The Operational Context Gap

Rethinking OT Cyber Risk Beyond CVSS

Velos Shield Point of View Whitepaper Series • August 2026

Executive Summary

Industrial organisations can now see more of their operational technology than at any point in the last two decades - yet CISOs and plant leaders still cannot say, with confidence, which cyber conditions actually threaten their operation right now. This is not a visibility failure. CISA has published over 12,000 known OT/ICS vulnerabilities since 2010, and severity keeps climbing. The problem is that severity is not consequence, and no amount of scanning changes that: a critical finding on an isolated, redundant display and an identical finding on a single-point-of-failure safety controller are not the same risk.

Velos Shield calls this divide the Operational Context Gap - the distance between what cybersecurity tools can technically describe and what operations, engineering and executive leadership need to know before that data becomes a decision. Closing it means deliberately building the context that has always lived only in experienced engineers' heads: what an asset does, what depends on it, and what would genuinely happen if it failed.

> *"Visibility tells an organisation what exists. Context tells it what matters. Intelligence tells it what to do about it."*

This paper examines why CVSS, alert volume and generic criticality all break down inside OT - a conclusion CISA itself reached in June 2026 when it formally retired mandated CVSS-based prioritisation. It walks through five documented incidents to show that context, not attacker sophistication, determined each outcome, and introduces the Operational Context Stack: a six-layer framework - Asset, Relationship, Process, Criticality, Threat, Consequence - for building that context, on the path to what Velos Shield calls Cyber-Aware Operations.

1. The OT Security Visibility Paradox

Over the past decade, OT security has largely solved the problem it set out to solve. Passive, protocol-aware monitoring now identifies most devices on a well-instrumented network automatically, and CISA's advisory pipeline has scaled from a standing start in 2010 to publishing over 500 ICS advisories in a single year for the first time last year.

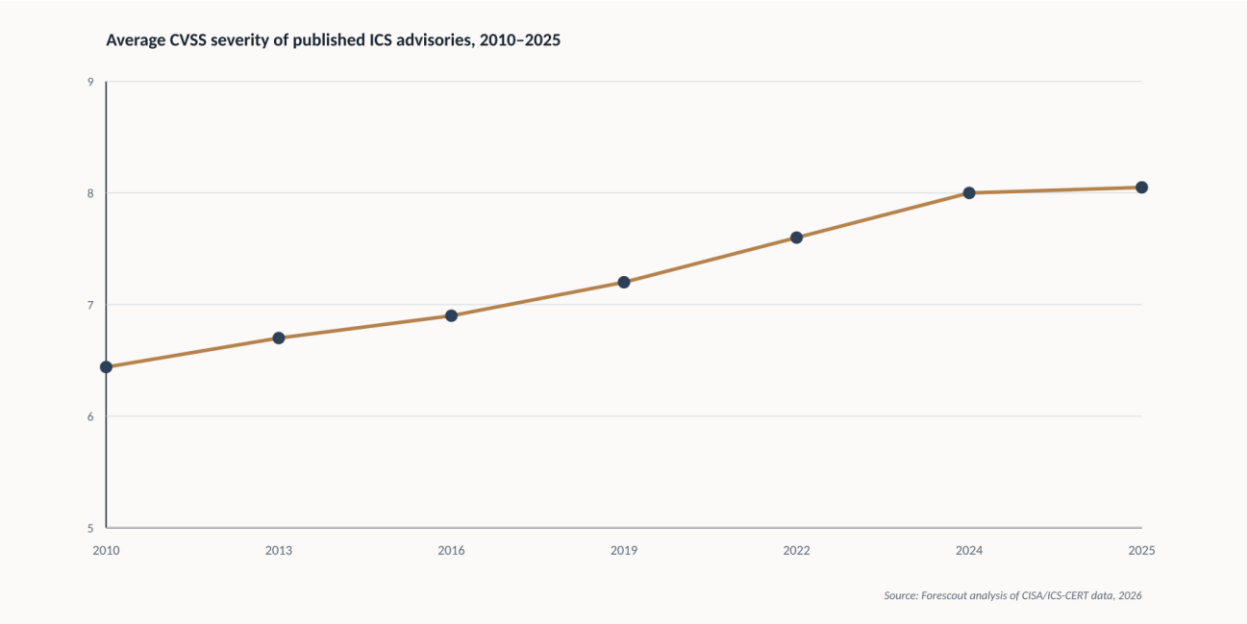

Figure 1 - Average CVSS severity of published ICS advisories has risen steadily since 2010.

Yet Independent OT security assessments find that roughly 45 percent of engagements still uncover fundamental visibility gaps. Both facts are true at once: visibility has genuinely improved in aggregate, while the gap between what is visible and what is understood has widened, because information has grown faster than any organisation's capacity to interpret it. This is the paradox - more sensors and advisories have not produced a proportional increase in knowing what actually matters.

2. Seeing an Asset Is Not Understanding an Asset

Five capabilities are routinely treated as one: Discovery (something is there), Identification (it has a name), Classification (it has a type), Contextualisation (it has a place), and Operational Understanding (it has a meaning).

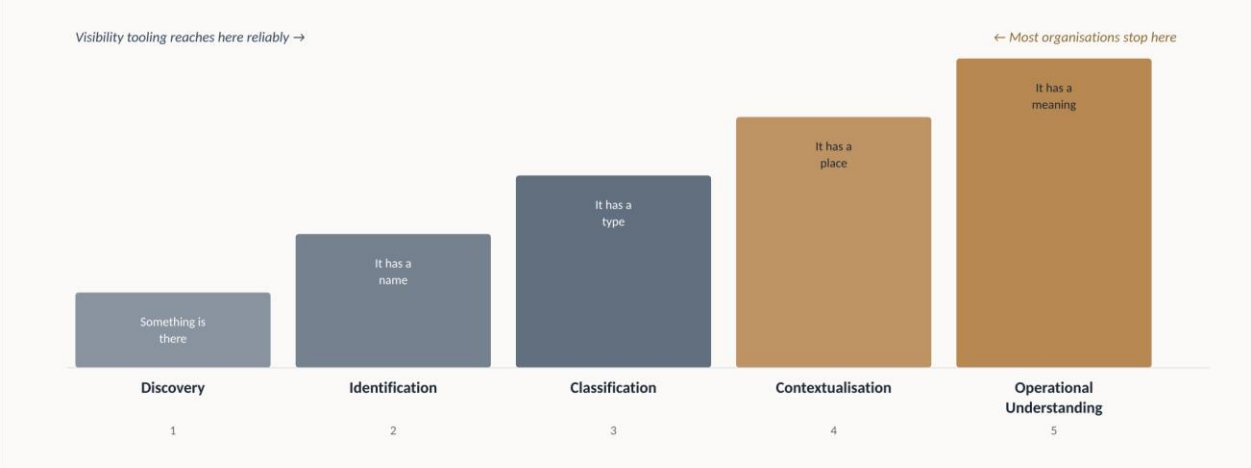

Figure 2 - Five distinct capabilities routinely mistaken for one: seeing an asset.

Modern tooling is genuinely strong at the first three. A sensor can identify a controller's exact model and firmware - but it cannot tell you, unaided, that this specific controller governs an emergency shutdown sequence with no manual backup, while an identical unit two buildings away controls parking-lot lighting. That gap between classification and operational meaning is where most OT risk programmes quietly fail.

3. The Operational Context Gap

The Operational Context Gap is the distance between what cybersecurity data can technically describe and what operations, engineering and executive functions need to know before that description becomes a decision. It exists because the information that resolves it - process role, failure mode, redundancy - has always lived in engineering drawings and operator memory, not in any system security tooling can query.

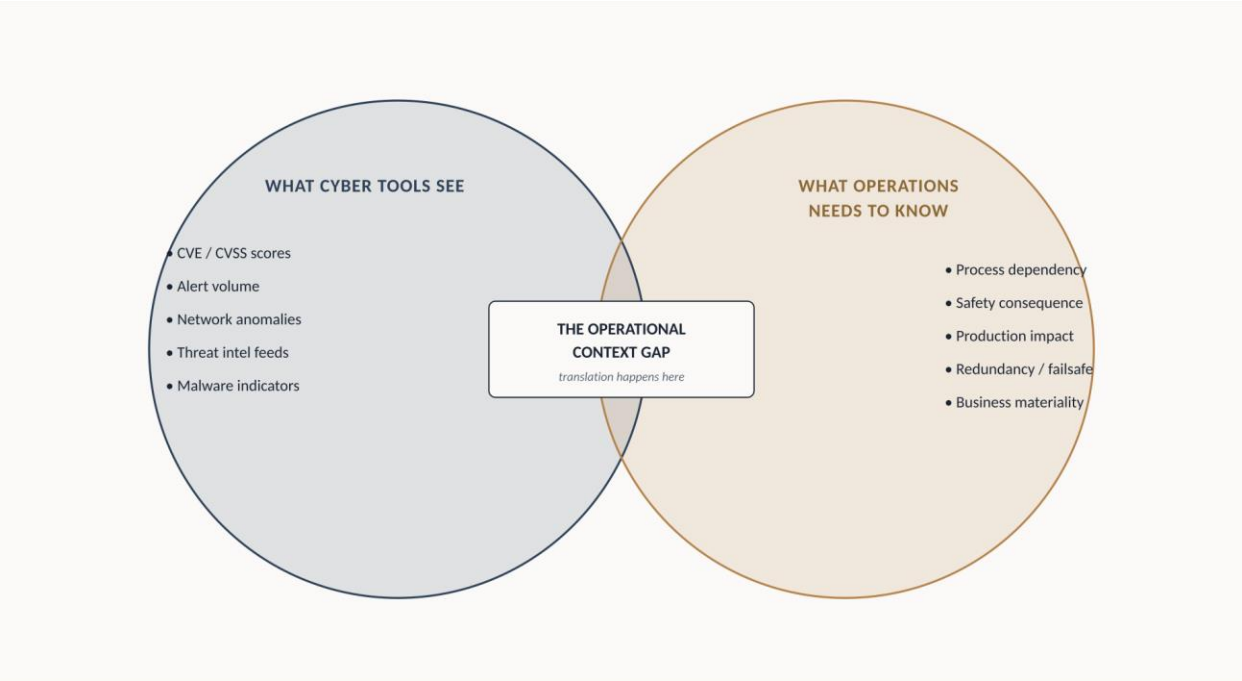

Figure 3 - The Operational Context Gap: the missing translation layer between cybersecurity data and operational reality.

Each stakeholder feels the gap differently: CISOs cannot translate a vulnerability count into a board-level business case; SOC analysts face alert fatigue with no reliable way to separate routine anomalies from real threats; plant managers see patch requests disconnected from operational reality. Each is a symptom of the same condition - more cyber data than operational meaning attached to it.

4. Why Traditional Risk Models Break Down in OT

None of the standard mechanisms are wrong - each solves a real problem. Each simply becomes insufficient the moment it is asked to carry a full OT prioritisation decision alone.

CVSS and severity scores

CVSS describes theoretical severity, independent of environment. In June 2026, CISA's Binding Operational Directive 26-04 formally retired mandated CVSS-based prioritisation for federal vulnerability management, replacing it with exposure, exploitation and mission-impact questions - the same questions operational context is built to answer.

Vulnerability counts, generic criticality, and threat intel

Raw counts measure exposure trends but not what to do today - Industry analysis found only about 3 percent of one year's ICS vulnerabilities required truly immediate action once context was applied. Static criticality ratings miss that an asset's importance shifts with production schedule and redundancy. And sector-level threat intelligence says who is out there, not whether they can reach a specific asset or whether it would matter if they did.

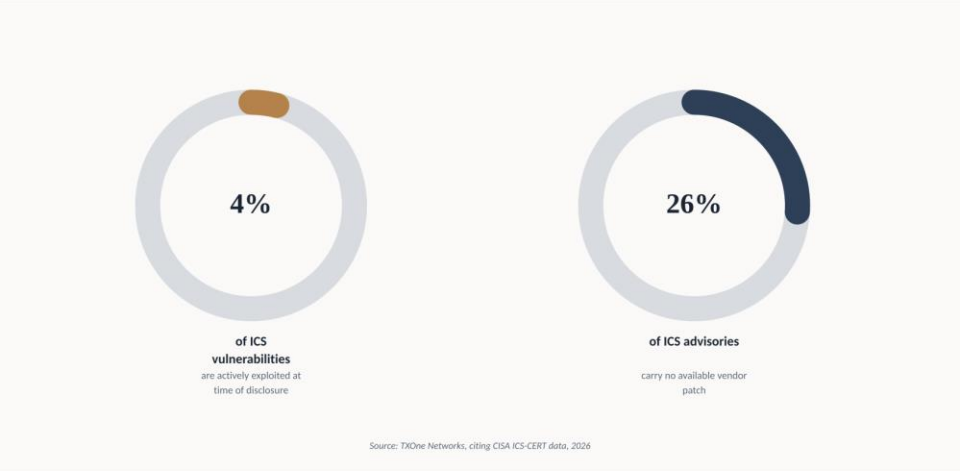

Figure 4 - A small fraction of ICS vulnerabilities are actively exploited; over a quarter have no available patch. Severity alone cannot triage this.

5. Cyber Risk vs. Operational Risk

Hold a technical finding constant and let operational context change, and the gap becomes obvious. A CVSS 9.8 on an isolated HMI duplicated by physical gauge panels is a low-priority finding. The identical CVSS 9.8 on an HMI with vendor remote access, on a flat network shared with IT, and no manual override for a pressure-relief valve, is an emergency.

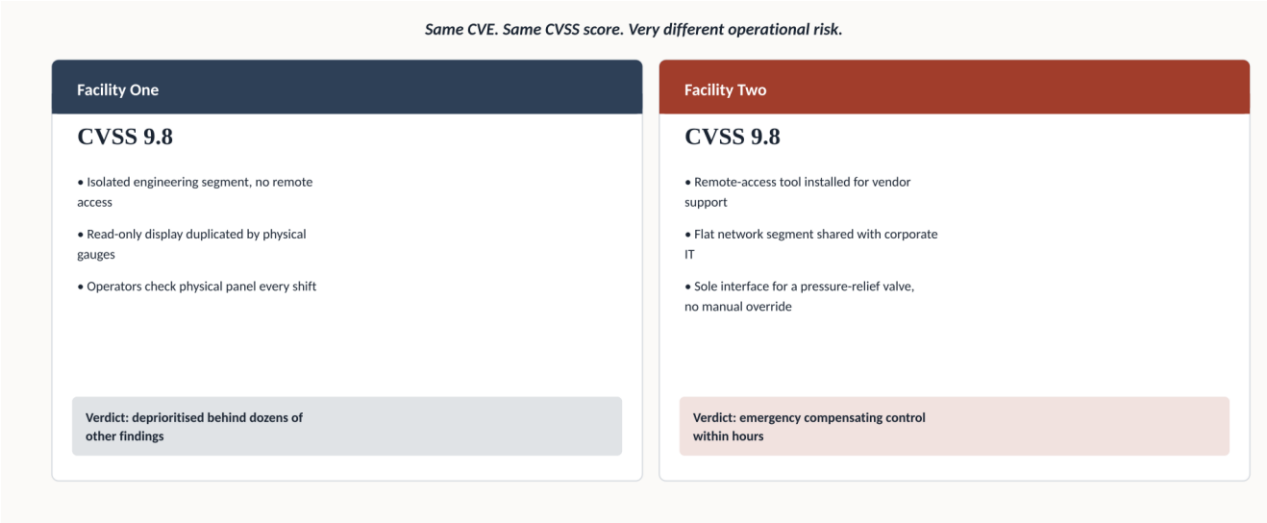

Figure 5 - Same CVE, same CVSS score, very different operational risk.

Two assets can carry identical technical findings and warrant entirely different responses - and that difference is invisible to any tool that stops at the technical layer.

6. The Operational Context Stack

The Operational Context Stack is Velos Shield's framework for closing the gap: six layers - Asset, Relationship, Process, Criticality, Threat, Consequence - ordered from what is easiest to automate today to what requires the most cross-functional judgment.

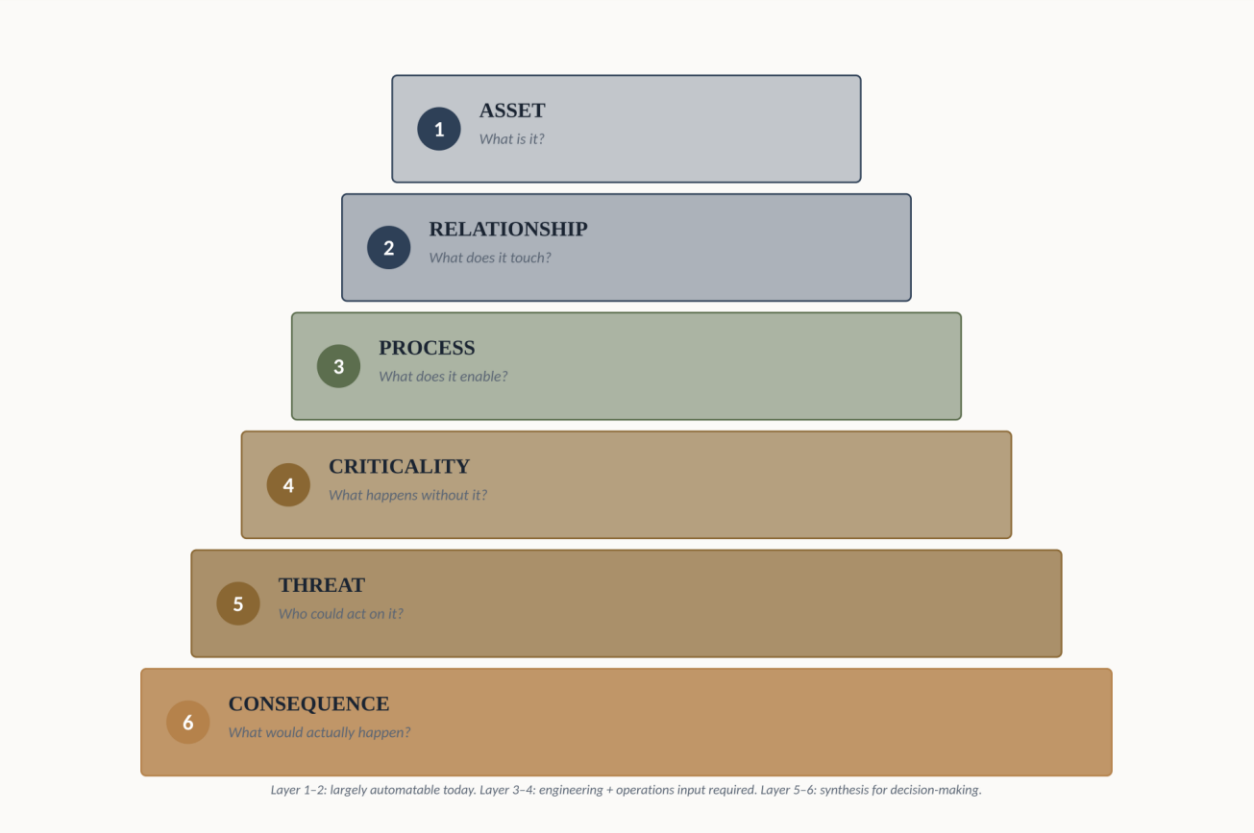

Figure - The Operational Context Stack: Velos Shield's six-layer framework.

Most organisations can answer confidently through Relationship, because network topology and reachability - what can logically reach a given controller, and how - are the layer most existing tooling is already built to answer. Very few can answer Consequence, in board-ready terms, without days of manual work per asset. That gap between layer two and layer six is the Operational Context Gap in structural form - and closing it is the subject of the rest of this paper.

7. From Technical Signals to Operational Intelligence

Most organisations already own the raw material they need - OT telemetry, asset inventory, engineering documentation, threat intelligence, business criticality - sitting in different systems, owned by different teams. Simply collecting more data widens the gap rather than closing it; each new feed adds another dashboard for an analyst to manually reconcile.

Figure 6 - Ten data sources fused into one operational context layer, distributed as role-specific views.

Data fusion, not data collection, is the actual unlock: a persistent layer that automatically joins telemetry to the process it belongs to, so a new finding inherits its operational meaning the moment it appears, rather than waiting for a human to reconstruct it from scratch.

8. The Human Context Problem

Even perfect data fusion would not fully close the gap, because part of it is organisational. SOC, cybersecurity, engineering, operations and executive leadership routinely look at the same environment and see materially different things, trained and measured on different metrics.

Figure 7 - Five functions, five vocabularies for describing the same industrial environment.

Multiple research shows the pattern: OT incident response plans are often absent, engineers rarely trained in cyber response, and assumed IT/OT segmentation frequently fails penetration testing. None of this is a technology gap - it is a shared-vocabulary gap between the people who understand the process and the people who understand the threat.

Two lessons follow directly. First, IT and OT teams do not primarily need more shared data; they need shared insight built around what an asset does operationally, not just what a scan reports - the distinction that separates a vulnerability count from an operational decision. Second, security tooling itself has to evolve to feed structured operational context into these layers as a native output, rather than leaving each team to reconcile spreadsheets and dashboards on its own after the fact.

9. A New Model for OT Risk Prioritisation

Rather than a false-precision scoring formula, Velos Shield proposes Operational Materiality: a qualitative, banded model moving the central question from "what is technically severe" to "what could materially affect the operation," built directly on the Stack's layers - reachability, process role, failure consequence, plausible adversary, and timing.

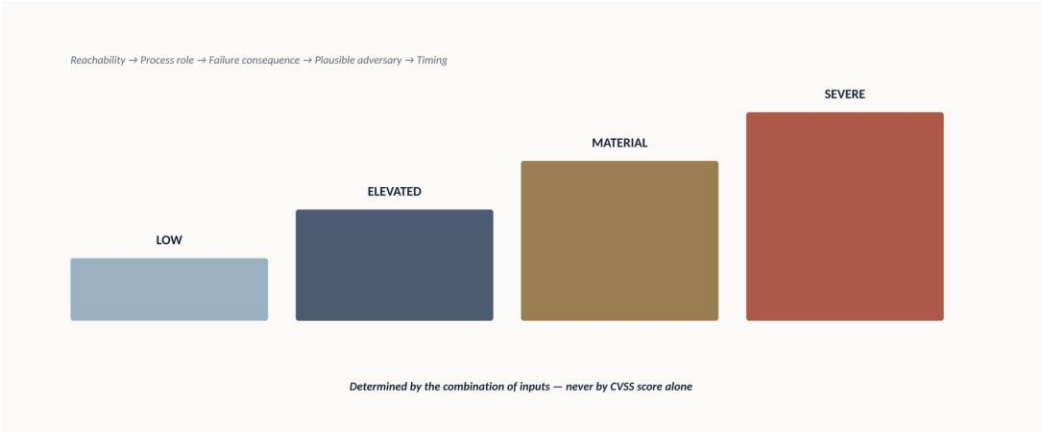

Figure 8 - Operational Materiality bands, determined by combined context, never by CVSS alone.

Findings land in one of four bands - Low, Elevated, Material, Severe - based on the combination of answers, not any single input. This keeps everything useful about CVSS and threat intelligence as supporting evidence, while refusing to let either stand in for the full judgment.

10. From OT Security to Cyber-Aware Operations

Cyber-Aware Operations is the destination state: continuous, automatic understanding of what is happening, where, what it affects, why it matters now, and what to do about it - without reconstructing that picture manually after every incident.

> *"Cyber-Aware Operations is not a bigger SOC. It is an operation that understands its own cyber risk the way it already understands its own safety risk."*

The comparison to safety culture is deliberate: industrial organisations don't treat safety as a periodic audit, they build it into daily operating rhythm. Cyber-Aware Operations argues operational cyber risk deserves the same structural treatment - a continuously available layer of understanding, not a quarterly report.

11. What the Future OT Security Architecture Should Look Like

Five principles, deliberately at the level of architecture rather than product: context as a first-class, persistent layer rather than a spreadsheet generated for audits; fusion over collection; bidirectional flow so process knowledge moves from engineering into the context layer as naturally as threat data moves the other way; role-specific views drawn from one shared foundation rather than competing versions of "risk"; and design for degradation - recognising when a compensating control already exists, rather than treating every technical exposure as equally undefended.

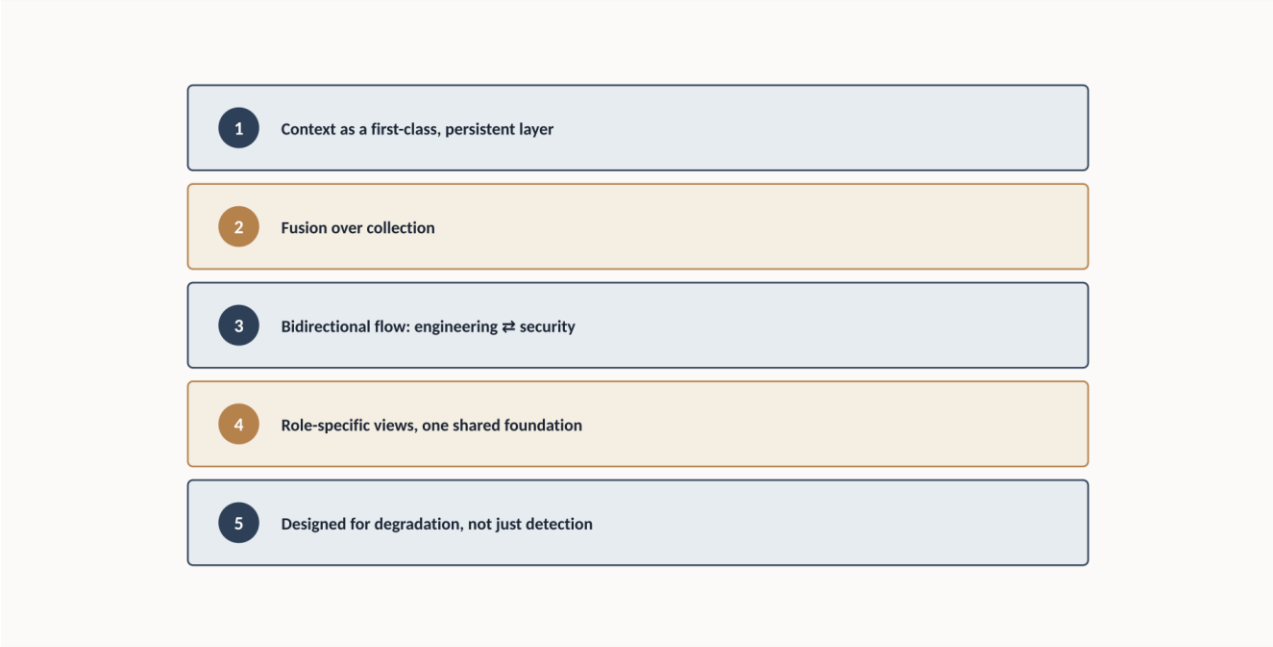

Figure 9 - Five architectural principles for a context-first OT security programme.

12. A Practical Roadmap for Industrial Organisations

A realistic, five-phase path: Assess maturity against each Stack layer (months 1–2); Contextualise the smallest set of assets carrying the largest share of risk (months 2–6); Prioritise using Operational Materiality bands even manually at first (months 4–8); Operationalise into a persistent, automatically updating context layer (months 8–16); and Continuously Learn - treating every incident and near-miss as an input back into the context layer, indefinitely.

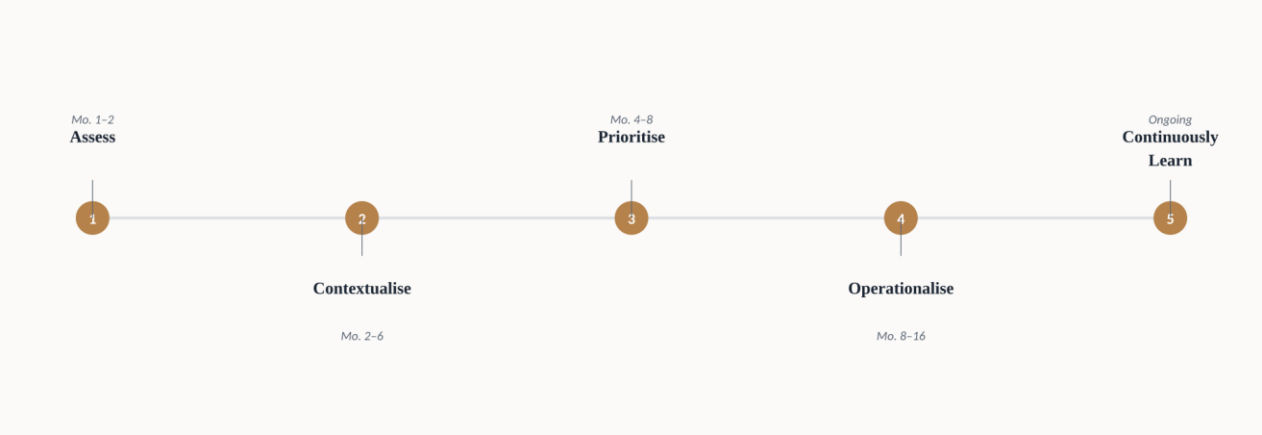

Figure 10 - A realistic five-phase roadmap from initial assessment to continuous operational context.

Organisations that skip the deliberate Contextualise phase consistently underestimate how much of this work is organisational rather than technical.

Case Studies: Five Documented Incidents Through the Operational Context Lens

Reconstructed from named, independently verifiable public sources. Documented facts are presented as such; interpretive analysis of the operational-context dimension is Velos Shield's own reading.

Figure 11 - Five documented OT incidents, 2015–2021, each analysed through the operational-context lens.

Colonial Pipeline (2021)

An IT-side ransomware event that shut down 45 percent of the U.S. East Coast's fuel supply - without the attackers ever reaching the pipeline's control systems.

- May 6–7, 2021: attackers used a compromised VPN credential lacking multi-factor authentication to access Colonial Pipeline's IT network and exfiltrate roughly 100 GB of data within a two-hour window (U.S. Department of Energy; TechTarget).
- The DarkSide ransomware group then encrypted IT systems including billing and accounting, but did not compromise the operational technology network that physically moves fuel (TechTarget; National Security Archive).
- Colonial proactively shut the pipeline down as a precaution - partly because billing systems needed to invoice deliveries were down, partly to ensure the attack could not spread into OT - halting transport of more than 100 million gallons of fuel per day (DOE; case-file synthesis of public reporting).
- The company paid roughly 75 bitcoin (about $4.4 million); the U.S. Department of Justice later recovered roughly 63.7 bitcoin, worth about $2.3 million after a price decline (public DOJ reporting, as summarised in incident retrospectives).
- A federal emergency declaration followed on May 9; the pipeline resumed full operations May 12–13, 2021 (Department of Energy; CHDS timeline).

What happened

A single leaked VPN password, without MFA, gave attackers a foothold in Colonial's IT environment. They exfiltrated data and deployed ransomware against business systems. The pipeline's OT network - the part that actually moves fuel - was never directly compromised.

What technical signals existed

The technical signal was, in the narrowest sense, an IT ransomware event: encrypted billing and accounting systems, a compromised remote-access credential, and no MFA on that access path. Nothing in the initial technical picture described a physical-consequence event.

What operational context mattered

The operational context that mattered was almost entirely about dependency and precaution rather than direct compromise: Colonial could not reliably bill for fuel delivered, and could not be certain the ransomware would stay confined to IT. In the absence of a way to state confidently "OT is segmented and unaffected," the only responsible operational decision was to shut down a pipeline moving nearly half of the East Coast's fuel supply. The consequence was driven not by what the attackers actually reached, but by what the organisation could not immediately prove they hadn't.

What could have changed the outcome

An organisation able to state, in real time and with confidence, exactly how segmented its OT environment was - and to demonstrate that segmentation was holding - would have had a materially different decision to make than one forced to assume the worst. This is a direct illustration of the Relationship layer of the Operational Context Stack: the shutdown decision was, in effect, a decision made in the absence of reliable relationship context between IT and OT.

The lesson

Consequence can flow from an IT-only event when operational confidence in segmentation is itself uncertain. Closing the Operational Context Gap is not only about understanding OT assets; it is about being able to prove, quickly and credibly, what an incident did not reach - which requires the same relationship and process context this paper argues for, applied defensively during a live event rather than only during routine prioritisation.

Triton / Trisis (2017)

Purpose-built malware that reprogrammed a petrochemical plant's last line of physical defence - and went undetected for two months because the first warning sign was investigated as a mechanical fault, not a cyber event.

- Triton (also known as Trisis or HatMan) targeted Triconex safety instrumented systems (SIS) - controllers whose sole function is to force a plant into a safe state when a process reaches dangerous conditions (Wikipedia; FBI/CISA joint advisory).
- A first, unrecognised incident in June 2017 knocked an emergency shutdown system offline; the plant's vendor-led investigation treated it as a mechanical or engineering fault rather than a cybersecurity event (Dark Reading, citing incident responder Julian Gutmanis).

- A second incident in August 2017 triggered a genuine plant shutdown when Triconex controllers detected an anomaly caused by a bug in the attackers' own malware - the plant survived not because the attack failed, but because the malware itself malfunctioned (FBI/CISA advisory).
- Later reinvestigation found six compromised safety controllers, not the two initially reported, and established the attackers had been present on the corporate network for roughly a year before reaching the SIS (Dark Reading; independent forensic investigators).
- The United States later publicly attributed the malware to TsNIIKhM, a Russian state research institute, and unsealed an indictment of an affiliated individual in 2022 (FBI PIN advisory).

What happened

Attackers spent roughly a year moving from initial IT access through the corporate network into the OT environment, ultimately reaching engineering workstations that communicated with the plant's safety instrumented systems, and installed malware built specifically to reprogram those controllers.

What technical signals existed

The June 2017 shutdown was itself a technical signal - an emergency shutdown system had gone offline unexpectedly - but it was diagnosed through a mechanical and engineering lens, not a cybersecurity one, because the investigating team did not treat SIS anomalies as a category that warranted cyber forensic review.

What operational context mattered

The decisive context here is organisational, not technical: the plant had engineering and safety expertise capable of investigating a shutdown, and cybersecurity expertise capable of investigating an intrusion, but no established process connecting the two - no one asked whether a safety-system anomaly might also be a security incident. This is the Threat and Process layers failing to meet: the people who understood the safety process did not have a working channel to the people who understood adversary behaviour, and vice versa.

What could have changed the outcome

A documented process requiring cyber-forensic review of any unexplained safety-instrumented-system event - treating SIS anomalies with the same cross-functional urgency as a safety incident - would very plausibly have surfaced the June intrusion two months earlier, before the attackers achieved the persistence and access that led to the August event.

The lesson

Safety systems are not exempt from cyber risk because they are safety systems; if anything, their consequence profile makes them the highest-priority candidate for the Consequence layer of the Operational Context Stack. And the case is a direct illustration of Chapter 8's argument: the human, organisational gap between engineering and cybersecurity can be exactly as costly as any technical vulnerability.

Ukraine Power Grid Attacks (2015–2016)

Two attacks, one year apart, on the same national grid - the first a proof that a cyberattack could turn off power at scale, the second a demonstration that purpose-built malware could do it without a human at the keyboard.

- December 23, 2015: attackers who had used BlackEnergy3 malware, delivered via spear-phishing months earlier, remotely operated three Ukrainian distribution utilities' SCADA systems to switch off roughly 30 substations, cutting power to approximately 230,000 customers for one to six hours - the first publicly acknowledged successful cyberattack on a power grid (Wikipedia, sourced to ICS-CERT and SANS ICS reporting).
- The attackers also disabled backup power at operator control centres and used destructive KillDisk malware and a denial-of-service attack on the utilities' call centres to slow the human response and delay customer awareness of the outage (Military History Wiki, sourced to ICS-CERT analysis).
- December 17, 2016: a second attack against a Kyiv transmission substation used Industroyer (also called CrashOverride), malware built with native knowledge of grid-control communication protocols, allowing it to directly manipulate substation equipment without relying on the utility's own operator software - cutting about a fifth of Kyiv's power for roughly an hour (Wikipedia; independent OT threat analysis).
- Independent analysis of the 2016 event concluded it was not simply intended to cause a short outage but appeared to be a deliberate, large-scale capability test, including an attempt to cause longer-lasting physical damage when operators restored power (Wikipedia, citing independent security research and Wired reporting).

What happened

In 2015, human operators, using stolen credentials and remote access gained through months of prior reconnaissance, manually operated grid-control software to open breakers and cut power, then worked to slow the utilities' recovery. In 2016, purpose-built malware operated grid protocols directly, without needing to imitate a legitimate operator at all.

What technical signals existed

In 2015, the technical signal was a spear-phishing campaign and BlackEnergy3 malware present on utility networks for months prior to the attack - reconnaissance activity that, with the right monitoring and threat context, was in principle detectable well before the outage itself. In 2016, the signal was a malware framework engineered specifically for grid-control protocols, a category of technical capability that did not previously need to be defended against because it had not previously existed in the wild.

What operational context mattered

The operational context that mattered most across both events was process and consequence-layer understanding of grid topology and cascading dependency: which substations, if switched off simultaneously, produce the largest customer impact, and what redundancy or manual restoration capability exists if automated systems are compromised. Ukrainian utility operators' ability to manually restore power -

reportedly faster than utilities with more automated, less manually operable systems might have managed - was itself a form of operational resilience that blunted the 2015 attack's ultimate impact.

What could have changed the outcome

Earlier detection of the BlackEnergy3 reconnaissance phase - months of dwell time before the 2015 attack - represents the clearest missed opportunity; this is squarely a Threat-layer and Temporal-layer gap, where the presence of a known malware family on utility networks was not connected to the elevated geopolitical threat context of the period.

The lesson

Purpose-built ICS malware, once demonstrated as viable in 2016, becomes a template other adversaries can adapt to other grids and other sectors. Threat context cannot be treated as static once a capability has been proven in the wild; the Consequence layer of any similarly structured facility should be revisited the moment a comparable attack succeeds anywhere in the world.

Norsk Hydro (2019)

A ransomware attack that never touched a single piece of industrial control equipment directly - and still cost one of the world's largest aluminium producers tens of millions of dollars by forcing a return to manual operations.

- LockerGoga ransomware began infecting Norsk Hydro's systems on March 18, 2019, and the company disclosed the attack publicly the next day, describing it as impacting operations across several business areas (SecurityWeek; AXA XL).
- Hydro isolated all plants and switched to manual operations and procedures across roughly 160 plants in 40 countries, instructing its approximately 35,000 employees not to connect devices to the company network (AXA XL).
- The company stated it had strong backups in place and would not pay the ransom - an unusually transparent public posture that included regular press conferences disclosing operational and financial impact as it unfolded (published incident reporting; BankInfoSecurity).
- Hydro's own disclosed estimates put the first week's financial impact at $35–41 million, with the company's Extruded Solutions division - the hardest hit - running at 70–80 percent capacity a week later; cumulative first-half impact approached $71 million by some later estimates (SecurityWeek; Computer Weekly; public incident reporting).
- The company's core profit fell 82 percent in the first quarter of 2019, though the loss was contained enough that it fell within a range some analysts had feared could be worse, and Hydro shares rose following the disclosure (SecurityWeek).

What happened

Ransomware spread across Hydro's IT environment, and - as a defensive measure rather than because the malware itself reached production equipment - Hydro isolated its plants and reverted large portions of its operations to manual control to avoid any risk of the ransomware reaching OT systems.

What technical signals existed

The technical signal was file-encrypting ransomware behaviour consistent with LockerGoga, spreading across IT infrastructure, with no public reporting suggesting the malware itself was designed to operate against industrial control systems.

What operational context mattered

The decisive operational context was Hydro's own honest assessment of its IT/OT relationship: rather than assume segmentation would hold, the company chose to isolate plants proactively and accept the cost and disruption of manual operations rather than risk finding out the hard way that segmentation had a gap. This is the same Relationship-layer caution visible in the Colonial Pipeline case, but exercised with more operational preparedness - Hydro's plants, unlike many organisations', retained the trained personnel and documented procedures needed to actually run manually at scale.

What could have changed the outcome

The case is notable less for what could have prevented the attack and more for what limited its consequence: strong backups, a genuine ability to operate manually, and transparent, fast internal communication about the scope of the incident. An organisation without well-rehearsed manual operating procedures facing the same ransomware event would very plausibly have faced a longer, more damaging shutdown.

The lesson

Operational resilience - the ability to fail gracefully into manual, process-safe operation - is itself a form of Consequence-layer risk mitigation, and belongs inside any serious Operational Context Stack assessment. The gap this paper describes is not only about detecting and prioritising threats; it is also about knowing, in advance, how much operational shock absorption an organisation actually has.

Oldsmar, Florida Water Treatment Intrusion (2021)

The clearest documented example in this set of context - specifically, an alert human operator - succeeding where technology alone would not have been enough.

- On February 5, 2021, an intruder gained remote access to the Bruce T. Haddock Water Treatment Plant's SCADA workstation using remote-access software shared among plant staff, and briefly accessed the system that morning without incident (Idaho National Laboratory CyOTE case study; Washington Post).
- At approximately 1:30 p.m., the same intruder returned, and a plant operator watched his mouse cursor move on its own for three to five minutes as the intruder opened various software functions before changing the plant's sodium hydroxide (lye) dosing setpoint from 100 parts per million to 11,100 parts per million (Washington Post; Pinellas County Sheriff's Office, via CBS News).
- The operator, who was actively watching his screen, immediately reversed the setpoint change and notified his supervisor; the plant's remote-access software was disabled shortly afterward (CyOTE case study; Manufacturing.net).

- Officials stated that even without the operator's intervention, redundant pH alarms and the physical time required for the chemical change to reach the treated water supply - estimated at 24 to 36 hours - would have provided additional opportunity to catch and correct the change before any water reached the public (CyOTE case study; Logical Systems Inc. lessons-learned report).

What happened

An intruder used a shared remote-access credential to reach the plant's human-machine interface twice in one day, and on the second attempt manipulated the sodium hydroxide dosing setpoint to a level officials described as potentially dangerous if it had reached consumers.

What technical signals existed

The technical signal was, in the most literal sense, a moving mouse cursor with no operator input behind it - an extremely low-fidelity indicator with no automated detection system flagging it as anomalous. The plant's technology stack did not detect or block the unauthorised access or the setpoint change; a human watching the screen did.

What operational context mattered

The context that mattered was entirely at the Consequence and Criticality layers, and it worked exactly as designed: the plant had redundant pH alarms independent of the compromised HMI, and a physical process delay between a chemical dosing change and its arrival in the treated water supply, both of which existed specifically because water treatment engineers had long understood sodium hydroxide dosing as a safety-critical process requiring layered, non-digital safeguards.

What could have changed the outcome

This is, notably, the case study in this set where less needed to change to produce a good outcome, because the operational context - layered physical and procedural safeguards around a known-critical process - was already built in by engineering discipline that predates modern OT cybersecurity entirely. What could have made the outcome more reliable, rather than dependent on one operator happening to be watching a screen, is automated monitoring tied to the same Criticality context that already justified the physical safeguards - turning a lucky catch into a designed one.

The lesson

Operational context does not have to arrive through a cybersecurity tool to work. Engineering disciplines like defence-in-depth process safety already encode Criticality and Consequence context, often more rigorously than most cybersecurity programmes do. The task for OT security is not to invent this thinking from scratch, but to connect to it - treating decades of process-safety engineering as a primary source of operational context, not a separate discipline to be reinvented.

Conclusion

Every incident examined here had a technical signal somewhere in it - a leaked credential, a malfunctioning piece of malware, months of dwell time, a moving mouse cursor. In not one case was the outcome determined by whether the signal existed. It was determined by whether the organisation around it had the context to interpret it correctly - and, at Oldsmar, by whether decades-old engineering discipline had already built consequence-aware safeguards into the process itself.

The OT cybersecurity industry has invested enormously, and successfully, in seeing more. But visibility was never going to be the finish line, because it answers what exists, not what matters. In June 2026, the U.S. federal government's own vulnerability-management policy moved toward the questions operational context is built to answer: is this reachable, is this being exploited, and what would actually happen if it were.

> *"The industry does not have a data problem anymore. It has a meaning problem. And meaning cannot be scanned for - it has to be built."*

Velos Shield's contribution is the Operational Context Stack: a structured, teachable way to build the six layers of understanding that separate a technical finding from an operational decision, on a maturity path toward Cyber-Aware Operations. Organisations that internalise this distinction now - ahead of the regulatory and threat pressure this paper documents - will be the ones able to say, with confidence, what a given vulnerability actually means to their operation. That question is the territory worth owning for the next decade of industrial cybersecurity.

A position, not a product claim

This paper describes our perspective on how OT monitoring can evolve. It intentionally doesn't present every idea here as an already-shipped capability, and it doesn't claim this is the only valid architecture.

LET'S TALK

How you think about OT Security, we'd like to hear it. Reach out to us in contact@Velos Shield.com to continue the conversation.

About Velos Shield

Velos Shield builds OT monitoring designed around trustworthy evidence, operational context and persistent memory - so security and engineering teams spend less time re-discovering what they already know, and more time acting on what matters.