

The Cybersecurity Fear Machine

The Cybersecurity Fear Machine

How Inflated Threats Hijacked America's Critical Infrastructure Policy

Daniel Ward, Ph.D.



Fountain & Key Press · Cambridge, MA

The Cybersecurity Fear Machine: How Inflated Threats Hijacked America's Critical Infrastructure Policy

Copyright © 2026 Daniel Ward, Ph.D.

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

ISBN (Paperback): 978-1-972842-01-0

ISBN (Ebook): 9781972842003

First Edition

Published by Fountain & Key Press

Cambridge, Massachusetts

fountainandkeypress.com

Artificial intelligence tools may have been used to support the writing, editing, research, and/or production of this work. All content has been reviewed and approved by the author, who is responsible for the final text.

Any third-party websites referenced in this work are not under the control of the publisher, which assumes no responsibility for their content or availability.

Printed in the United States of America

CONTENTS

1	THE UNEXAMINED NARRATIVE: OPERATIONAL TECHNOLOGY UNDER SIEGE?	I
2	THE HISTORICAL RECORD: SCARCITY OF CATASTROPHE	7
3	THE MAZE OF U.S. GOVERNANCE: OVERLAPPING MANDATES AND PARALYSIS	31
4	THE VENDOR ECOSYSTEM: PROFITING FROM FEAR	49
5	POLICY FRAMEWORKS BUILT ON MYTH: A CRITICAL EXAMINATION	75
6	THE MYTH OF IMMINENT CATASTROPHE: A STATISTICAL REALITY CHECK	103
7	INTERNATIONAL PERSPECTIVES: DIVERSE APPROACHES TO OT SECURITY	131
8	THE ILLUSION OF CONTROL: SYSTEMIC VULNERABILITIES UNMASKED	155
9	TOWARDS A FEDERATED MODEL: ENHANCING SITUATIONAL AWARENESS	177
10	IMPROVING INCIDENT COORDINATION AND RESPONSE	203
11	REDUCING SYSTEMIC VULNERABILITIES THROUGH GOVERNANCE REFORM	229

12	THE DELICATE BALANCE: STATE AUTHORITY, INDUSTRY, AND CIVIL OVERSIGHT	249
13	BEYOND FEAR: SHIFTING FROM PERCEIVED THREATS TO SYSTEMIC RESILIENCE	261
14	CONCLUSION: SECURING AMERICA’S FUTURE IN A COMPLEX WORLD	297
	Glossary of Terms	313
	Selected Bibliography and Further Reading	315
	Notes	317
	THE UNEXAMINED NARRATIVE: OPERATIONAL TECHNOLOGY UNDER SIEGE?	317
	THE HISTORICAL RECORD: SCARCITY OF CATASTROPHE . . .	317
	THE MAZE OF U.S. GOVERNANCE: OVERLAPPING MANDATES AND PARALYSIS	324
	THE VENDOR ECOSYSTEM: PROFITING FROM FEAR	325
	POLICY FRAMEWORKS BUILT ON MYTH: A CRITICAL EXAMINATION	330
	THE MYTH OF IMMINENT CATASTROPHE: A STATISTICAL REALITY CHECK	331
	INTERNATIONAL PERSPECTIVES: DIVERSE APPROACHES TO OT SECURITY	333
	THE ILLUSION OF CONTROL: SYSTEMIC VULNERABILITIES UNMASKED	334
	REDUCING SYSTEMIC VULNERABILITIES THROUGH GOVERNANCE REFORM	335
	BEYOND FEAR: SHIFTING FROM PERCEIVED THREATS TO SYSTEMIC RESILIENCE	335
	CONCLUSION: SECURING AMERICA’S FUTURE IN A COMPLEX WORLD	336
	About the Author	339

CHAPTER I

THE UNEXAMINED NARRATIVE: OPERATIONAL TECHNOLOGY UNDER SIEGE?

On May 7, 2021, an employee at Colonial Pipeline discovered a ransom note on a company computer. The company, which by its own account supplies roughly 45 percent of the fuel consumed on the U.S. East Coast, was hit by a ransomware attack. The attackers had not breached the operational technology (OT) controlling the pipeline itself, but had compromised the company's information technology (IT) billing systems. Fearing the infection could spread and unable to bill customers, Colonial's leadership made the extraordinary decision to shut down the entire 5,500-mile pipeline — the first full system shutdown in the company's history attributable to a cyber incident.

The effects were immediate. Panic buying led to widespread gasoline shortages, flight cancellations, and a spike in fuel prices. The roughly five- to six-day shutdown was not caused by a sophisticated cyber-physical attack, but by an IT breach that prompted an operational shutdown out of caution. The incident illustrates the book's central theme: while the discourse on OT security focuses on catastrophic physical destruction,

the most significant impacts often stem from failures in governance, resilience, and the blurred lines between IT and OT.

A potent narrative dominates cybersecurity policy: America's critical infrastructure is on the precipice of catastrophic cyber-physical attack. This book argues that the discourse is misaligned with the evidence. It overemphasizes catastrophic attacks while underappreciating more frequent threats like governance failures and poor resilience, directing resources toward improbable scenarios instead of foundational vulnerabilities.

To deconstruct this narrative, we must first define "Operational Technology" (OT). Unlike Information Technology (IT), which focuses on data processing, communication, and business operations, OT systems are designed to monitor and control physical processes. Think of the Supervisory Control and Data Acquisition (SCADA) systems that manage power grids, the Distributed Control Systems (DCS) that operate chemical plants, or the Programmable Logic Controllers (PLCs) that automate manufacturing lines. These are the digital interfaces to the physical world, where failures have direct, tangible consequences. A compromised power grid can lead to widespread blackouts, a disrupted water treatment facility can endanger public health, and a halted manufacturing plant can disrupt supply chains. The security of OT is a distinct challenge due to its real-time operational needs, long system lifecycles, and the direct physical impact of failure.

The Genesis of the Narrative

The narrative of imminent cyber-physical catastrophe did not emerge in a vacuum; it evolved from technological advancements, geopolitical events, media portrayals, and the rise of a powerful cybersecurity industry. Early, often sensationalized, cyber incidents etched a sense of vulnerability into the public consciousness. The initial reporting of computer viruses and network intrusions, while often lacking detail about the actual impact on physical systems, planted the seed of digital danger. As the internet became more pervasive and interconnected, the potential attack surface expanded, and so did the perceived threat. Geopolitical tensions further amplified this narrative, with accusations of state-sponsored cyber espionage and sabotage becoming commonplace, even when attribution was uncertain or politically motivated. This historical trajectory, which we will explore in more detail in subsequent sections, laid the fertile ground for the "fear machine" to take root.

What is striking about this narrative is its age. The prediction of civilizational collapse by keystroke is not a product of the ransomware era; it is older than the commercial internet itself. The specter of an “Electronic Pearl Harbor” was invoked as early as 1991, long before most Americans had ever sent an email.¹ Two decades on, the warning had migrated from think-tank papers to the highest levels of government. On October 11, 2012, in a speech to Business Executives for National Security aboard the USS Intrepid, Secretary of Defense Leon Panetta warned that a ‘cyber Pearl Harbor’ could derail passenger trains, or even more dangerous, derail trains loaded with lethal chemicals, and could ‘contaminate the water supply in major cities, or shut down the power grid across large parts of the country.’² These were not abstractions; they were specific, vivid, catastrophic scenarios attached to a deadline that was always imminent. And yet the central, uncomfortable fact this book sets out to examine is that more than three decades after the phrase was coined, and well over a decade after Panetta’s warning, there has been no cyber-induced derailment of a chemical train, no cyber-caused mass poisoning of a city’s water supply, and no U.S. grid felled by an adversary’s keystroke at civilizational scale — though intrusions or suspected intrusions into water utilities (the disputed Oldsmar, Florida incident in 2021; the CyberAv3ngers PLC compromise at Aliquippa, Pennsylvania in 2023) show the surface is real even where consequences have been contained, and Russian-attributed cyberattacks caused localized grid outages in Ukraine on December 23, 2015 (affecting regional distribution utilities) and on the night of December 17–18, 2016 (the Pivnichna transmission substation north of Kyiv), demonstrating that limited, targeted disruptions are possible even if the catastrophic civilizational version has not materialized. We will return to this gap between prophecy and record again and again, because it is the thread from which the entire fear machine is woven.

Understanding the “fear machine” requires us to identify its architects and their motivations. It is not a monolithic entity but an ecosystem of actors who amplify perceived threats. Several groups contribute to this ecosystem:

- **Cybersecurity Vendors:** Have a vested interest in highlighting threats to sell products and services.
- **Consultants:** Often emphasize worst-case scenarios to justify risk management engagements.
- **Think Tanks:** Can shape policy discourse with reports, sometimes funded by industry or government, that create demand for specific security measures.

- **Government Agencies:** May use perceived threats to justify budgets, expand mandates, and demonstrate their role in national security. Recognizing these diverse motivations is crucial for dissecting the “fear machine” and understanding how commercial interests and bureaucratic imperatives can shape public perception and policy decisions, potentially diverting resources from more pressing, fundamental issues.

The scale of the commercial interest at stake is not incidental to this story; it is the engine of it. Industry analysts project the global cybersecurity market will reach roughly \$322 billion by 2029.³ A market of that magnitude does not merely respond to fear; it has every reason to cultivate it. We will hold that number in mind throughout what follows, because the question of who profits from the narrative of perpetual siege is one no honest accounting of OT security can avoid.

Several characteristics distinguish OT from IT and contribute to its security challenges. Firstly, OT systems often have extraordinarily long lifecycles, frequently measured in decades, far exceeding the three- to five-year refresh cycles typical for much enterprise IT hardware. This longevity means many OT systems are built on outdated technologies, operating systems, and communication protocols that were designed with little or no consideration for cybersecurity. Patching and upgrading these systems are often complex, costly, and disruptive, requiring extensive downtime or specialized engineering expertise that may not be readily available. Secondly, OT systems are characterized by stringent real-time operational requirements. Unlike IT systems where a slight delay in data processing might be acceptable, in OT environments, even milliseconds of latency can have severe consequences, potentially leading to equipment damage or unsafe operating conditions.

This makes traditional IT security measures, such as frequent reboots or intrusive scanning, impractical or even dangerous. Finally, the direct link to physical processes means that cybersecurity failures in OT can result in immediate, tangible, and potentially catastrophic physical damage, injury, or environmental harm. This differs from most IT breaches, which typically involve data theft or business disruption. It is this direct physical impact that fuels much of the fear surrounding OT cybersecurity, and it is precisely this aspect that has been heavily leveraged in the construction of the “fear machine.”

Government agencies also participate. Justifying large budgets and expanded mandates often relies on demonstrating a significant threat. Cybersecurity has become a

vital component of national security, and highlighting potential vulnerabilities allows agencies to secure resources and assert their importance in protecting the nation. This can create a feedback loop: perceived threats are amplified to secure funding, which fuels demand for security products and reinforces the narrative. Understanding these motivations, the commercial imperative of vendors, the economic incentives for consultants, the influence of think tanks, and the bureaucratic needs of government, is essential for recognizing how the “fear machine” operates and how it can shape policy and public perception in ways that may not always align with a pragmatic assessment of actual risks. The real vulnerabilities, as we will explore, often lie in less sensational, but more fundamental, areas of governance and operational practice.

Our analytical framework for deconstructing the “fear machine” is grounded in a commitment to empirical data, rigorous historical analysis, and a critical examination of policy and industry claims. We reject speculative threat assessments and sensationalized narratives in favor of evidence-based reasoning. Our methodology relies on several key pillars. Firstly, we undertake a forensic examination of documented cyber incidents targeting U.S. OT infrastructure. This involves scrutinizing the scope, impact, and attribution of notable events to establish a factual baseline of actual damage, rather than relying on hypothetical scenarios or vendor-driven threat intelligence. We will assess the frequency and severity of past attacks to challenge the notion that catastrophic, nation-state-level attacks are common. Specific, verifiable incidents will be detailed, analyzing their true consequences and the lessons learned, or perhaps, not learned.

Secondly, we draw a crucial distinction between simulated exercises, vendor demonstrations of potential capabilities, and actual, successful breaches that resulted in significant disruption or damage. It is a common tactic within the “fear machine” to conflate theoretical vulnerabilities or exercises with historical realities. We will highlight how “near-misses” or the potential for exploitation, while important for awareness, are frequently presented as de facto successes or imminent threats, thereby inflating perceived risk. This analysis aims to clarify what constitutes a genuine cybersecurity failure versus a demonstration of potential risk, underscoring the importance of empirical evidence over hypothetical scenarios.

The interaction between these actors, commercial vendors, consultants, think tanks, and government agencies, creates a complex and often self-reinforcing ecosystem. Each has a vested interest, whether financial, political, or institutional, in maintaining and, at times, amplifying the perception of an escalating OT cybersecurity threat. This does not negate the reality of OT threats themselves, which are significant

and evolving; what we contest is the catastrophic framing, not the existence of risk. However, understanding these underlying motivations is crucial for dissecting the “fear machine” that has become so deeply embedded in the public and policy discourse surrounding operational technology. It allows for a more critical evaluation of the information presented, enabling stakeholders to distinguish between genuine security needs and the commercial or political imperatives that may shape the articulation of those needs. By recognizing who benefits from the narrative of constant siege, we can move towards a more balanced, evidence-based approach to securing OT systems, one that prioritizes effective risk management over uncritical acceptance of amplified anxieties.

An Evidence-Based Approach

This book’s analysis is grounded in empirical evidence. The methodology rests on three pillars:

- **Forensic Examination of Incidents:** We analyze documented OT incidents to establish a factual baseline, distinguishing verifiable events from hypothetical scenarios.
- **Distinguishing Capability from Impact:** We distinguish demonstrated capabilities (e.g., proofs-of-concept) from successful attacks causing disruption, as the two are often conflated.
- **Root Cause Analysis:** We investigate governance failures, human error, and poor maintenance as frequent causes of OT disruption, comparing their impact to malicious cyber intrusions.

This thesis would be challenged by clear evidence of widespread, catastrophic OT incidents caused by external adversaries.

What to Expect

The book is organized into four parts. Part One establishes an empirical baseline for OT incidents and governance. Part Two dissects the “fear machine” and its incentives. Part Three tests the catastrophe narrative against evidence. Part Four proposes reforms for a more resilient, evidence-based security posture.

CHAPTER 2

THE HISTORICAL RECORD: SCARCITY OF CATASTROPHE

The historical record of cyberattacks on Operational Technology (OT) is not one of constant catastrophe. While the potential for severe disruption is real, documented events reveal targeted intrusions and disruptive cybercrime rather than catastrophic outcomes. This chapter scrutinizes notable incidents to establish a factual baseline.

But establishing that baseline requires more than counting incidents. It requires reading each one twice: once as it is *commonly told*, and once as the *documented record* actually supports. The gap between those two readings is the subject of this chapter. Almost every famous OT cyber incident arrives in public discourse pre-loaded with a dramatic shape, a villain, a near-apocalypse averted at the last second. And almost every one, examined against primary sources, resolves into something narrower, stranger, and more instructive than the headline. The pattern is not that these events were fabricated. The pattern is that they were *mislabeled*, with the load-bearing details quietly inverted. A criminal extortion becomes a national security strike. An IT outage becomes a physical hack. A single contested phone call becomes a permanent fact. A laboratory test becomes documentary footage of a real attack. Where the previous chapter framed the problem in the abstract, this chapter does the forensic work case by case, distinguishing the narrative we inherited from the evidence we can actually point to.

A Framework for Analysis

This chapter uses a consistent framework to assess the historical record: a severity rubric, a table of documented incidents, and an acknowledgment of analytical limits. For each headline case, it also applies a two-column discipline, *narrative as commonly told* versus *what the documented record shows*, so that the reader can see exactly where the popular account departs from the sourcing. The goal is not to declare that nothing happened. It is to insist that what happened be described in the words the evidence will bear, and no stronger.

A Note on Methods and Limits Analyzing OT cyber incidents presents several challenges: **Underreporting:** Many incidents are never publicly disclosed due to reputational risk, legal concerns, or a lack of reporting requirements. The available data is an incomplete subset of all events. **Attribution Uncertainty:** Pinpointing the responsible actor is difficult, as evidence can be obscured and attribution often relies on classified intelligence or geopolitical inference rather than definitive technical proof. **Classification:** The line between an IT and OT incident can be blurry, especially when an IT breach (like ransomware) forces an OT shutdown as a precautionary measure.

The incidents in this chapter are drawn from primary sources (company disclosures, government alerts) and reputable secondary sources (cybersecurity firm reports, investigative journalism). While every effort is made to rely on verifiable data, some details regarding impact or attribution remain subject to the limitations outlined above.

One further caveat cuts the other way, and intellectual honesty requires stating it plainly. Some respected practitioners, the control systems engineer Joe Weiss prominent among them, argue that OT physical-impact incidents are *under-counted* rather than over-counted, precisely because OT environments often lack the forensic logging that would let an investigator distinguish a cyber cause from an equipment fault.¹ This is a fair methodological point and it is taken seriously here. The response is not to dismiss it but to be precise about scope: most of the incidents Weiss and others classify as under-counted are equipment-level malfunctions or insider mistakes, not the scenario the public fears, a remote adversary reaching across a network to seize physical controls and cause harm. Both things can be true at once. The long tail of murky, unlogged process upsets may indeed be larger than the public record suggests, *and* the specific, headline-grade, remote-attacker-destroys-infrastructure event remains conspicuously rare in the documented record. This chapter concerns the second claim.

OT/ICS Incident Severity Rubric

Level	Category	Description	Example
1	Nuisance	Minor operational anomalies or data exposure with no disruption to physical processes or safety.	Unauthorized network scan; temporary loss of view on an HMI.
2	Disruption	Localized and temporary disruption of a physical process; contained financial loss.	A single production line halted for several hours; minor data breach.
3	Significant Disruption	Widespread or prolonged disruption to a key service or production; significant financial loss; localized public impact.	Colonial Pipeline (2021): IT ransomware attack forces multi-day precautionary shutdown of OT.
4	Severe Disruption	Major, prolonged outage of critical infrastructure with regional impact; risk of injury or localized environmental damage.	Ukraine Grid Attacks (2015/2016): Coordinated attacks cause power outages for over 200,000 people for hours.
5	Catastrophic	Widespread, long-term failure of critical infrastructure resulting in mass casualties, severe economic collapse, or a national security crisis.	<i>No documented historical examples directly caused by a cyberattack.</i> Classification requires verifiable public evidence from national authorities, international bodies, or corroborated reporting from multiple credible journalistic sources. Speculation does not meet this threshold.

Date	Sector	Type	Impact	Downtime	Source(s)
2010	Nuclear	OT-Native (Sabotage)	Physical destruction of ~1,000 centrifuges.	N/A (Degradation)	Symantec, Langner
2014	Manufacturing (Germany)	IT→OT	Physical damage to a steel mill blast furnace.	Unspecified	BSI (Germany)
2015	Electric Utility (Ukraine)	OT-Native (Disruption)	Power outage for ~225,000 customers.	~6 hours (manual recovery)	E-ISAC, SANS
2016	Electric Utility (Ukraine)	OT-Native (Disruption)	Power outage in Kyiv.	~1 hour	Dragos, ESET

THE HISTORICAL RECORD: SCARCITY OF CATASTROPHE

Date	Sector	Type	Impact	Downtime	Source(s)
2017	Petrochemical (Saudi Arabia)	OT-Native (Sabotage)	Targeted SIS malware that could have enabled unsafe conditions; discovered after a fail-safe shutdown; no explosion, casualties, or public physical harm.	None (tripped to safe)	FireEye, Schneider Electric
2021	Water & Wastewater (USA)	IT→OT	Attempted manipulation of water chemical levels; later unconfirmed; no public impact.	None	Pinellas County Sheriff; FBI (later unconfirmed)
2021	Pipeline (USA)	IT→OT (Ran-somware)	Precautionary shutdown of major fuel pipeline; public evidence does not show OT compromise; panic-buying drove shortages.	~5 days	Colonial Pipeline, FBI, DOJ
2023	Water & Wastewater (USA)	IT→OT	Defacement of internet-exposed PLCs (default password); no public-health impact and no community loss of drinking-water service, though one booster station shifted to manual operation.	Minimal	CISA, Multiple Utilities
2021	Food & Agriculture (USA)	IT→OT (Ran-somware)	Shutdown of meat processing plants, disrupting food supply chain.	~1 day	JBS, FBI
2022	Manufacturing (Japan/Global)	IT→OT (Ran-somware)	Toyota suspended all 14 domestic plants (28 production lines) after supplier Kojima Press. Industry was hit by ransomware.	~1 day	Reuters, Feb 28, 2022; Toyota corporate statement

The table above lists the events most often cited as proof that the sky is falling. What follows is the careful reading of each, the part the citations usually skip. Read in this two-column way, the canonical set of “infrastructure cyberattacks” tells a consistent story: real intrusions, real expense, real lessons, and almost no instance of a remote adversary actually wresting physical control of a critical process and causing public harm.

Stuxnet (2010): the superpower's scalpel, sold as everyone's knife. The Stuxnet worm, discovered in 2010, brought OT cybersecurity into stark relief, and it remains among the most technically accomplished cyber-physical operations ever made public. *As commonly told*, Stuxnet is the proof that malicious code can reach out and physically wreck industrial equipment, an open-ended warning that what happened in Iran could happen to any plant with a programmable logic controller. *What the documented record shows* is narrower and, for the purpose of assessing commodity threats, almost the opposite of the popular lesson. Stuxnet was a nation-state operation, widely reported as a joint U.S.–Israeli program codenamed “Olympic Games,” developed over years at a cost that only a government could absorb.² It chained together multiple Windows zero-day vulnerabilities, used stolen, legitimate digital certificates to sign its drivers and appear trustworthy, and was engineered to jump an air gap into a network with no direct internet connection, reportedly via removable media. Its payload was not generic. It searched specifically for Siemens S7 controllers driving a particular configuration of centrifuge cascade, and ignored everything else. Its target was singular: the bunkered, air-gapped enrichment hall at Natanz. There it damaged on the order of 1,000 of the roughly 9,000 IR-1 centrifuges then installed at Natanz while replaying recorded, falsified “normal” sensor readings back to operators, so the machines tore themselves apart while the control room saw nothing wrong.³ Read correctly, Stuxnet is evidence of what a superpower can do to *one* fixed, high-value target, with enormous effort, bespoke intelligence, and years of preparation. That is a poor template for the commodity threats that ordinary utilities and manufacturers actually face. The barriers to entry that Stuxnet revealed, the zero days, the stolen certificates, the physical intelligence about a specific cascade, the air-gap-jumping logistics, are precisely why it was *not* the opening shot of a global sabotage wave. In the fifteen years since, no cyber-physical weapon of comparable sophistication and destructive effect has been documented against civilian infrastructure. The worm is rightly a landmark; its abuse comes from treating a state's scalpel as if it were a knife any criminal could pick up.

Ukraine grid (2015 and 2016): the famous “blackout hacks” that lasted hours because operators drove to the substations. The Ukrainian power grid attacks of 2015 and 2016 are the first widely confirmed cyberattacks to cause electrical outages, and remain the canonical examples, and they are state-sponsored work, attributed to the Russian military intelligence group commonly tracked as Sandworm, carried out against a country with which Russia was already at war.⁴ *As commonly told*, they are the proof that a grid can simply be switched off by remote attackers, the dress rehearsal for an

American blackout. *What the documented record shows* is a far more sobering lesson for anyone selling grid-collapse scenarios. In December 2015, attackers used stolen credentials to operate the distribution utilities' own control software, opening breakers and cutting power to roughly 225,000 customers; to slow the response, they also overwrote firmware on serial-to-Ethernet devices and flooded the call center with automated phone calls so customers could not report the outage.⁵ And yet customer power was restored within roughly one to six hours (though full automated operation was not restored for months). Not because the defenders out-hacked the attackers in real time, but because operators fell back to *manual* control: they physically drove to the affected substations and closed the breakers by hand.⁶ The 2016 follow-on, using the purpose-built ICS malware variously called Industroyer or CrashOverride, was more automated and more ambitious in design, capable of speaking grid protocols directly, yet in practice it knocked out a single transmission substation north of Kyiv for roughly an hour and a quarter.⁷ The headline writes itself: the most famous grid hacks in history, executed by a nation-state's military hackers, in the middle of a shooting war, kept the lights off for *hours*, because the people running the system could still throw the switches themselves. That manual fallback, the very "backwardness" that grid-modernization marketing often treats as a liability, was the resilience that bounded the damage.

TRITON / TRISIS (2017): "the world's most murderous malware," discovered because it failed safe. In 2017, malware later named TRITON (also TRISIS) was found inside a petrochemical facility in Saudi Arabia, where it targeted Schneider Electric Triconex Safety Instrumented Systems, the dedicated controllers whose only job is to bring a plant to a safe state if a process runs out of bounds. Because tampering with a *safety* system is the closest any documented attack has come to courting human casualties, TRITON has been described in the press as "the world's most murderous malware." The attribution points to a nation-state: the threat intelligence firm FireEye (now Mandiant) assessed that the activity was supported by a Russian government research institute, the Central Scientific Research Institute of Chemistry and Mechanics (CNIIMH), and in October 2020 the U.S. Treasury sanctioned CNIIMH in connection with the attack.⁸ *As commonly told*, TRITON is the moment the nightmare nearly arrived, an attacker reaching for the kill switch on a hazardous plant. *What the documented record shows* is that the operation was discovered only because it *failed*, and failed *safe*. The intruders' code inadvertently tripped the Triconex controllers into a protective shutdown, halting the plant. That shutdown, an engineering nuisance, was what exposed the entire campaign.⁹ The "disaster," in other words, was the safety system doing exactly what it

was built to do. The lesson is not that safeguards are futile; it is that even a state actor sophisticated enough to reach a safety controller was defeated by the layered, fail-closed design philosophy that governs how hazardous facilities are engineered.

Colonial Pipeline (2021): the pipeline that shut itself off because the billing system was down. In the United States, major OT cyberattacks causing catastrophic damage have been notably absent. Instead, the historical record shows a prevalence of less sensational, though still significant, types of intrusion, and no case illustrates the gap between narrative and record more cleanly than the May 2021 Colonial Pipeline incident. *As commonly told*, hackers, often imagined as a hostile government, seized control of the largest fuel pipeline on the U.S. East Coast and shut it down, triggering fuel shortages from a direct strike on critical infrastructure. *What the documented record shows* dismantles nearly every load-bearing element of that account. The attackers were the Dark-Side group, a financially motivated ransomware-as-a-service criminal operation, not a nation-state; President Biden stated publicly that there was “no evidence” the Russian government was responsible.¹⁰ The intrusion did not breach the pipeline’s operational technology at all. It encrypted Colonial’s *IT* and billing systems. The pipeline’s control systems were never hacked, Colonial shut the line off *itself*, preemptively, in substantial part because the ransomware had locked up the billing system and the company could not invoice for the fuel it moved.¹¹ The single point of failure was mundane: entry came through one legacy VPN account, protected by a password and no multi-factor authentication.¹² The fuel shortages that dominated the coverage were driven substantially by *panic-buying*, drivers topping off tanks and filling jerricans, rather than by any physical scarcity the shutdown alone had yet created. And the ending rarely makes the retelling: the Department of Justice subsequently clawed back 63.7 of the roughly 75 bitcoin Colonial had paid, recovering most of the ransom in bitcoin terms (though a smaller share of the dollar value originally paid, owing to the bitcoin price decline).¹³ The sharpest possible summary is this: *the pipeline’s control systems were never hacked, Colonial shut it off itself, largely because ransomware had locked up the billing system.* This was a serious, costly, genuinely disruptive crime. It was not a cyber-physical attack on a pipeline, and the distinction is the whole point. Critical infrastructure can suffer severe consequences from the *interconnectedness* of IT and OT and from the business decisions that interconnection forces, even when the OT itself is never touched. The event rightly prompted new federal directives for pipeline operators and a renewed focus on segmenting IT from OT, on backups, and on the unglamorous matter of turning on multi-factor authentication.

Oldsmar, Florida (2021): the water-poisoning that senior officials later called a “non-event.” The water and wastewater sector, despite being critical infrastructure, has also generated headline scares that shrink under scrutiny, and the 2021 incident at the water treatment plant in Oldsmar, Florida is the cautionary tale about how a single dramatic account hardens into permanent “fact.” Here the reality-check framing must be handled with particular care, because the honest position is not that we know what happened, but that an external attack was never established and the original account was later walked back by the officials closest to it.

As commonly told, an unknown intruder remotely accessed the plant’s controls and tried to poison the town’s water by raising the level of sodium hydroxide, lye, to dangerous concentrations; an alert operator, watching his screen, saw the cursor move on its own and dragged the setting back down just in time. That vivid story traces to essentially *one source on one day*: a press conference by Pinellas County Sheriff Bob Gualtieri shortly after the event.¹⁴ *What the documented record shows* in the years since is a striking retreat from that narrative. Roughly two years later, Oldsmar’s city manager, Al Braithwaite, publicly characterized the episode as a “nonevent,” saying the FBI had found “no evidence of any access from the outside” and that it was “likely the same employee” who had been “actually banging on his keyboard,” rather than an intruder.¹⁵ CyberScoop, reviewing the matter, reported that “the FBI was not able to confirm that this incident was initiated by a targeted cyber intrusion”.¹⁶

Two points must be stated precisely, in both directions. First, this book does *not* assert that Oldsmar was definitively operator error. The FBI’s conclusions were never publicly released in full, and certainty is not available; what can be said is that an external intrusion was never established, that senior officials later walked the dramatic version back, and that the matter remains, at best, unproven. Second, even taking the original, most alarming telling entirely at face value, no public harm was ever remotely close. By the sheriff’s own account, the altered sodium-hydroxide setpoint “would have taken from 24 to 36 hours to take effect,” and the change “would have triggered” the plant’s pH alarms long before any tainted water reached a customer.¹⁷ No water was contaminated. No one was endangered. The engineered safeguards, time delay, pH alarming, downstream monitoring, made public harm effectively impossible regardless of who moved the setting or why. Oldsmar thus does double duty as a cautionary tale: it warns about exposed remote-access on control systems, and it warns, just as loudly, about how a single uncorroborated account can become an unkillable data point in the threat narrative.

Cyber Av3ngers and Unitronics (2023–24): defacing PLCs with the password “1111.”

A more recent water-sector scare deserves the same treatment because it was widely reported as foreign hackers attacking American water systems. In late 2023 and into 2024, a group calling itself Cyber Av3ngers compromised Unitronics programmable logic controllers at water and other facilities. *What the documented record shows* is that the affected devices were internet-exposed and were reachable using the Unitronics default password, “1111”; the intruders changed the human-machine interface screens to display a political message.¹⁸ It was closer to defacement than to sabotage: beyond changing the HMI screens, the intruders disrupted a booster station’s pressure-regulating logic, forcing a switch to manual operation — a real but contained OT impact, with no documented loss of water service to any community. The episode is a real indictment of basic hygiene, devices that should never face the open internet, default credentials left unchanged, but it is not the sophisticated assault on the water supply that the headlines implied.

The German steel mill (2014): the canonical “physical destruction” case is a single anonymous paragraph. If any incident is treated as the bedrock proof that a cyberattack can physically destroy heavy industrial equipment, it is the German steel mill. *As commonly told*, hackers breached a steel plant, took control of the systems, and caused a blast furnace to be shut down improperly, inflicting “massive” physical damage, the closest thing on the books to a deliberate cyber-physical wrecking of a factory. *What the documented record shows* is far thinner than that reputation, and the gap here is the most important in the chapter. The core public evidence for the steel mill case still rests primarily on BSI’s sparse account: a *single anonymous, undated, untechnical paragraph*, roughly page 31, of the German Federal Office for Information Security’s (BSI) 2014 annual report.¹⁹ That paragraph names no company. It includes no malware sample, no forensic report, no indicators of compromise, no quantified figure for the damage, and no independent verification. In the eleven-plus years since, no public forensic report, malware sample, indicators of compromise, or quantified damage figure has ever publicly surfaced; later secondary analyses, including DOE/INL’s CyOTE work, have inferred a likely victim and timeline but have not produced a full technical writeup.

The intellectually honest conclusion is *not* that the event did not happen. It may well have. The conclusion is about the *thinness of the evidence*: the canonical proof that hackers can physically destroy industrial equipment is, to this day, a single anonymous paragraph in a government report, unaccompanied to this day by any public primary forensic report, malware sample, or quantified-damage figure. A claim that load-

bearing deserves load-bearing evidence, and this one has never had it. That an entire genre of “they can melt your factory” argumentation rests on one unsourced paragraph is itself a finding about the state of the evidence, and it is exactly the kind of soft spot that the broader threat narrative is built to keep readers from noticing.

Aurora (2007): a laboratory test, replayed for years as if it were an attack. A final, instructive entry is not an incident at all. The Aurora Generator Test was a 2007 experiment conducted by the Idaho National Laboratory, in which researchers used a cyber technique to open and reclose a diesel generator’s breakers out of phase, destroying the machine in a cloud of smoke and metal under controlled conditions.²⁰ The footage of that generator shuddering itself to death is genuinely arresting, which is precisely why it has been recycled for years in presentations and broadcasts, sometimes presented, or simply allowed to be understood, as if it were documentation of a real attack on the grid. It was a *test*, a demonstration of a real vulnerability class staged by defenders to study it, not an adversary striking a live target. Treating lab footage as field evidence is the demonstration-versus-incident confusion in its purest visual form.

Volt Typhoon (2023–24): real, current, state-backed, and not the same thing as an attack. The most contemporary entry must be handled with the same precision, in the opposite direction from the cases above, because here the danger is *under-reacting* to something real. Beginning in 2023 and continuing into 2024, U.S. authorities warned that People’s Republic of China state actors, tracked as Volt Typhoon, were “pre-positioning” themselves on the IT networks of U.S. critical infrastructure organizations, maintaining quiet, persistent access apparently to enable disruption at a future moment of Beijing’s choosing, per the joint CISA advisory AA24-038A.²¹ This is not fake, not overstated, and not a vendor demonstration. It is a genuine, ongoing, nation-state campaign and it belongs in any sober threat assessment. The discipline this chapter demands is simply to describe it in the words the evidence supports, and the operative distinction is *access versus impact*. Pre-positioning and persistent access are not the same as disruption. As of this writing, no physical effects, no outages, and no destructive attacks have been attributed to Volt Typhoon. The campaign is correctly understood as alarming *preparation*, the establishment of a foothold, and it warrants serious defensive response on exactly those terms. Reporting it as a blackout that has already happened would repeat the very error this chapter is cataloguing; ignoring it because no blackout has happened yet would be the mirror-image mistake.

Two further categories round out the documented record without changing its shape. The first is the theft of intellectual property and operational data from OT-

adjacent environments, frequently attributed to state-sponsored espionage, intrusions aimed at *learning*, not breaking. Reports have recurred for years of state-linked groups targeting U.S. energy companies, chemical manufacturers, and other industrial entities to exfiltrate proprietary information over long periods.²² The long-term costs to economic competitiveness and national security are real and hard to quantify, but the activity is intelligence-gathering, not the immediate, widespread physical destruction the word “catastrophe” conjures. The second is the cluster of *ransomware-driven IT→OT shutdowns* that fill out Table 2.1, JBS (food processing, ~1 day) and Toyota’s Japanese plants halted after a ransomware hit on its supplier Kojima Industries (~1 day).²³ Each was disruptive and costly. In each, as at Colonial, the harm flowed from a business decision to halt operations after an *IT* compromise, not from an adversary manipulating physical controls.

Why is catastrophe so rare in the documented record? Several factors contribute:

- **System Design:** Many OT systems remain segmented from IT networks or use specialized protocols, making remote access difficult. They are also built with physical safety systems, redundancies, and manual overrides designed to prevent catastrophic failure. The Ukrainian operators who drove to their substations, and the Triconex controllers that tripped TRITON’s own plant into a safe state, are not anecdotes; they are the mechanism by which severe attacks were bounded into hours-long or no-harm outcomes.
- **Attacker Requirements:** A successful cyber-physical attack requires immense resources, expertise, and specific intelligence about the target. This sophistication is typically limited to a few nation-state actors, who risk attribution and retaliation. Stuxnet is the proof of the price of admission: years of work, multiple zero-days, stolen certificates, and bespoke knowledge of one cascade, spent against a single fixed target.
- **Adversary Motives:** The most common cyber threats are driven by financial gain (ransomware) or intelligence gathering (espionage), not physical destruction. These motives lead to disruptive or costly incidents, but not necessarily catastrophic ones. The criminals behind Colonial and JBS wanted to be *paid*; breaking a pipeline or a slaughterhouse would have destroyed the leverage they were selling back.

- **Defensive Efforts:** While imperfect, cybersecurity defenses are continuously improving. Increased awareness, investment in security, and government regulation have hardened many potential targets, though, as Oldsmar's exposed remote access and Cyber Avengers' "1111" passwords show, the gap between what we know to do and what is actually done remains wide.

In conclusion, the historical record reveals a landscape of persistent threats, disruptive incidents, and significant espionage, but a notable absence of catastrophic cyber-physical attacks on U.S. critical infrastructure. Read case by case, the famous exceptions prove the rule: the superpower's scalpel at Natanz, the wartime grid hacks undone by hand, the "murderous" malware caught failing safe, the pipeline that shut itself off, the water-poisoning that senior officials later called a nonevent, and the "physical destruction" case that survives only as one anonymous paragraph. This evidence counters narratives that overemphasize doomsday scenarios and supports security strategies focused on resilience against observed threats, not just hypothetical cases.

The discourse surrounding cybersecurity, particularly concerning critical infrastructure and Operational Technology (OT) systems, is often characterized by a stark dichotomy: either a state of pervasive, imminent catastrophe or an underestimation of genuine threats. Within this narrative, a crucial, yet frequently blurred, line exists between simulated exercises, vendor demonstrations of potential capabilities, and actual, successful breaches that have resulted in tangible disruption or damage. This distinction is not merely semantic; it is foundational to understanding the empirical reality of OT cybersecurity and informing effective policy and investment. The frequent conflation of "near-misses" and theoretical vulnerabilities with historical realities can lead to misallocation of resources, undue panic, and a skewed perception of risk. A rigorous examination is required to clarify the difference between a genuine failure and a demonstration of potential risk, underscoring the importance of empirical evidence over speculation.

Distinguishing Reality from Simulation

Cybersecurity vendor demonstrations and red team exercises, while valuable for illustrating potential threats, operate in controlled environments. A simulated attack against a test system, where defenders are aware of the exercise, bears little resemblance to a real-world breach. These exercises reveal potential vulnerabilities, not historical

precedent; they show what *could* happen, not what *has* happened. The Aurora Generator Test is the cleanest illustration of why the line matters: a controlled laboratory experiment whose footage has been replayed for years as if it documented a real attack on the power grid.²⁴ The smoke was real; the attack was staged.

These demonstrations understate the complexity of a real-world attack, which involves overcoming unexpected network behaviors, diverse legacy hardware, and active security measures. Furthermore, real-world attackers are not necessarily motivated by the same objectives as a vendor showcasing a product. They may prioritize stealth for espionage, rapid financial gain through ransomware, or strategically timed disruption, and their methods will adapt to these goals, often involving prolonged reconnaissance and exploiting human or procedural weaknesses in addition to technical ones.

Similarly, penetration testing and “red teaming” exercises, when conducted properly, are invaluable tools for assessing an organization’s defenses. These activities involve skilled professionals attempting to breach systems using methodologies that mimic real adversaries. However, the outcomes of these exercises are, by definition, controlled and contained. The participants are typically authorized, their objectives are defined, and the scope of their actions is delineated. A successful penetration test might demonstrate that a particular vulnerability exists and can be exploited, leading to unauthorized access or control. This is critically important for improving defenses. However, a red team’s success should not be equated with a successful external attack causing widespread damage.

The very nature of a red team engagement means that the organization is aware of the exercise, has likely provided intelligence to the testers about the scope, and has established protocols for responding to a simulated breach. The absence of genuine catastrophe in such controlled scenarios, despite simulated success in gaining access, often points to the inherent resilience of OT systems and the effectiveness of layered security and operational safeguards. The fact that a red team can bypass certain defenses and gain access to a control network does not automatically imply that a nation-state actor or sophisticated criminal group could do the same in a live, unannounced attack, nor that they would achieve the same level of disruptive impact without detection. The critical difference lies in the element of surprise, the sustained operational tempo of a real adversary, and the potential for cascading failures in a complex, interconnected system that a red team might not be authorized or equipped to fully explore.

A related confusion sits one rung above the red team: the gap between an adversary’s *access* and an adversary’s *impact*. The Volt Typhoon warnings of 2023–24 are the