



# ICIT

## Beyond Compliance: Engineering Cyber Resilience for Critical Infrastructure

Author: Dr. Ed Harris

# Table of Contents

|                                                                 |     |
|-----------------------------------------------------------------|-----|
| I. When did compliance last stop a real attack?                 | 3   |
| II. The Compliance Gap: Necessary but Not Sufficient            | 4   |
| III. The Resilience Imperative by the Numbers                   | 5   |
| IV. Modeling Cyber-Physical Risk: Where Bits Meet Steel         | 5-6 |
| V. Prioritizing by Business Impact, Not Asset Class             | 6-7 |
| VI. The Supply Chain as the New Perimeter                       | 7   |
| VII. Leading Through Crisis: The First Four Hours               | 7-8 |
| VIII. The Ecosystem of Resilience: Public-Private Collaboration | 8   |
| VI. Stakeholders and Their Role in the Ecosystem                | 8   |
| IX. Designing Systems That Fail Gracefully                      | 9   |
| XI. Closing                                                     | 10  |



## About ICIT

The Institute for Critical Infrastructure Technology (ICIT) is a nonprofit, nonpartisan, 501(c)3 think tank with the mission of modernizing, securing, and making resilient critical infrastructure that provides for people's foundational needs. ICIT takes no institutional positions on policy matters. Rather than advocate, ICIT is dedicated to being a resource for the organizations and communities that share our mission. The views and opinions expressed in this essay are solely those of the author(s) and do not necessarily reflect the official policy or position of ICIT. Any assumptions made within the analysis are not reflective of the position of any entity other than the author(s). To learn more, please visit [www.icitech.org](http://www.icitech.org)

## The Institute for Critical Infrastructure Technology

# I. When did compliance last stop a real attack?

That question looms over every security leader and operational engineer responsible for critical infrastructure. It is an uncomfortable question to open with. Yet, it is the right one for anyone who runs security on the systems that keep power flowing, water clean, and pipelines pressurized. For years, the critical infrastructure sector has measured itself against audit frameworks—[NERC CIP](#), [IEC 62443](#), [NIST CSF](#)—and has treated a passed assessment as evidence of safety. The data tells a different story. Point-in-time audits, checkbox attestations, and prevention-first mindsets create accountability, but not immunity. In complex operational technology (OT) and industrial control system (ICS) environments, compromise is not a possibility organizations can eliminate; it is an inevitability they must absorb, adapt to, and recover from rapidly. The rapid convergence of information technology (IT) and operational technology (OT) networks—driven by remote access demands, predictive analytics, and enterprise data integration—has rendered traditional perimeter defense insufficient.

For decades, industrial operators treated cybersecurity as a structural barrier: build sufficiently high perimeter walls, pass regulatory audits, and physical operations will remain undisturbed. The stark reality facing modern operators is that regulatory compliance no longer guarantees operational continuity. Organizations regularly pass comprehensive audits that certify full compliance with recognized control standards, only to experience catastrophic downtime weeks later. The underlying flaw is conceptual: compliance measures static controls at a fixed point in time, whereas industrial threats operate in dynamic, persistent, and adaptive cycles.

To close this gap, security leaders and engineering executives must shift from a standard prevention mindset to an *active engineering-for-continuity framework*. This strategic evolution acknowledges a foundational operational truth: organizations must assume the adversary has breached the environment and measure resilience by how seamlessly a system contains physical impacts and maintains core operations under stress.

This article reframes cybersecurity for critical infrastructure operators around a single imperative: resilience. Drawing on practitioner insights, recent incident data, and established frameworks, we outline how organizations can move beyond the compliance floors that standards such as [NERC CIP](#), [IEC 62443](#), and the [NIST Cybersecurity Framework](#) establish, and how to build systems and cultures that keep essential services running.

## II. The Compliance Gap: Necessary but Not Sufficient

Compliance remains foundational. It establishes minimum controls, creates audit trails, and satisfies regulatory obligations. Yet the data reveals a persistent and dangerous gap between passing an audit and achieving operational security. Regulatory frameworks such as NERC CIP, IEC 62443, and the NIST Cybersecurity Framework (CSF) establish necessary baseline controls. However, treating these frameworks as an ultimate security target creates a dangerous false sense of safety. Compliance is binary and checklist-driven; resilience is continuous and operational.

Industry data highlights this gap. According to the [Dragos 2024 ICS/OT Report](#), 43 percent of critical infrastructure organizations experienced a significant breach within 12 months of passing an audit. Average OT incident dwell time stands at 197 days, meaning attackers often persist undetected for more than six months ([IBM X-Force Incident Response 2024](#)). These figures are not anomalies; they are structural. During this period, attackers often leverage legitimate credentials and standard administrative tools, bypassing traditional boundary detection mechanisms completely.

The compliance mindset prioritizes:

- Point-in-time assessments
- Checkbox attestation
- Audit-driven investment cycles
- Prevention as the primary goal

*Success is measured by controls passed.*

A resilience mindset flips the model:

- Continuous threat posture awareness
- Operational stress-testing
- Risk-driven investment
- Recovery as a core design principle

*Success is measured by time to recover and impact radius.*

[Colonial Pipeline](#) illustrated the distinction starkly. The company operated under applicable [TSA Pipeline Security Guidelines](#), yet a ransomware attack forced a six-day shutdown of fuel delivery across the Eastern United States. Compliance created accountability, but resilience would have determined how quickly operations could resume and how far the impact spread. The practical question for any security leader isn't "would we pass an audit tonight?" It's "what does our recovery time look like if OT is compromised tonight?" That reframing—from prevention to resilience—is the organizing idea of everything that follows.

# III. The Resilience Imperative by the Numbers

Recent data underscores why the shift matters:

- 74 percent of ICS/OT operators experienced at least one cyberattack that impacted production in 2024 ([Claroty 2024](#)).
- Average recovery time for OT incidents without pre-planned resilience procedures reaches 22 days ([SANS ICS 2024](#)).
- 58 percent of critical infrastructure breaches originated in the supply chain or third-party ecosystem ([CrowdStrike 2025](#)).
- The average cost of a critical infrastructure data breach is \$8.4 million—20 percent higher than the cross-industry average ([IBM Cost of a Data Breach Report 2024](#)).

Resilience isn't about eliminating risk. It's about ensuring systems keep operating, or recover rapidly, when risk materializes—through three capacities: the ability to absorb a shock without cascading failure, adapt operations under degraded conditions, and recover to normal function on a known timeline.

# IV. Modeling Cyber-Physical Risk: Where Bits Meet Steel

Bridging the operational divide requires mapping digital vectors directly to physical outcomes. Leveraging the MITRE ATT&CK for ICS framework alongside NIST CSF 2.0 Governance guidelines, organizations must model how cyber intrusions manifest as operational downtime, equipment destruction, or environmental hazards.

When ransomware compromises business network domain controllers, secondary dependencies—such as automated engineering workstations or historian data feeds—often trigger safety shutdowns even if Programmable Logic Controllers (PLCs) remain uncompromised.

Effective resilience begins with rigorous cyber-physical risk modeling and a comprehensive dependency mapping across IT and OT boundaries. Digital events do not remain digital in OT environments. They translate into physical consequences: pressure loss, cascading grid failures, water contamination, or pipeline over-pressurization.

A practical five-step workflow provides structure:

1. **Asset Inventory and Dependency Mapping** — Maintain a living OT/IT asset registry that captures every cyber-physical touchpoint and upstream dependency.
2. **Threat Scenario Development** — Define adversary tactics, techniques, and procedures relevant to the sector (examples include [TRITON/TRISIS](#), [Volt Typhoon](#), and ransomware lateral movement into control networks).
3. **Consequence Modeling** — Explicitly map digital compromises to physical outcomes. This step requires participation from safety engineers and operators, not solely the security operations center.
4. **Control Effectiveness Assessment** — Evaluate existing controls through operational simulation rather than paper compliance.
5. **Risk Prioritization Matrix** — Score scenarios by probability × consequence × detectability, directing investment first toward the highest physical-impact risks.

Recommended foundations for this work include the [MITRE ATT&CK for ICS](#) matrix and the Govern function of [NIST Cybersecurity Framework](#) 2.0. Consequence modeling also directly supports emerging regulatory obligations, including material cyber incident disclosure requirements under [SEC rules](#).

## V. Prioritizing by Business Impact, Not Asset Class

Not every system carries equal weight. A business impact analysis (BIA) built around asset categories tends to under-protect the systems that matter most. A better model prioritizes by consequence: which systems gate revenue and regulatory compliance, with downtime cost calculated in advance; which govern life safety — SCADA and safety-instrumented systems chief among them; and which carry mandatory reporting obligations under CIRCIA, NIS2, or SEC disclosure rules, where a compliance failure compounds an operational one.

Every critical system needs a documented recovery time objective and recovery point objective. If that documentation doesn't exist, it's the first gap to close—and interdependency analysis needs to trace cascading failure paths not just across the IT-OT boundary but across organizational boundaries, into utilities, municipalities, and vendors. A simple tiering model helps allocate resources in sequence: Tier 1 for life safety, Tier 2 for mission-critical operations, Tier 3 for supporting systems. A Business Impact Analysis (BIA) is not a one-time exercise—it needs review after every major infrastructure change, acquisition, or shift in the threat landscape.

A BIA must drive prioritization:

- Identify OT systems that directly gate revenue or regulatory compliance and pre-calculate downtime cost per hour.
- Segregate life-safety systems—SCADA and [safety instrumented systems](#) (SIS)—for highest protection.
- Map systems to mandatory reporting regimes such as [CIRCIA](#), NIS2, and [SEC rules](#), recognizing that compliance failure compounds operational loss.
- Establish documented recovery time objectives (RTO) and recovery point objectives (RPO) for every critical system. Absence of these metrics is itself a critical gap.
- Analyze interdependencies across IT-OT boundaries and organizational boundaries (utilities, municipalities, vendors).
- Apply a tiered investment model: Tier 1 for life safety, Tier 2 for mission-critical operations, and Tier 3 for supporting systems.

BIA is not a one-time exercise. It requires review after every major infrastructure change, acquisition, or shift in the threat landscape. Misalignment between system owners and leadership on RTO knowledge is a material risk in itself.

## VI. The Supply Chain as the New Perimeter

Modern industrial facilities do not operate in isolation. Supply chain and vendor ecosystem access represent the primary entry point for major critical infrastructure compromises, accounting for 58% of recorded breaches. Supply chain compromises have repeatedly demonstrated that third-party access equals first-party risk.

- SolarWinds is a nation-state compromise of a software build pipeline that affected approximately 18,000 organizations, including federal agencies and energy operators, with a nine-month dwell time.
- Kaseya VSA enabled ransomware deployment to roughly 1,500 downstream organizations within hours via managed service provider access.
- The XZ Utils incident showed that open-source dependencies themselves can become vectors when a malicious contributor nearly inserted a backdoor into critical authentication components.

CISA Binding Operational Directive 23-01 requires federal agencies to enumerate internet-exposed assets and continuously manage third-party vendor exposure. The practical starting point is straightforward, as Critical infrastructure operators should adopt similar discipline: implement software bill of materials (SBOM) requirements for OT vendors, demand SOC 2 Type II or equivalent assurance, conduct quarterly risk reviews of Tier 1 suppliers, and treat any vendor with privileged OT access as a potential initial access vector. The working assumption CISA recommends is that any vendor with privileged access to an OT environment is a potential initial access vector and should be segmented, monitored, and limited accordingly. The harder question is the one that tends to quiet a room of practitioners: how many vendors have unmonitored remote access to your OT environment right now?

## VII. Leading Through Crisis: The First Four Hours

When an incident occurs, executives face compressed decision windows. Pre-authorization is the single greatest accelerator of effective response. Decisions made under pressure without pre-approved frameworks take three to seven times longer ( SANS ICS 2024 ).

Five critical decisions dominate the first four hours:

1. Isolate — Which systems shut down and which remain running? Pre-authorized isolation authorities prevent paralysis.
2. Communicate — Who is notified first—board, regulators (CISA and sector-specific agencies), or critical customers? The Cyber Incident Reporting for Critical Infrastructure Act (CIRCI) imposes a 72-hour reporting for covered entities and 24-hour reporting for ransomware payments.
3. Operate — Can the organization shift to manual or degraded-mode operations? Has the workforce trained and recently tested analog fallback procedures?
4. Preserve — Activate legal hold immediately. Evidence preservation and forensics must run concurrently with recovery, not afterward.
5. Recover — Execute a tested, board-approved incident response plan written with OT engineers, not a document last updated years earlier.

Boards ultimately ask two questions after every incident: Were you prepared? Did you communicate clearly? Organizations that can answer yes to both typically survive the crisis with less damage.

## VIII. The Ecosystem of Resilience: Public-Private Collaboration

Critical infrastructure operators consistently underuse the free federal and sector resources built specifically for them. No operator faces these challenges alone.

Underutilized free resources exist:

- CISA's Joint Cyber Defense Collaborative (JCDC) provides real-time threat sharing.
- Sector Information Sharing and Analysis Centers (ISACs) and Information Sharing and Analysis Organizations (ISAOs)—including E-ISAC, WaterISAC, FS-ISAC, and H-ISAC—deliver sector-specific intelligence.
- CISA offers free CISA Cybersecurity Performance Goals assessments, resilience planning support, and on-site exercises.
- Sector-specific agencies (Department of Energy, Environmental Protection Agency, Department of Homeland Security, Department of Transportation, and others) maintain specialized resources.

Pre-incident relationship building with CISA, the FBI field office cyber coordinator, and the relevant sector agency accelerates response by days. Cyber Incident Reporting for Critical Infrastructure Act (CIRCI) obligations should be understood thoroughly before an incident forces rapid compliance.

# IX. Designing Systems That Fail Gracefully

Resilience architecture rests on four progressive stages:

1. **Redundancy** — Eliminate single points of failure through hot-standby controllers, diverse communication paths (fiber, cellular, satellite), and offline historian backups isolated from the IT network.
2. **Fail-Safe Design** — Define the safe state for every critical system and verify it holds under digital attack. Safety instrumented systems must operate independently of basic process control systems; that independence requires testing.
3. **Recovery Planning** — Document and regularly test restore procedures from known-good offline images. Tier 1 systems should target an RTO of four hours or less. Validate through both tabletop and live drills.
4. **Culture and Training** — Resilience is ultimately an organizational effort. Operators who understand threat scenarios make better real-time decisions. Quarterly OT-specific drills and career-development emphasis on resilience close the gap between technology and execution.

Organizations that conduct OT-specific tabletop exercises reduce mean time to contain by 47 percent ([Dragos 2024 ICS/OT Year in Review](#)). Yet only 34 percent of critical infrastructure operators test OT recovery procedures annually ([SANS ICS Survey 2024](#)).

## Five Actions for This Week

Resilience begins with deliberate decisions followed by execution. Consider these immediate steps:

| 01                                                                                                                                | 02                                                                                                                                                                                            | 03                                                                                                                                                           | 04                                                                                                                         | 05                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Audit the business impact analysis.</b><br>Confirm every Tier 1 OT system has a documented RTO, RPO, and named recovery owner. | <b>Map vendor access.</b><br>Enumerate every third party with remote access to the OT environment. Remove unnecessary access, segment the remainder, and enforce multi-factor authentication. | <b>Build the speed dial.</b><br>Add CISA (1-888-282-0870), the sector ISAC, and the local FBI cyber coordinator to the incident response contact list today. | <b>Schedule a tabletop.</b><br>Conduct an OT-specific exercise within 90 days that includes both executives and operators. | <b>Train for analog fallback.</b><br>Identify at least three critical processes that can operate manually and verify that the required skills still exist on the floor. |

## XI. Closing

Resilience is not a project with an end date. It is a posture—one built through continuous risk modeling, impact-driven prioritization, supply-chain discipline, pre-authorized crisis decision frameworks, active public-private collaboration, graceful failure design, and a culture that treats recovery as a core competency.

The floor is set by compliance. The ceiling is defined by how quickly and safely critical systems absorb, adapt, and recover when compromise occurs. That ceiling is built one deliberate decision at a time. Adversaries plan around recovery time. Organizations that plan faster retain control.

*Assume compromise. Engineer for continuity.*

# Author



Dr. Ed Harris  
Chief Information Security  
Officer, Mauser Packaging  
Solutions

---

Dr. Harris, has a Doctorate in Information Assurance and Cybersecurity and 35+ years of network and security experience. Dr. Harris is a Certified Information Systems Security Professional and has spent 12 years in academia teaching Computer Security at the University of North Carolina. Dr. Harris is a founding member of the Retail Cyber Information Sharing Center. He currently serves on the Cyber Security Scholarship Review Board for ISC2 and is an active member of the FBI InfraGard. He is the former President of the ISC2 Chapter in Greensboro, NC, and a former member of the National Security Council at the US Chamber of Commerce. Dr. Harris is a security evangelist, author, and international speaker. He has spoken both academically and professionally for more than 32+ years.



# ICIT

[www.icitech.org](http://www.icitech.org)