matter to plant operators:

- Warning of exploited vulnerabilities. Since September 11, 2026, manufacturers must report actively exploited vulnerabilities and inform impacted users, with mitigation information and, where appropriate, in a machine-readable format.
- A declared support period. It must reflect how long the product is expected to be in use, with a minimum of five years, and its end date must be stated at the time of purchase. The CRA recognizes that products used in industrial settings, such as industrial control systems, often remain in service far longer.
- Security updates. These must be provided without delay and, where technically feasible, separately from feature updates, so that operators no longer have to accept a feature release to obtain a fix. They are free of charge unless the manufacturer and a business user agree otherwise for a tailor-made product.
- A software bill of materials. Manufacturers must maintain an SBOM in their technical documentation. They are not required to hand it to customers, so access is a contractual matter.
- Coordinated vulnerability disclosure. Manufacturers need a disclosure policy and a channel for receiving vulnerability reports.

The timing matters. The CRA's actively exploited vulnerability reporting obligations also apply to products placed on the EU market



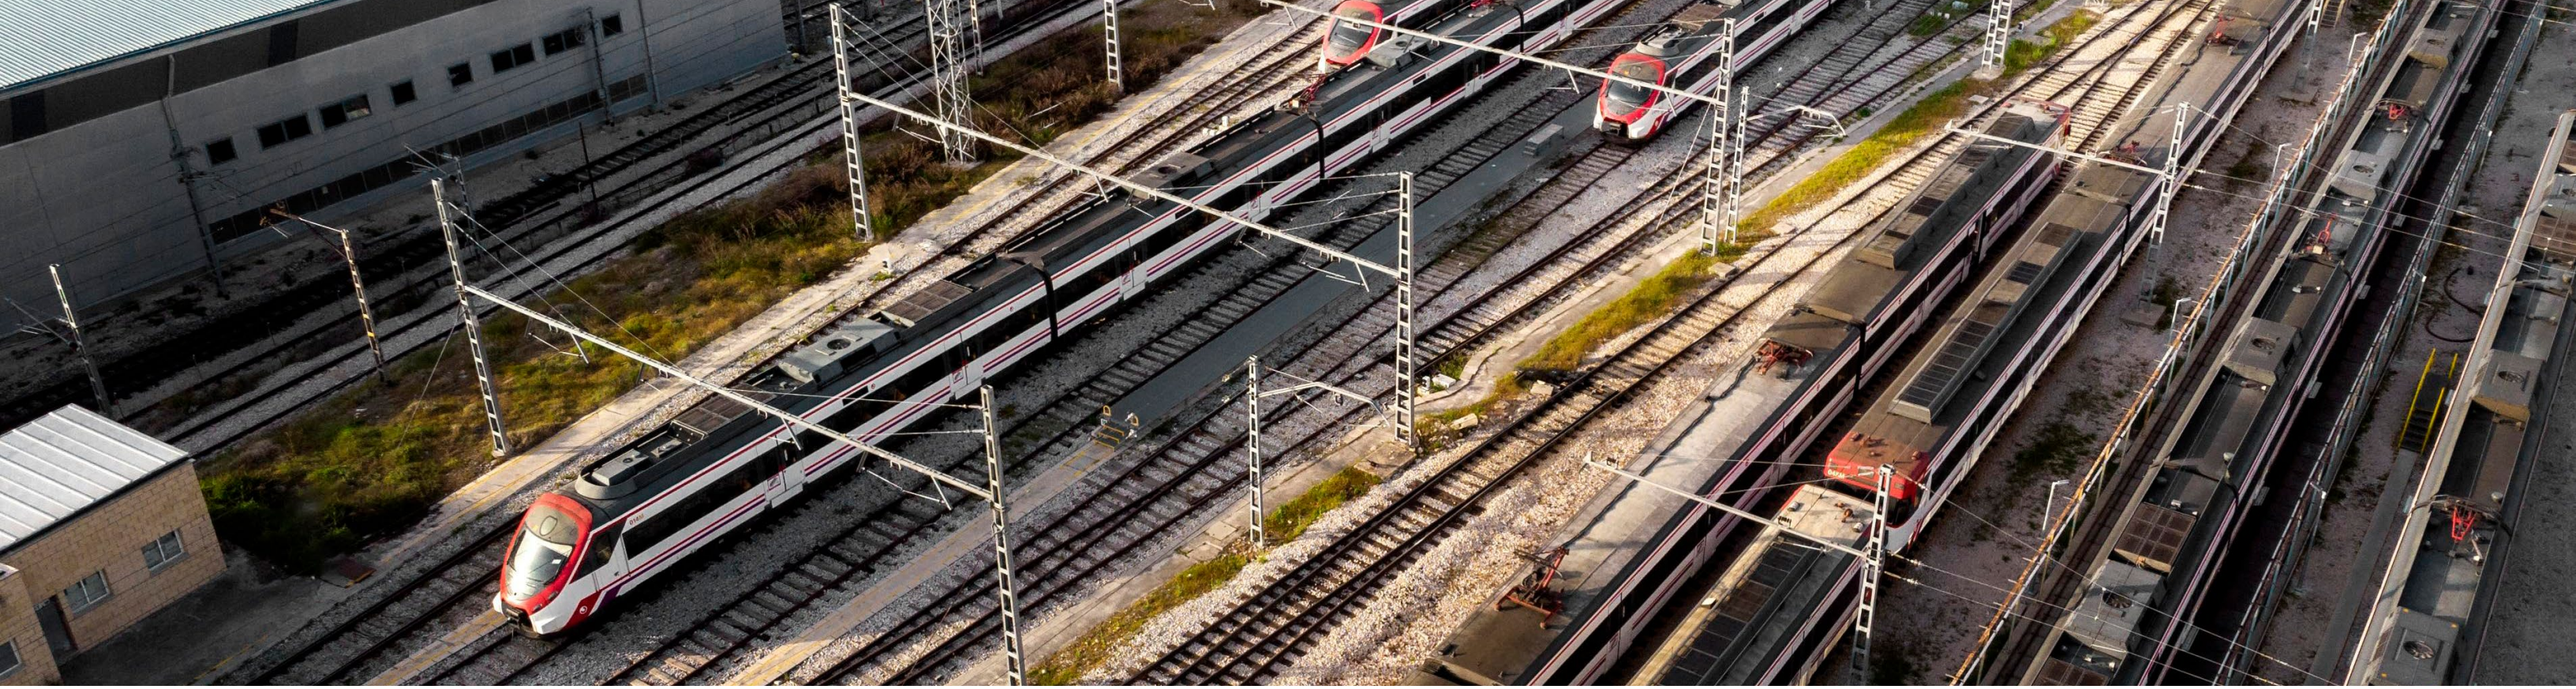

before December 11, 2027, including products already installed, where the applicable reporting conditions are met. Most other requirements attach only to units placed on the EU market from December 11, 2027, even if they are identical to units already in service. Identical spare parts for repairs are out of scope. An installed product falls under those requirements only if it is substantially modified after that date. Security fixes are generally not treated as substantial modifications; feature changes can be. For years, much of the installed base will live on the warning alone.

What NIS2 asks in return

NIS2 asks covered entities for evidence. They must address security in their supply chain, taking into account the quality of their suppliers' products and security practices, handle vulnerabilities in what they acquire and maintain, and report significant incidents through early warning, notification, and final report stages within 24 hours, 72 hours, and one month. The CRA makes vendors produce much of that evidence: the warning with mitigations now, and from December 2027 CE marking that covers cybersecurity, a declared support period, a disclosure policy, free security updates, and a contact that must take the operator's report. For the first time, the vendor's compliance team is on the operator's side. Whether that evidence is enough remains the operator's call.

One vulnerability, two regimes

Consider a substation operator that receives an advisory for its protection relays. A denial-of-service flaw sits in an open-source web server component, reachable over the network without authentication. The vendor advisory arrived four months after the underlying vulnerability was published. It is rated high severity, with no public exploit and no known exploitation. The fix exists only inside a major firmware upgrade, which for a protection relay means protection testing and an outage that may be months away. With no intrusion prevention signature available, the operator restricts the relay's web service to the engineering workstation and accepts the residual risk until the upgrade. On the facts available, it is a reasonable decision.

Now rerun the case under the CRA, with a unit of the same relay model placed on the EU market in 2028, so every obligation applies. Assume the flaw is already being exploited at another utility, and the vendor knows. With an SBOM delivered under contract and matched against new vulnerabilities, the operator learns about the web server flaw on day one. The vendor must warn it, with mitigations, and the security fix ships on its own, free of charge, without the major upgrade.

The decision changes. Leaving an exploited vulnerability open for months becomes hard to justify, while an unplanned outage may be unnecessary if tighter controls can hold for a few weeks. The least-risk action this cycle is to tighten the access rule, hunt for signs of compromise, test the fix on a reference relay, and install it within weeks. A security-only fix is still a firmware change on a protection device, and it still needs testing and an outage. Questions about reachability and what an outage means for the plant remain the operator's to answer. The CRA brings facts and options; the decision stays with the plant.

The CRA owes installed relays the warning, not the fix. If the model is still sold after December 11, 2027, however, the vendor must provide fixes for new buyers, and those fixes will often install on existing units running the same firmware line. That is not guaranteed. Planning a major upgrade onto a firmware line the vendor will continue to sell improves the odds. If the hunt finds exploitation in the operator's own network, two clocks start.

The CRA's definition of an actively exploited vulnerability covers exploitation in a system without its owner's permission, so the operator's

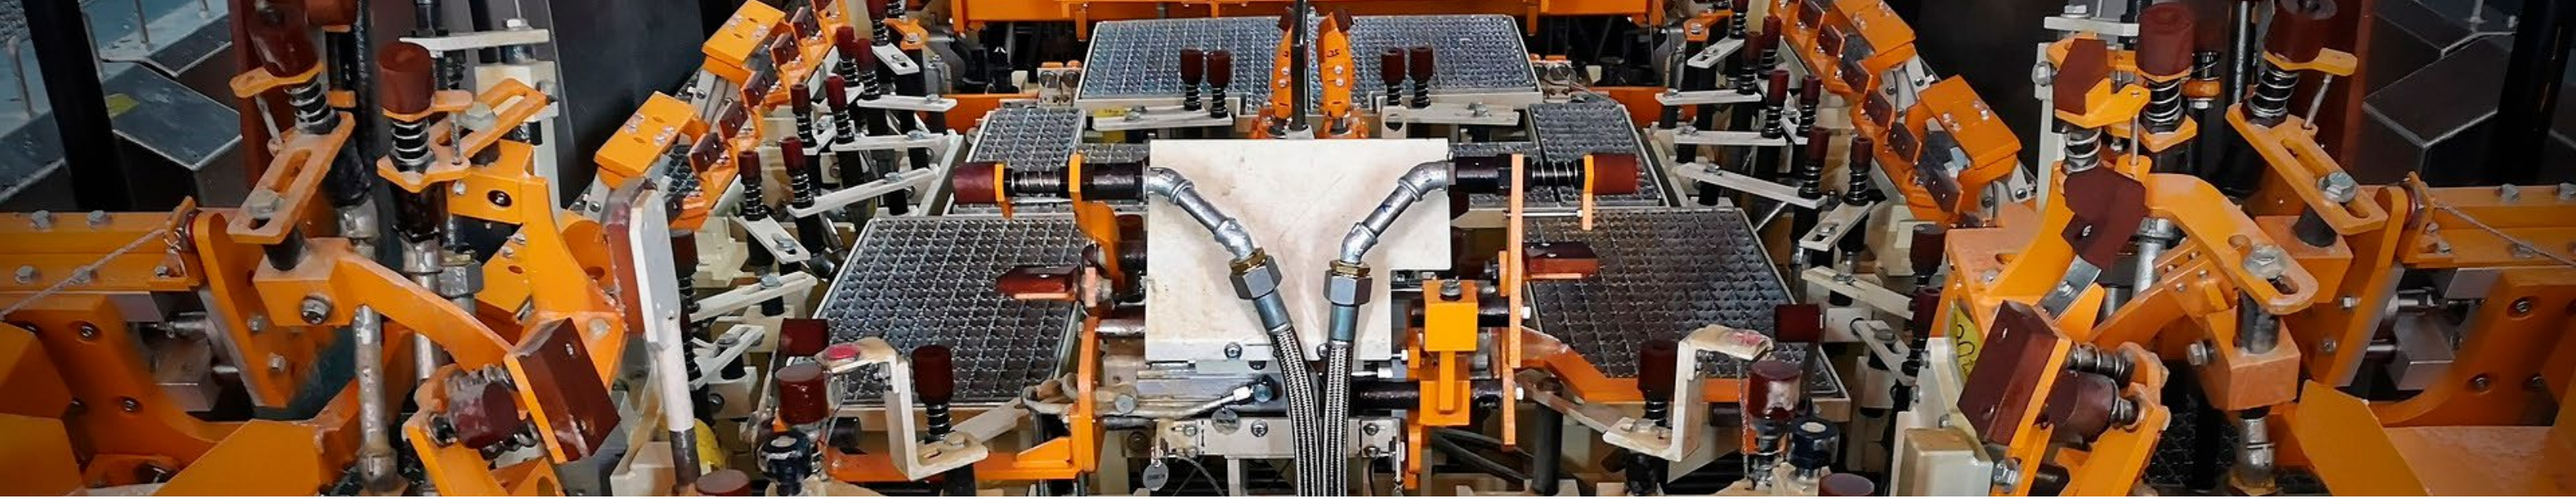

detection may be the evidence the manufacturer needs. Once the manufacturer becomes aware of an actively exploited vulnerability, it must submit an early warning within 24 hours and a vulnerability notification within 72 hours. A final report is due no later than 14 days after a corrective or mitigating measure becomes available.

The manufacturer must also inform affected users, and if it does not do so in time, the coordinating CSIRT may inform them itself. The operator, if the event is a significant incident under NIS2, runs its own 24-hour, 72-hour, and one-month sequence. One event produces two reports, often prepared by the same people.

Where a vulnerability cannot be patched before the next maintenance window, NIS2's risk-based approach expects a deliberate decision: meaningful compensating measures and a documented rationale for accepting the remaining risk.

Operators outside the EU

Outside the EU, the CRA covers only what vendors place on the EU market. Its warnings and support commitments reach other plants through the vendor's single product line, if at all, so they are best written into contracts.

Secure by Demand guidance for OT owners and operators, published in January 2025 by CISA and international partners including the European Commission, lists what to ask vendors for, including defined support periods, free patches, SBOMs, and a vulnerability disclosure program. The guidance tells operators what to ask for; the CRA requires vendors serving the EU market to deliver much of it. Where the two differ, as with delivery of the SBOM, the contract decides.

New risks, and five lines for the next contract

The CRA also creates new exposure. Plant extensions may receive next-generation equipment that does not fit safety cases built around older units, although units first placed on the EU market before December 11, 2027 keep their status regardless of installation date. An integrator that substantially modifies and supplies a product may become its manufacturer, for the modified part or, if the change affects the product's security as a whole, for all of it.

Some vendors will likely retire legacy lines rather than bring them into compliance, affecting plants everywhere. Every critical vendor should be asked where it stands on the CRA and which installed products it will carry beyond December 2027; a vendor unaware of the CRA is a warning sign.

Five provisions belong in the next procurement contract:

- SBOM delivery: format, depth, and cadence, with a new SBOM for every update.
- Support period: an end date aligned with the planned operating life, plus a notice period.
- Exploited vulnerabilities: a service level and a named channel in both directions.
- Security fixes: available for units already in service, without feature upgrades, and validated for the operator's configuration.
- Discontinuation: advance notice, a last-time-buy window, and a compatible successor.

For manufacturers, these five lines are what their customers will ask for next.

In the long run, the CRA should raise the floor for what plants buy: products built secure by design, shipped without known exploitable vulnerabilities, and supported for a declared period. That benefit arrives as the installed base turns over, which in OT takes years. The CRA puts product security on the manufacturer and the key facts on the operator's table. The decisions stay with the plant, and they do not get easier.

Key takeaways

- Expect CRA obligations to reach plants outside the EU through vendors' product lines, but secure them contractually rather than assuming them.
- Use the CRA's exploited-vulnerability warnings, support periods, separate security updates, and disclosure channels as evidence for NIS2 supply chain obligations.
- Recognize that installed products are owed the warning but not the fix, and plan firmware upgrades onto lines the vendor will continue to support.
- Treat CRA information as an input to a risk decision that remains the operator's, including testing and outage planning for security-only fixes.
- Prepare for one event to trigger both manufacturer and operator reporting clocks, often handled by the same team.
- Add SBOM, support period, exploited-vulnerability, security-fix, and discontinuation terms to procurement contracts now.

Backed Up, But Can You Recover?

Why successful backups do not guarantee the restoration of critical operations.

Authored by Macrium



Introduction: Protected is not the same as recoverable

For years, critical infrastructure security concentrated on keeping threats out and making sure backup procedures existed. Both remain essential. A series of high-profile incidents has made clear that prevention alone is not enough. When an attack gets through, or a critical system simply fails, what matters is whether operations can be restored within the time the process can tolerate.

The distinction is between protecting data and restoring an industrial process. A completed backup job proves that data was copied. It does not prove that a control system can be rebuilt, that the people who know how to do it will be available, or that the restored system will run the process again.

Measuring what is proven, not what is in place

Organizations often measure preparedness by what they have installed. Consider a site with a modern disaster recovery solution, fully documented procedures, and a calculated recovery time of one hour. By most frameworks, it appears well prepared. Then suppose the one or two people who know how to carry out the recovery are hundreds of miles away from a remote or offshore facility, where there is often no dedicated IT support on-site. The one-hour estimate no longer holds.

Recovery time calculations go wrong when they ignore practical constraints like these. Closing the gap requires hands-on exercises, regular recovery testing, complete backup coverage, and cross-training so that procedures do not depend on particular individuals. Recovery plans should be built around the real limitations of the specific site and sector, not an idealized version of them.

Macrium's 2026 manufacturing benchmark study found that only 18% of manufacturers met or exceeded their recovery-time targets in full recovery tests. A recovery target is only useful if it has been proven achievable in the conditions the site will actually face.

Backup and recovery are also frequently retrofitted into environments that were never designed for them, leaving systems that cannot be restored. Recovery by design

reverses that. Backup and restoration are considered from the first stages of a project through to production, so that recovery capability is built in rather than added later.

Coverage beyond the obvious systems

A common misconception in industrial environments is that backing up critical servers and engineering workstations protects the critical systems. However, a control system depends on much more, including PLC programs, third-party application servers, and network switches carrying specific configurations, apart from those on industrial networks. Each can be as critical to restoring operations as the main servers.

The consequences can be severe. In one case, the failure of a single network switch brought down operations on an offshore facility. The switch's configuration had never been backed up, and nobody knew where any copy might be stored. The main servers were fully protected; the process still stopped.

Not every asset needs the same backup frequency. A configuration that changes daily needs more frequent backups than a largely static system. What matters is that every subsystem has a backup, and that each backup has been shown to be recoverable.

Coverage also depends on change control. Equipment is replaced and systems are modified, but backup procedures are not always updated to match. An asset list that is not maintained as the environment changes produces recovery procedures written for equipment that is no longer there.



Documentation that works under pressure

Recovery must be repeatable by someone other than the person who designed it. Organizations sometimes inherit backup arrangements with no documentation at all and must reconstruct how they work. Documentation should be treated as a living record, updated as equipment and people change. It should capture not only what is protected but why it is configured that way, such as the rationale behind different backup frequencies.

It should also cover failure scenarios that differ from the rehearsed ones. What happens if the network is down during recovery, or if the backup server itself is unavailable? These situations occur, and documentation that addresses them allows someone other than the original owner to achieve the same result.

Good documentation carries its own risk. The information that helps recover a system can help an attacker compromise it: network diagrams, backup server names, hostnames and IP addresses, system configurations, domain details, software versions, and privileged account information. Organizations handle this in different ways. Some keep complete documentation under strict access control. Others separate the procedures from the sensitive details, holding the latter in a secure register with its own controls, so that a single leaked document does not reveal everything. The choice depends on a risk assessment and the organization's risk appetite.

There is a balance to strike in volume as well. Thousands of pages are of little use in the middle of an incident, and every additional document adds maintenance overhead. The aim is the right amount of documentation for the environment.

Validating recovery without overwhelming the team

Confidence in a backup comes from validation, not from faith in the software. Validation methods sit on a spectrum of effort and assurance. At the lower end, hash verification confirms that backup data has not been corrupted, and restoring to a virtual machine shows that a system and its critical applications will start. Neither proves that a system can be restored to dissimilar hardware when an identical replacement is unavailable, or that the recovery can be performed without the person who set it up. Full recovery exercises and tabletop scenarios provide greater confidence but require more effort.

Few organizations have the time or budget to test every backup every time. A practical approach links test frequency to criticality and backup frequency. A critical application server backed up daily might be proof-tested weekly, while operator stations backed up monthly might be tested quarterly. Systems backed up more often are tested more often. Every asset, and every backup approach, should be tested at least once. A validation program that demands unrealistic effort is likely to lose management support.

The starting point is modest. Identify the systems in the environment and how critical they are, protect them, prove that recovery works, and maintain that process over time. Validating one backup today is better than having validated none.



Legacy systems and lifecycle planning
Many industrial systems have run reliably for 15 or 20 years or longer, on operating systems, hardware, and software that are no longer manufactured or supported. Upgrading them can involve significant engineering effort, revalidation, sometimes complete re-engineering, and downtime that is difficult to justify financially. Lifecycle management in OT rarely receives the priority it has in IT.

That is not a reason to leave obsolete systems in place indefinitely, since obsolescence expands the attack surface. Organizations need to plan for eventual replacement, and reliable backup and recovery provides a safeguard while that planning proceeds.

The question to ask is not whether a disaster recovery plan exists, but whether it is practical and proven to work. Even a proven plan has limits. When systems cannot be restored quickly, or outside help cannot reach the site, the challenge shifts from recovering technology to keeping essential services running without it.

Key takeaways

- Measure recovery readiness by what has been proven through testing, not by the tools and procedures in place.
- Build recovery time estimates around real constraints, including staff availability at remote sites and the need for cross-training.
- Extend backup coverage to PLC programs, third-party servers, and network device configurations, not only servers and workstations.
- Maintain asset lists and recovery procedures through change control, so they reflect the equipment actually installed.
- Keep documentation current and usable under pressure, while protecting the sensitive details it contains.
- Match recovery testing frequency to criticality, start with the most important systems, and test every asset at least once.

Keeping Essential Services
Running After Compromise

Preparing for the incident in which prevention fails and outside assistance is unavailable.

Introduction: Starting from the other end

Most security planning is about keeping the attacker out. Recovery planning assumes that systems can be restored once the attacker is gone. Neither covers the situation critical infrastructure increasingly has to plan for: an adversary inside the network, communications disrupted, and no one available to call for help. The question is no longer only whether systems can be recovered, but whether the services they support can continue in the meantime.

That shift underpins CI Fortify, guidance issued by CISA with international partners that focuses on isolation and recovery. Its aim is not to make operators autonomous indefinitely. It is to prevent essential services from dropping from full capacity to nothing, so that a minimum viable level of service supporting health, safety, economy, and national defense can continue until outside assistance arrives.

Why the planning assumptions are changing

Two developments drive this approach. The first is destructive activity against OT equipment itself. In late December 2025, attackers entered Polish energy sector networks through internet-exposed VPN gateways, uploaded corrupted firmware to RTUs, deleted essential files from protection relays and other controllers, and reset communications devices to factory settings, according to CERT Polska. An attempt to run wiper malware at a combined heat and power plant was blocked by endpoint protection on the machines it reached. Some legacy field devices lack firmware signature verification, making them vulnerable to malicious firmware installation. If equipment were damaged at scale, spare hardware from manufacturers could run short, and recovery could stretch from days to months or longer.

The second is telecommunications. State-sponsored actors have compromised backbone telecommunications equipment, and access of that kind enables disruption as well as espionage. Exercises have repeatedly shown that critical infrastructure does not operate well without phones and internet. Unless an organization owns its entire communications path, an outage at a third-party provider will affect it.

These threats also break a long-standing emergency planning assumption. Mutual aid works because a disaster in one region leaves others able to help. A widespread cyber event offers no such guarantee. The limited

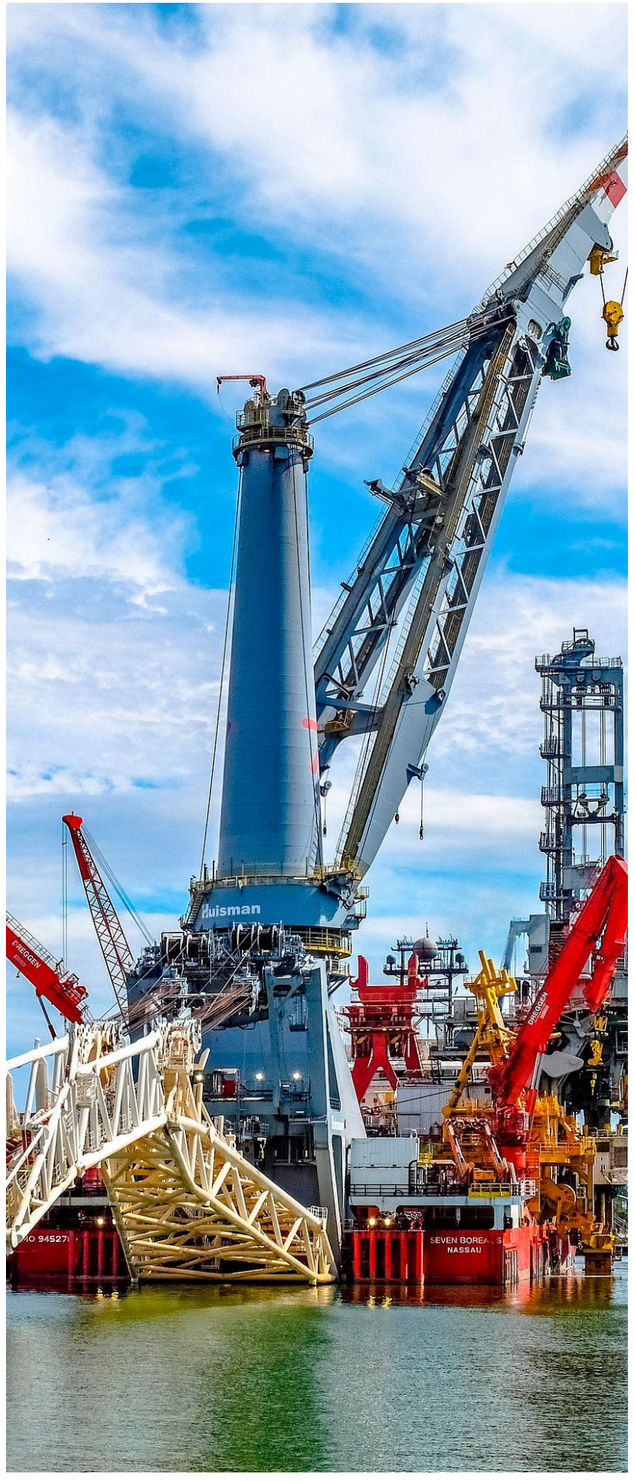

pool of OT incident responders, integrators, engineering firms, and managed service providers would face intense competition for their time. Help may come, but not necessarily quickly.

Isolation is a plan, not an air gap

Isolation in this sense does not mean permanently disconnecting OT. Remote access, third-party communications, integrators, and managed security services are normal business arrangements. The point is to keep within the organization the ability to operate when those arrangements fail or become untrustworthy.

Isolation also protects equipment that has not yet been damaged. When a device fails, engineers will typically replace it from spares long before anyone from cybersecurity arrives. If the adversary is still present and the replacement is configured the same way, it can be compromised immediately through the same route. Responders have seen replaced equipment re-compromised in exactly this way. If an attacker is deliberately destroying equipment, spares run out faster than expected. Separating IT from OT, and preparing engineers to recognize the signs and preserve remaining equipment, limits that damage without trying to change how engineers routinely work.

Decisions about isolation carry different costs. Disconnecting IT from OT may cost little for some organizations, while disconnecting cellular links to remote sites may be expensive. Knowing in advance which disconnections are easy, which are difficult, and who has the authority to order each allows an organization to respond to a warning from the government or an information-sharing body, or to news that a communications provider it relies on has been compromised.

Recovery when no one is coming

Recovery has a human dimension as well as a technical one. In a communications outage, staff need to reach each other, and the organization needs to reach its suppliers of treatment chemicals, backup fuel, and integration and security services. That requires planning: alternative communication methods with contacts already exchanged, and agreed default behavior when contact is lost. A chemical supplier that keeps delivering on the usual schedule if it cannot reach its customer may avoid a far worse outcome than an unplanned invoice.

Technical recovery depends on assumptions that often go untested. Many organizations increasingly rely on cloud services, vendor connections, licensing servers, and shared drives on the IT network. Backups that require any of these may not work during a communications outage. Some sites lack PLC logic or configurations on site, engineering software on site laptops, or staff trained to reconstitute SCADA without external help. In one recent incident, a compromised, internet-exposed PLC could only be modified by a single-person system integrator who was unavailable for two weeks. The operator waited, with the adversary's access still in place.

Endurance has limits too. One water utility with two weeks of backup power found that its real constraint was three days. Without SCADA, water quality sampling had to be done manually, and only three trained staff could do it before fatigue set in. Its solution was to train call center staff, who would have little to do during a communications outage, to take samples at remote sites.

Questions every operator should be able to answer

The CI Fortify approach centers on a series of practical questions:

- Vital customers: Who depends on the service, and where would its loss cause cascading emergencies?
- Vital sites: Which sites and control processes serve those customers?
- Degraded operation: Can those processes continue if communications and internet connectivity are lost or cannot be trusted?
- Isolation: Can that infrastructure be practically isolated when a warning arrives, and who has the authority to decide?
- Independent recovery: Can the systems be restored without outside support?
- Time to failure: How long can operations be sustained, and what is the real bottleneck?

From tabletop to functional test

Tabletop exercises are not enough. A functional test brings together cybersecurity, engineering, emergency management, IT, OT, and leadership to compare their understanding of how the system actually works, which often reveals that each group describes a different system. It then tests what happens. It does not need to cover an entire site to be useful.

Three areas deserve testing. For isolation, disconnect IT from OT and both from the internet, run SCADA over an alternative path such as radio or microwave, and disconnect third-party and cellular links. For recovery, establish which backups work without internet access and whether recovery depends on a single license dongle or a single person. For communications, apply primary, alternate, contingency, and emergency planning to remote sites and confirm that emergency communications work as expected. Something usually breaks during a test, and it is better to find it now.

Government guidance can help justify these exercises to leadership and secure budget, and the benefits extend beyond cybersecurity to emergency planning as a whole. Simply failing outright is not an acceptable outcome for infrastructure on which others depend.

Even well-prepared organizations cannot eliminate every risk. Some exposures will remain because the cost of removing them is too high or the means do not yet exist. How leadership decides which of those risks to retain, and whether that decision would stand up to scrutiny, is the next question.

Key takeaways

- Plan for an incident in which the adversary remains present, communications are disrupted, and outside help is delayed.
- Treat isolation as a prepared capability with defined decision rights, not as a permanent air gap.
- Protect spare equipment by isolating affected environments before replacing failed devices, so that replacements are not compromised the same way.
- Agree communication methods and default behaviors with critical suppliers before an outage occurs.
- Confirm that backups, engineering software, and configurations can be used on site without internet access or a single external expert.
- Run functional tests of isolation, recovery, and communications to find the real time to failure and its bottlenecks.

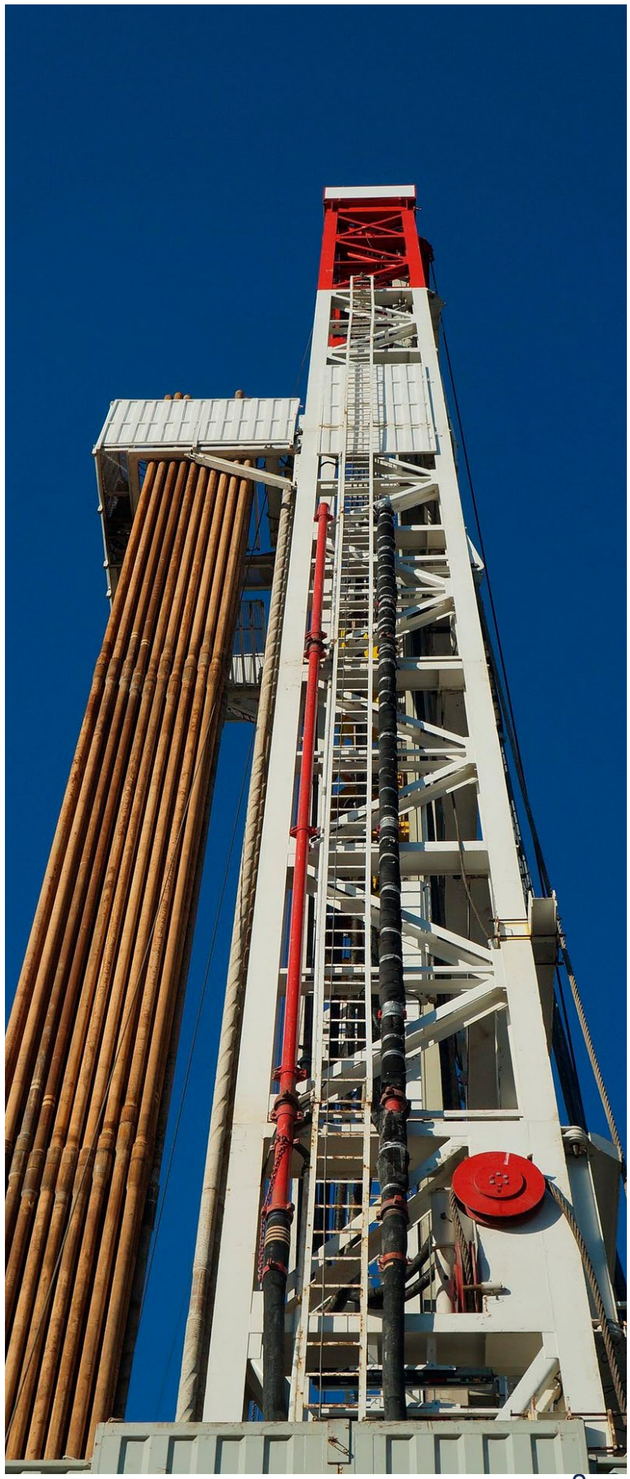

Making Residual Risk Defensible

What boards, regulators, and insurers should expect from industrial cybersecurity decisions.

Introduction: The risk that remains

No critical infrastructure organization can eliminate every cyber-physical risk. Budgets, outage windows, legacy equipment, and supplier dependencies all leave some exposure in place. The question that follows is whether the decision to retain it would stand up to scrutiny from a board, an insurer, a regulator, or a court.

Agentic AI provides a particularly demanding test of that question, exposing weaknesses in how organizations identify, document, and govern risks they cannot fully eliminate. Systems that act on their own introduce risks that do not depend on an adversary at all. Because these systems are probabilistic, drift, hallucination, and unexpected behavior are not remote possibilities but expected characteristics, and they arrive alongside adversaries who can use the same technology to attack. Few utilities are yet positioned to track these intrinsic risks. How organizations govern AI in or near operations shows, more clearly than almost any other issue, what makes a retained risk defensible.

Oversight is a duty that cannot be delegated

Legal precedent sets a clear expectation for corporate boards. Under Delaware's Caremark doctrine, reinforced in Marchand v. Barnhill, corporate directors have a duty to make a good-faith effort to establish board-level monitoring and reporting systems for mission-critical compliance risks, and respond appropriately to significant warning signs. They are not required to understand every technical detail. They are required to ensure that information reaches them regularly and that they act on it. Public and municipal utilities answer to different legal frameworks, but the principle carries across: those who govern must know what is happening inside the operation.

Execution can be handed to an automated system. Fiduciary duty cannot. Many boards and councils have limited understanding of agentic AI's current state or likely trajectory, and that gap does not reduce their responsibility.

Oversight also cannot wait for a deliberate decision to deploy. IT and OT platform providers are adding agentic capabilities to version upgrades, often because they are competing with each other, so an organization that has chosen to proceed cautiously may find AI a single upgrade away. Security agencies have cautioned that agents in operations expand the attack surface, require broad system access to be useful, behave unpredictably, and make operational decisions harder to trace and verify.





Unlike companies in competitive markets, utilities are not in a race. They can wait, watch, and learn from early adopters in other sectors, start with lower-risk business applications, and move toward operations only when they understand what they are taking on.

What defensible oversight looks like

A defensible position rests on a few disciplines. Oversight of any agent that touches operations is mandatory. Directors and senior leaders need familiarity rather than mastery, which short, focused briefings can provide. Organizations should anticipate what could go wrong before attempting formal assessment, then play those scenarios out.

That includes liability drills. Directors and officers can imagine being sued after an AI-related failure and work through how they would respond before anything happens. A realistic scenario starts with success. Operators review each agent recommendation carefully at first, occasionally rejecting one. Over a year of good results, review time shrinks from minutes to seconds, and when the model drifts, an error passes straight through. The complacency described in Chapter 4 becomes a governance failure, and the question for the board is whether it could show that it knew the risk existed and had put controls around it.

Evidence decides whether a decision holds up
A decision is defensible only if the organization can later show what it decided, why, and what happened. AI makes that harder. After an event involving agents, can the organization reconstruct the sequence? Are there logs? Is there someone with enough knowledge of the affected process to explain it to leadership, regulators, and possibly a court? What if the evidence is inconclusive, or attribution is impossible? These questions need answers before an incident, not after it.

Detection belongs in the same plan.
Organizations need ways to identify intrinsic failures such as drift, hallucination, and emergent behavior, not only adversarial activity, and many have not yet developed them.

Engineering limits on what can go wrong

Insurers are working through the same problem. Their concerns span output failures, automated decisions with real-world consequences, physical damage in OT environments, a larger attack surface for adversaries, and a long technology supply chain running from model developers through platform providers to the operator. Their scenario modeling ranges from a world in which defense keeps pace with AI-enabled offense to a systemic event that drives claims and severity sharply upward.

The most useful question insurers and operators can ask is which physical processes an agent could change. Cyber-informed engineering, developed at Idaho National Laboratory and gaining adoption in the water and electric sectors, applies directly. Engineered controls that prevent a damaging change, regardless of what the software instructs, provide stronger assurance than additional security software. Risk engineers with industrial backgrounds are well placed to help operators build that case.



Key takeaways

- Treat every retained risk as a decision that must withstand scrutiny from boards, insurers, regulators, and courts.
- Ensure that leadership has a working channel for information about operational risk, including AI, since oversight cannot be delegated.
- Assume agentic capabilities may arrive through vendor upgrades, and set oversight before deployment rather than after.
- Prepare the evidence a defensible decision requires, including logs, process expertise, and the ability to reconstruct events.
- Use engineered controls to limit what any automated system can physically change, rather than relying on software alone.
- Build a defensible position now instead of waiting for regulation, and hold integrators to the same expectations.

Regulation will not settle the question

Regulators are unlikely to define defensibility for operators soon. In the United States, federal policy has so far emphasized AI adoption over new regulation. Sector regulation moves slowly: electric reliability standards took years to address cloud computing, and the federal agency responsible for the water sector has limited ability to mandate cybersecurity measures. In the EU, amendments adopted in 2026 postponed the AI Act's high-risk obligations to December 2027 for stand-alone systems, including those used in critical infrastructure, and to August 2028 for AI embedded in regulated products, while narrowing some of their scope. South Korea's AI Basic Act, in force since January 2026, treats energy as a high-impact sector, with a grace period before fines apply. In the meantime, operators will have to set their own standard for defensible AI use.

Some of the earliest movement is at the state level. On September 17, 2026, the New York Public Service Commission ordered electric, gas, and water utilities to inventory their AI use cases and disclose the policies governing them within 60 days, with recurring reports to follow. State utility commissions may become the first practical regulators of operational AI. Until rules mature, organizations must build their own defensible position.

Smaller utilities face a particular challenge. Many have no dedicated cybersecurity or IT staff and rely on a system integrator serving many clients. These questions should be put to that integrator directly. If it is unwilling or unable to engage, finding another provider becomes part of the risk decision.

Much of the risk that remains is inherited, built into the products and platforms organizations buy, sometimes arriving uninvited in the next upgrade. Making that risk defensible ultimately depends on what manufacturers build into their products in the first place.

Secure by Design Is a **Shared Responsibility**

What asset owners should expect from industrial technology manufacturers throughout the product lifecycle.

Introduction: One term, two expectations

Much of the risk critical infrastructure operators manage is inherited. It arrives inside the controllers, relays, platforms, and software they buy, and it stays for as long as those products remain in service, which in industrial environments is often decades. Reducing that inherited risk depends on how products are designed, built, and supported, and on what operators do with them once installed.

"Secure-by-design" appears on almost every product data sheet, but buyers and sellers often mean different things by it. An asset owner may read it as a promise that the product is secure and needs no further attention. A manufacturer may mean that the product was developed through a secure process, but still has to be deployed and operated correctly. Integration and configuration can also introduce vulnerabilities that were not present in the product as shipped, so a secure product can become an insecure system in the field. That gap in understanding creates a gap in expectations, and plenty of room for frustration on both sides.

Where the manufacturer's role ends and the owner's begins

For a manufacturer, secure by design means building security in from the development phase: testing products, remediating findings before delivery, and supporting the product afterward through vulnerability disclosure and incident response assistance. It also means providing clear guidance on what a secure deployment looks like.

Delivery is not the end of the security work. Once a product is installed, the asset owner is responsible for securing the environment around it, including the network, access controls, and operating practices, across the asset's life. Because industrial products are now networked and digital, neither party can do its part in isolation. The relationship should continue throughout the product's supported life, with explicit arrangements for what happens when support ends. Asset owners also need to be involved early, ideally before a product ships, so that both sides understand where their responsibilities meet.

Asking for evidence, not assurances

Most procurement still relies on certifications and contractual clauses. Both matter, but neither shows how a product was actually built or how it will be supported. Asset owners can ask for more as part of a shared responsibility model, including evidence of secure development requirements, product testing, and vulnerability handling.

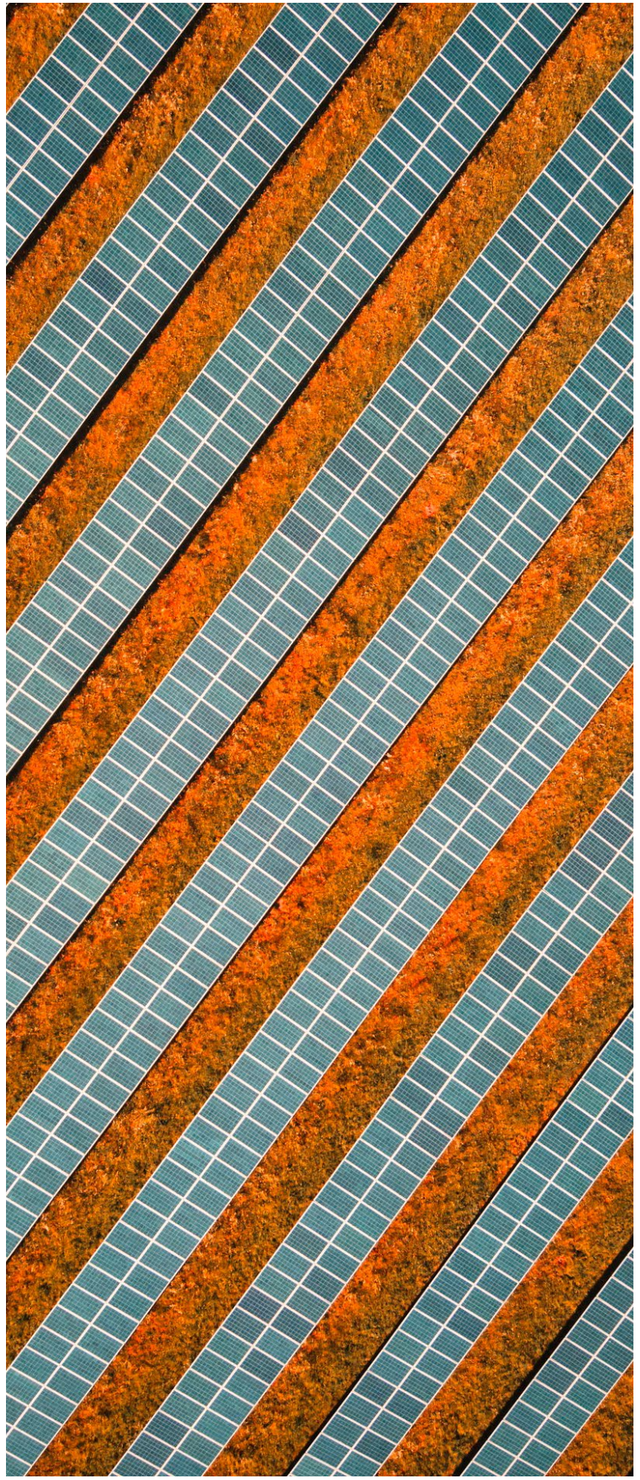

Supply chain security is among the hardest parts of product security, largely because of legacy products and long lifespans. Realism about what a supplier can deliver is essential. Two characteristics indicate a supplier worth working with. The first is a risk-based approach that covers legacy products as well as current ones. The second is transparency: a willingness to provide SBOMs, to state limitations openly, and to commit to timelines for meeting expectations. Certifications indicate maturity, but transparency is the more practical signal.

Contract requirements can go undelivered, not because a supplier is acting in bad faith, but because the capability does not yet exist. Simply requiring adherence to a set of cybersecurity requirements is not enough. An open line of communication, with demonstrable progress over time, is more useful. Transparency can also be misread. A manufacturer that publishes many vulnerability advisories may be among the more open suppliers rather than among the least secure.

A credible supplier response shows the core elements of a product security program, scaled to the supplier's size:
Governance: written policies and standards that are actually followed.
Threat modeling: an understanding of how the product could be attacked.
Secure development: security built into the development process, with testing and remediation before release.
Lifecycle support: a clear, committed support period, including for newer technologies that may not stand the test of time.
SBOMs: a willingness to share them.

Size matters on both sides

Smaller suppliers cannot do everything larger

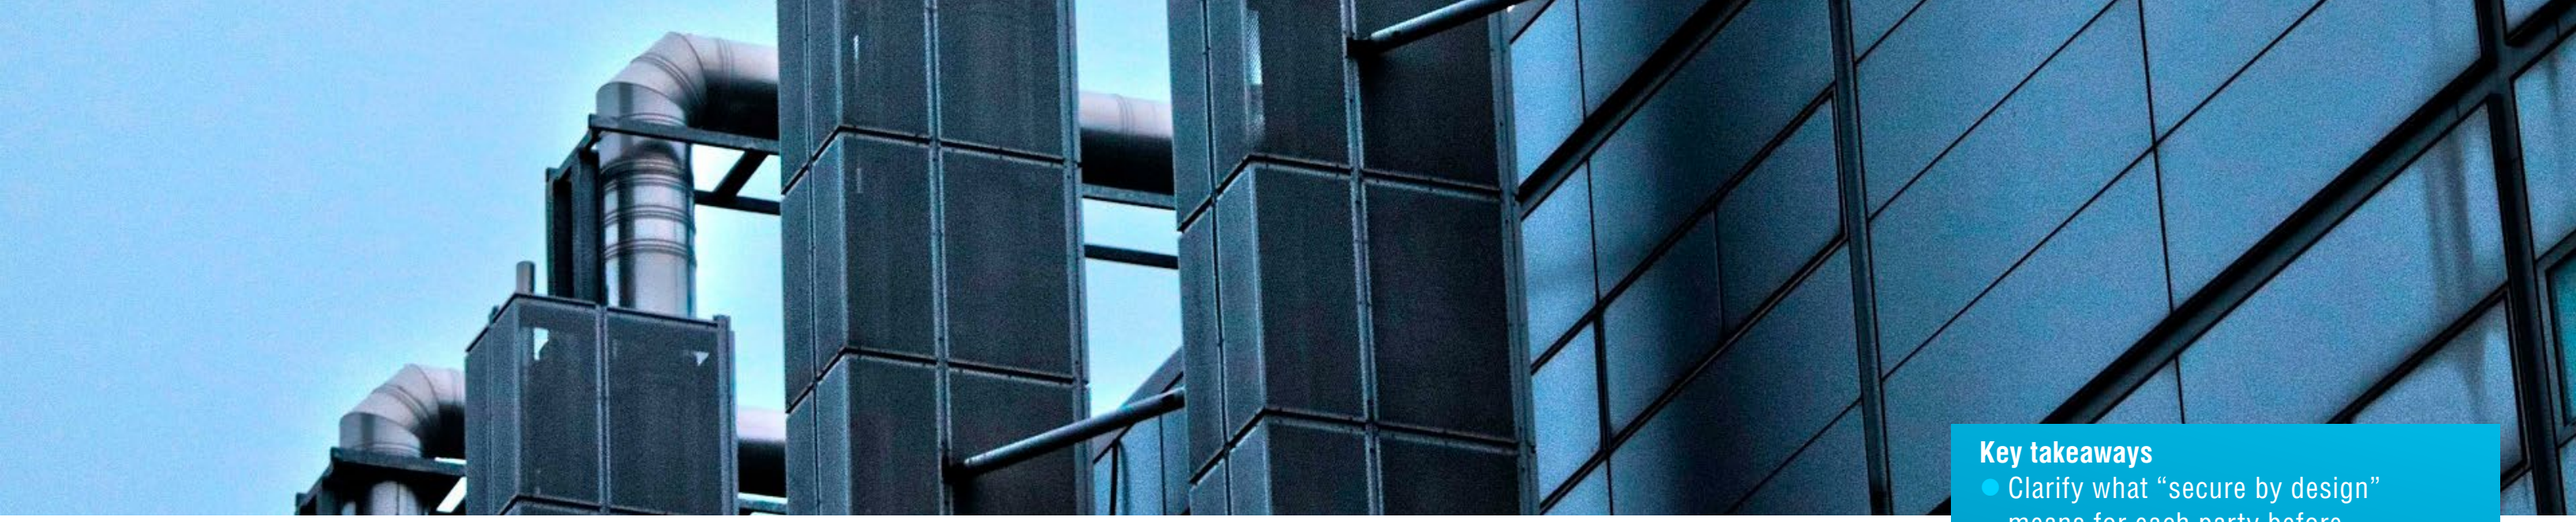

ones can, and smaller operators lack the purchasing leverage of large asset owners. A risk-based approach helps on both counts. Operators should assess how critical a supplied component is to their system and whether the supplier's controls align with their own risk appetite. Where a critical supplier falls short, particularly a smaller one with limited resources, working with it to raise its capability is often more effective than replacing it.

Supply chains in industrial sectors can be long, and they are only as strong as their weakest link. Clarity about who is responsible for what, agreed in advance among manufacturers, suppliers, integrators, and asset owners, is the foundation. Standards and guidance increasingly draw those lines, including the Supply Chain Cybersecurity Principles published by the U.S. Department of Energy in 2024.

Regulation raises the baseline

Regulation is changing economics. Meeting regulatory requirements should be a result of a strong security posture, not a goal pursued for its own sake. The EU Cyber Resilience Act, discussed in Chapter 6, establishes mandatory cybersecurity requirements for covered products with digital elements placed on the EU market. For global manufacturers, meeting those requirements can influence product development and security practices across their wider portfolios. Features that customers could once choose to pay for are becoming part of the baseline.

The uplift is substantial. A manufacturer with several business units producing different products must apply the same requirements consistently across all of them, even where some processes, such as security testing and reporting, already exist. Customers gain assurance, but some of the cost of that uplift is likely to be passed on. Whether the market comes to accept that cost, as it eventually accepted the cost of safety, remains to be seen. Public incidents that affect consumers tend to accelerate the shift, turning board-level attention into funding and reframing the conversation around how much risk exists and what it costs to reduce.

A relationship, not a feature

Gaps close fastest through collaboration. Manufacturers can engage their most important customers and suppliers to understand where each stands, identify gaps, and offer support. Neutral organizations, including government agencies and international forums, help by bringing asset owners and manufacturers together to discuss challenges that are not unique to any one organization. Because many manufacturers share the same suppliers and customers, that collaboration scales. Smaller asset owners can take part as well, both by contacting their manufacturers directly and by joining these forums.

Arrival of more capable AI models strengthens the case. More findings, more threats, and more opportunities to improve defenses all demand closer working relationships. Trust established before an incident also makes incidents easier to manage, because the parties already know how to work together when it matters.

Product security is not a feature delivered at the point of sale. It is a relationship maintained across engineering, procurement, deployment, operation, maintenance, and eventual retirement. Asset owners can ask more of their suppliers, regulation will ask more of manufacturers, and manufacturers benefit from hearing more from their customers. Critical infrastructure resilience depends on all three.

Key takeaways

- Clarify what "secure by design" means for each party before purchase, so that manufacturer and owner responsibilities are understood and agreed.
- Treat product security as a lifecycle relationship that continues after delivery, through vulnerability disclosure, support, and joint incident response.
- Ask suppliers for evidence of governance, threat modeling, secure development, lifecycle support, and SBOMs, rather than relying on certificates and contract clauses alone.
- Value transparency, including openly published vulnerabilities and candid statements of current limitations, as a sign of supplier maturity.
- Apply a risk-based approach to suppliers based on component criticality, and help critical smaller suppliers improve rather than simply replacing them.
- Use neutral industry and government forums to build collaboration across shared suppliers and customers, and establish trust before an incident requires it.

Decisions That Hold **Under Pressure**

This handbook opened with a simple question: what does a security decision protect? Ten chapters later, the answer is clearer, and so is how much depends on asking it first.

The chapters describe one discipline rather than ten separate topics. Knowing which outcomes cannot be lost gives meaning to the question of who decides. Those decisions are only as sound as the organization's view of its environment, including systems beyond the traditional OT boundary and the AI now entering it. Control of external pathways and clear supplier obligations narrow the exposure. Recovery and continuity planning accept that prevention will sometimes fail. Defensible risk decisions account for what remains, and secure design reduces what the next generation of systems will carry. Together, these decisions form a more coherent approach to protecting critical infrastructure.

Several tensions remain unresolved. Responsibility is increasingly shared among owners, operators, integrators, and manufacturers, yet every critical decision still needs a named owner. Automation promises efficiency while making meaningful oversight harder to sustain. Regulation raises the baseline but cannot substitute for operational control, and compliance alone can still breed complacency. No single standard or product will settle these.

What organizations can do is ask better questions. The following are a useful starting point:

- Which outcomes can we not afford to lose, and does our security program reflect them?
- Who has the authority to act when security and operations disagree, and is it written down?
- What is in our environment that we cannot see, including AI that arrived without approval?
- Who can reach our operations from outside, and could we disconnect them without stopping?
- Have we proven that we can recover, and how long could we keep essential services running if we could not?
- Could we explain to a board, regulator, or insurer why we accepted the risks we still carry?

Answering them takes time, budget, and cooperation across teams that do not always share a vocabulary. It also takes cooperation across organizations. The relationships described in these pages, between owners and suppliers, operators and integrators, practitioners and regulators, are what turn individual effort into resilience across the sector.

Our thanks go to every speaker, panelist, sponsor, and participant who shared what has worked and what has not, and to the community partners who help bring practitioners together. Industrial Cyber will continue that conversation in the year ahead, because the systems we depend on will keep changing, and so will the decisions needed to protect them.

Anna Ribeiro, Editor, Industrial Cyber

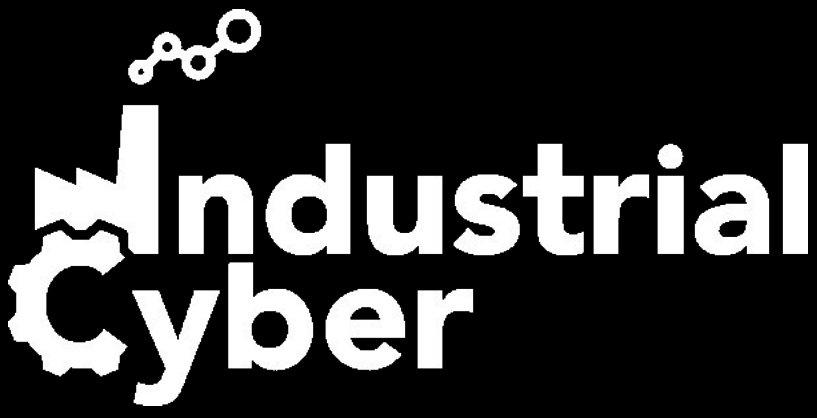

Critical Infrastructure Handbook 2027

DRAGOS DISPEL VYRON Macrium Software

Industrialcyber.co