ARTICLE



IM_

Built to bend

Resilience starts before the control system goes live

For many organisations, operational technology (OT) has traditionally been managed through two separate lenses. Engineering has focused on reliability and safety, while cybersecurity teams have focused on protecting systems from compromise. Both remain essential, but neither, on its own, tells us whether a critical operation is genuinely resilient.

A water utility can have highly reliable equipment and still be vulnerable to cyber manipulation. An industrial site can have mature security controls and still struggle to maintain essential services when equipment fails or to restore a critical system within the time the operation requires. And a well-protected control system may depend on identity services, telecommunications, cloud platforms, suppliers, or specialist engineers far outside the traditional OT boundary.

Whether the critical operation can continue safely when something goes wrong, and whether it can be restored within an acceptable time, are the questions that matter most now. A reliable and secure OT environment is where the answers begin to materialise.

Robust, secure, and recoverable

Critical operations can be disrupted by anything from a worn component to a hostile state. Resilience requires the ability to withstand and manage disruption anywhere along that range. Three capabilities need to come together.

The operation needs to be *robust*: able to tolerate faults and disruptions while keeping an acceptable level of service, whether through redundancy or reduced functionality. It needs to be *secure*: protected against cyber compromise, and able to limit the operational consequences of a breach. And it needs to be *recoverable*: able to be restored to safe running within a time the business can accept, whatever the cause of the disruption.

To this end, reliable equipment and systems remain fundamental. Robustness adds the ability to sustain essential operation when individual components or dependencies fail. Cybersecurity contributes to that robustness and deserves explicit attention because connected systems can be deliberately manipulated or compromised.

Safety is the boundary around all three. In critical infrastructure, resilience means continuing to operate within safe limits for people and the environment, or stopping safely when that is not possible. Bringing the three capabilities together may sound straightforward. In practice, however, they are often managed by different functions, with different methods, priorities, suppliers, and accountabilities. A disruption crosses all of those lines of responsibility at once, which makes that split increasingly hard to defend.

Business impact analysis starts from the operation itself and asks what the business most needs to keep running. It identifies which operations are critical, what their disruption costs over time, and the minimum service and recovery they must meet. That gives cyber risk assessment its starting point: which compromise scenarios could disrupt those operations, how well existing controls hold, and where a proportionate treatment is worth the effort. The link also runs the other way. A cyber scenario can expose a dependency that the business impact analysis missed, and that finding should go straight back into the analysis.

Beyond the fence line

What does it take to keep a critical operation running? Once, the answer sat inside the plant. Today, the operation runs through systems and services that may belong to someone else, and that dependency, more than connectivity itself, is the true source of the new risk.

A critical operation may now rely on enterprise IT and identity systems, telecommunications, cloud services, remote access by original equipment manufacturers (OEMs), software suppliers, spare parts, and external specialists, many of them outside the direct control of the people responsible for operating the asset. For that reason, it is no longer enough to simply map the OT environment. Organisations also need to understand what the operation depends on, what happens when those dependencies fail or can no longer be trusted, and what alternatives exist.

Two incidents from 2024 spell this out. In the first, attackers manipulated pump pressure at a Danish waterworks, disrupting the water supply and physically damaging pipes, and Danish intelligence later linked the group responsible to the Russian state.

The second came from a trusted supplier: a faulty update to CrowdStrike's security software disrupted organisations running critical services around the world. While these two examples had different causes, they raise the same management question, one that is fundamentally a resilience question more than strictly a cybersecurity one: what can we still operate when a critical dependency fails?

Day zero

One aspect of resilience deserves more attention: by the time an OT system goes live, many of the choices that determine its future resilience have already been made. Architecture, supplier selection, contractual responsibilities, remote access, support arrangements, integration, testing, and recovery requirements all shape what the organisation will later be able to protect, maintain, isolate, and restore. But once the system is operational, retrofitting some of those decisions can prove very difficult and expensive.

So resilience needs to enter the picture much earlier. When defining requirements for a new or modernised critical system, organisations should not only specify what the system needs to do, but also define what needs to happen when parts of it fail. What is the minimum acceptable level of service? Which dependencies must the operation be able to run without? Can it run locally or manually? Can parts of the environment be isolated without losing the entire operation? What is needed to restore the system safely?

Those requirements then need to survive the whole path from design to handover. A resilience requirement that disappears somewhere between the design document and the supplier contract is unlikely to help much when the operation is disrupted.

One OT lifecycle



Define the need and manage risk

- Business case
- Critical services and impact
- Risk assessment
- OT strategy and requirements

Turn objectives into obligations

- Technical requirements
- Security and resilience requirements
- Risk allocation
- Supplier obligations

Engineer secure, reliable solutions

- Architecture
- Engineering
- Security and resilience by design
- Design assurance

Build, integrate, and test

- Supplier management
- Integration
- FAT and SAT
- Commissioning

Ensure readiness for operation

- Acceptance
- Documentation
- Residual risk
- Operational and recovery readiness

Maintain and continuously improve

- O&M and reliability
- Vulnerability management
- Monitoring and response
- Recovery, change, and lifecycle

Four positions in the value chain

Product vendor
Build products that can be securely integrated and operated throughout their lifecycle.

EPC / contractor
Deliver compliant and reliable integrated systems on time, on budget, and in line with requirements.

Developer
Translate risk and operational needs into executable project requirements and commercial frameworks.

Asset owner and operator
Operate critical assets safely, reliably, and securely throughout their lifecycle.

Figure: Supporting the full OT lifecycle

Critical infrastructure is rarely designed, built, and operated by one organisation. Asset owners work with a chain of EPC (engineering, procurement, and construction) contractors, system integrators, OEMs, technology providers, and service partners. Responsibilities move between these parties throughout the lifecycle, making the interfaces between them particularly important.

A supplier may assume that the asset owner has defined the recovery requirement, while the owner may assume it is inherent in the technical specification. An integrator may test whether the system functions correctly without testing how it behaves during disruption. A remote-access arrangement that was useful during commissioning may quietly become a permanent part of the operating model. Each can look unremarkable on its own. Together, they determine what options the organisation has when something goes wrong.

Standards such as IEC 62443 treat cybersecurity as a shared lifecycle responsibility between the actors in the OT ecosystem. European regulation is moving the same way, with the NIS2 Directive, the Critical Entities Resilience Directive, and the Cyber Resilience Act all increasing the focus on continuity, supply chains, and recovery. Compliance can materially improve resilience when implemented well, but resilience goes further: it asks whether the critical operation can withstand disruption and recover.

Case example

Great technology is not enough

An asset owner investing in critical energy infrastructure based its business case on sustained production and predictable performance over the asset's lifetime. The EPC contractor and the supply organisation procured against engineering specifications, and a component vendor delivered technology that performed as required at commissioning. Yet the service, maintenance, security, and recovery arrangements needed to sustain that performance had not been sufficiently established.

The gap lay in the questions that went unasked across the value chain. The supply organisation needed to ask O&M what it would take to operate, maintain, secure, and restore the solution. O&M, in turn, needed to ask the asset owner which production objectives, acceptable interruptions, and lifetime costs the operating model had to meet. Component vendors needed to clarify the conditions and responsibilities required to sustain their technology in that operating model.

Those answers should have shaped the specifications, procurement packages, support agreements, and acceptance criteria. Without them, each party could fulfil its immediate scope while leaving essential conditions for long-term value unresolved.

The lesson is that **each actor needs to look beyond its own delivery and ask what the next part of the value chain needs to succeed.** Understanding the asset owner's business objectives is essential to designing the resilience needed to sustain them. That understanding must guide how the solution is engineered, procured, operated, secured, and restored, so that technology delivered at commissioning becomes value sustained in operation.

Recovery is an operational capability

Having backups is important, but a backup does not tell you whether a production line, water supply, energy system, or other physical operation can actually be restored.

Operational recovery raises different questions. Which physical process should be restored first? Is there a safe sequence for bringing the operation back online? Are the necessary configurations, controller logic, engineering workstations, licences, spare hardware, and specialist skills available? Will external suppliers be there when several customers need them at the same time?

The most important question is whether the organisation has demonstrated that it can recover within the required time, and leaders should keep in mind that a recovery objective on paper and a recovery capability tested under realistic conditions are two very different things.

Five questions for leaders

The implication is not that boards and executive teams should become OT security specialists. Rather, resilience should be discussed in terms of the operations for which they are already accountable.

Five questions are a useful place to start:

1. Which operations are genuinely critical, and what is the minimum acceptable service they need to provide?

2. What does each critical operation actually depend on, inside and outside the organisation?

3. If one of those dependencies fails or is compromised, can the operation fail safely or continue in a degraded mode?

4. What needs to be designed in, or procured, to give us the robustness and security we need – and how does this translate into verifiable supply requirements?

5. Who is accountable for resilience across internal functions, suppliers, and the lifecycle of the asset?

These questions cut across engineering, operations, cybersecurity, procurement, and leadership, which is precisely the point, because resilience cannot sit entirely within any single one of them.

Built before it is needed

No organisation can prevent every disruption, and resilience does not mean keeping every system running at all costs. Sometimes the resilient response is to isolate a system, operate manually or at reduced capacity, or shut an operation down safely and know how to bring it back. What matters is that these should be deliberately built capabilities rather than decisions invented for the first time during an incident.

For critical operations, resilience starts well before something goes wrong. It begins when requirements are defined and systems are designed, and it continues through procurement, contracts, engineering, commissioning, operations, maintenance, and change. The best time to build resilience into that lifecycle is before the control system goes live. For systems already running, the second-best time is now.



Contact

For more information, please contact:

Michael Vincentz Lundgren
Implement Consulting Group
+45 5221 5382
mvlu@implement.dk

Philip Özari
Implement Consulting Group
+45 5221 6262
phoz@implement.dk