The Game You Are Already Playing

A CISO's Field Manual for Strategic Cybersecurity in Critical Infrastructure

Sara Malik[1,*] and Naveed A.A.[1]

[1]PakCrypt NPO, Islamabad, Pakistan
[*]Corresponding author: smk@pakcrypt.org

April 22, 2026

Abstract

Cybersecurity in critical infrastructure is conventionally treated as a technical problem of closing gaps. This framing has reached the limits of its usefulness. Compliant organisations continue to be breached; mechanistic adversaries succeed against well-resourced defenders; and the hardest CISO problems are rarely technical. This paper argues that cybersecurity is more accurately understood as an unending strategic interaction between agents with conflicting goals, incomplete information, asymmetric resources, and the capacity to learn from one another. Drawing on three decades of academic work in game theory, security economics, and the lived experience of CISOs in energy, finance, and defence, we offer a field manual for playing the game deliberately. The paper introduces the **PAPI-T** frame (Players, Actions, Payoffs, Information, Time), a decomposition of "the adversary" into distinct player types with incompatible utility functions, and three power levers — Information Control, Commitment, and Mechanism Design — through which a CISO changes outcomes. We illustrate these constructs with canonical incidents from 2010 through 2025 across the three sectors, examine coordination problems with boards, vendors, insurers, and regulators, and conclude with principles for operational rhythm and for losing well. Security is a verb, not a state.

Keywords: game theory; cybersecurity; CISO; critical infrastructure; energy; finance; defence; mechanism design; commitment devices; Stackelberg security games.

Contents

Preface: Security is a Verb

There is a comforting fiction that circulates at cybersecurity conferences, in board decks, and in the marketing copy of every vendor you have ever met. It goes like this: *security is a state your organisation achieves, a plateau you reach after enough investment, a destination at the end of a roadmap.*

This is nonsense. It has always been nonsense. And if you are a CISO in critical infrastructure, believing it will eventually cost you your job, your organisation's operational continuity, or — in the sectors this paper addresses — human lives.

Security is not a state. Security is a game. More precisely, it is an unending series of strategic interactions between people with conflicting goals, incomplete information, limited resources, and the capacity to learn from each other. The attacker adapts. The regulator changes the rules. Your CFO re-prioritises. Your users find workarounds. Your vendors externalise their risk onto your balance sheet. And the game continues, whether you are paying attention or not.

The question is not whether you are playing. You are playing. The question is whether you are playing *well*.

This paper is for CISOs who suspect the checklist has reached the limits of its usefulness — who have watched compliant organisations get breached, incompetent adversaries succeed, and brilliant engineers fail to secure basic things because the incentives around them were wrong. It is for CISOs who have noticed that their hardest problems are not technical. The ransomware itself is trivially understood; the reason your organisation is vulnerable to it involves fifteen people you have never met and a contract you were not consulted on.

What follows is a field manual for playing the game deliberately. It borrows from game theory, from economics, from the hard-won lessons of deployed security systems, and — more than anything — from the accumulated experience of CISOs who learned these things the expensive way. It is organised not as theory followed by application, but as a sequence of moves you will actually make, in roughly the order a working CISO makes them.

Three warnings before we begin.

First, this is not a mathematical treatise. The academic game-theoretic literature on cybersecurity is rich and worth reading, but it is not what you need at 2 a.m. on a Tuesday when your SOC lights up. What you need is a vocabulary and a set of mental models. The math is scaffolding; we will show you the building.

Second, the frame assumes three sectors: energy, finance, and defence. The incidents we reference, the regulators we name, and the adversaries we describe come from those three domains. If you work in healthcare, water, transport, or telecom, the principles translate, but the specifics will differ. Translate with care.

Third, and most important: nothing in this paper will make your organisation unhackable. That is not the goal. The goal is to make your organisation's defeat expensive enough, unreliable enough, and reputationally dangerous enough for the attacker that rational adversaries go

elsewhere and irrational ones fail often enough to be detected. That is what winning looks like. It is less satisfying than the marketing copy. It is also achievable.

1 Seeing the Board

1.1 Why Checklists Fail

Every CISO has been handed a control framework. NIST CSF. ISO 27001. CIS Controls. NERC CIP. CMMC. NIS2's annexes. DORA's Regulatory Technical Standards. There are more frameworks than there are days in the compliance quarter.

Frameworks are useful. They codify things you would otherwise forget. They give auditors something to measure. They provide a common vocabulary when your Chief Risk Officer asks a question. They are also, taken alone, insufficient — and the reason they are insufficient is not that they are wrong but that they answer the wrong question.

A framework answers: *what should we do?* A strategic view answers: *what should we do, given what they will do, given what we have done?* The difference is not cosmetic. It is the difference between a chess opening book and actually playing chess. Openings are essential. Memorising openings is not the same as winning.

Consider Equifax. In 2017, Equifax lost the personal data of 147.9 million Americans, 15.2 million Britons, and roughly 19,000 Canadians. The breach began with an unpatched Apache Struts vulnerability, CVE-2017-5638, that had a patch available. It escalated because an SSL certificate had expired on an egress inspection device, rendering the organisation blind to the exfiltration for 76 days. Four members of the People's Liberation Army's 54th Research Institute were later indicted. Equifax settled for at least $700 million with the FTC and CFPB, $425 million with the states, and absorbed roughly $1.38 billion in total costs.

Equifax was SOC 2 compliant. It had a framework. It had been audited. It had a CISO. The organisation was, in every measurable checklist sense, *doing the right things*. What it was not doing was playing the game — noticing that a patch had been released, that an adversary was scanning for unpatched instances, that the cost of missing this patch was the value of the crown jewels, and that a blind egress monitor was worse than no egress monitor because it created false confidence.

The lesson is not that Equifax was uniquely incompetent. The lesson is that *the checklist does not see what the adversary sees*. A checklist sees line items. An adversary sees opportunities, incentives, and your weakest link. The work of a strategic CISO is to close that perceptual gap.

1.2 PAPI-T — The Five Things on the Board

Before you can play a game, you have to know what game you are playing. Every strategic interaction has five components. We will give you an acronym because acronyms are the tax we pay for memory, and call this one **PAPI-T**:

Players. The decision-makers whose choices affect outcomes. In cybersecurity, the players are not just you and the attacker. They include your board, your CFO, your employees, your vendors, your regulators, your insurers, and — critically — the other organisations in your sector, whose decisions shape your own. A CISO who models only themselves and the attacker has already lost half the game.

Actions. The moves available to each player. The attacker's actions include reconnaissance, initial access, lateral movement, persistence, exfiltration, encryption, and extortion. Your actions include prevention, detection, response, recovery, and — critically — signalling, negotiation, and regulation. Your CFO's actions include funding you, reducing your funding, firing you, or moving a dollar from security to sales. Your board's actions include calling for an external review, setting a risk appetite, or doing nothing.

Payoffs. The value each player assigns to each outcome. This is where most CISO thinking goes wrong. Payoffs are not dollar amounts alone. For the ransomware affiliate, the payoff is expected ransom minus expected cost of attack, discounted by the probability of successful laundering. For the Chinese state adversary conducting long-term access pre-positioning, the payoff is strategic optionality — an option to disrupt at a time of geopolitical choosing — for which no spot ransom payment can substitute. For your CFO, the payoff is enterprise value, which security affects in both directions. For your employee who shares a password with a colleague, the payoff is productivity plus convenience minus a small probability of consequences. The players' payoff functions are different, and if you imagine them to be the same as yours, you will continue to be surprised.

Information. What each player knows when they act. This is the single most underappreciated element of the game. The attacker sees your external attack surface. You see your internal configuration. Neither of you sees the other completely. Your board sees what you tell them; if you tell them only good news, they make decisions as if only good news exists. Your users see what they need to do to get their work done; they do not see your threat model. Information is not just an input to the game — it is the most valuable asset in the game, and the strategic CISO treats it as such.

Time. The dimension the academic textbooks leave out and the practitioners cannot afford to. Security games are not one-shot. They are repeated, continuous, and asymmetric in their temporal structure. The attacker can wait; you cannot. The attacker can invest in a zero-day for three years; you must respond in hours. The regulator moves on a different clock than the incident. Time is why patch cadence matters, why dwell time is the metric that actually matters, and why the FlipIt model — Ron Rivest's formalisation of continuous-time stealth-takeover games — is the right mental picture for advanced persistent threats. No deterministic rotation is optimal against an adaptive attacker. Randomisation beats schedule.

Every time you make a security decision, ask yourself the PAPI-T questions. Who are the players? What are their actions? What are the payoffs? What does each player know? What

is the time structure? You will be astonished how often the answer to one of these questions makes the right decision obvious, and how often the decision you were about to make becomes visibly wrong.

1.3 Knowing Your Adversary — The Decomposition

There is no such thing as "the attacker." There are attackers, plural, and they have almost nothing in common with each other. A CISO who defends against "the attacker" is defending against a phantom. A CISO who defends against *this* attacker, with *this* motivation, *this* capability, and *this* target selection logic, is defending against something real.

Here is the decomposition you need. It is not exhaustive, but it will cover most of what you face in energy, finance, and defence.

Ransomware affiliates. These are financially motivated small businesses. LockBit, Black Basta, ALPHV, Cl0p, the rotating cast of rebrands. They operate on roughly an 80/20 affiliate-to-developer revenue split. They care about dwell time only insofar as it lets them stage exfiltration before encryption. They are patient enough to wait weeks for the right moment and impatient enough to move on to a softer target if your defences cost them time. They respond to economics because they are an economic enterprise. Their median initial access broker listing on the underground runs $500 to $1,200. They appear on leak sites 23 to 36 days after initial access. If you can raise their cost by a factor of two, many of them will stop attacking you and attack someone else. This is the game where the weakest-link framing applies most cleanly.

State disruptors. These are adversaries whose utility function is political, not financial. Sandworm — Russia's GRU Unit 74455 — turning off Ukrainian substations in 2015 and 2016. XENOTIME deploying Triton against a Saudi petrochemical plant's safety systems in 2017 — the first malware designed, in the Dragos assessment, to potentially kill people. Volt Typhoon quietly establishing footholds in US critical infrastructure for what CISA's February 2024 advisory characterised as at least five years, with no exfiltration and no ransom, for the purpose of optionality in a future Taiwan contingency. These players do not respond to economic cost-raising. Raising their cost by a factor of two is rounding error against a state budget. What you can do is raise the *probability of detection* and the *time-to-effect*, both of which affect strategic value. A dormant access that must remain undiscovered for five years is a very different asset than one that must remain undiscovered for one year.

State spies. These are patient information collectors. APT29 — Russia's SVR, the SolarWinds actors, the HPE and Microsoft corporate email intruders. APT41, the Chinese MSS-affiliated dual-use espionage and financial operators. The teams that produced SUNBURST, Operation Cloud Hopper, and the Storm-0558 Microsoft cloud key compromise. Their time horizon is years. Their target selection is strategic: defence primes, government contractors, political institutions, and — increasingly — software supply chain upstream of all of the above. You cannot make yourself boring to them by being well-patched. If you are in the defence industrial base or adjacent to it, you are on their list whether or not you want to be.

Hacktivists. Ideologically motivated, low-capability-but-high-noise. They will deface your website, leak your documents, or DDoS you on a symbolic anniversary. They are real but usually not existential.

Insiders. The category everyone forgets until it happens. Sometimes malicious — Teixeira, Snowden, Manning at the state level; the Ubiquiti case at the corporate level. More often careless — the employee who ignores policy because the policy makes their job harder. The single most expensive cybersecurity event at your organisation over the next five years is substantially more likely to involve an insider, a vendor, or a misconfiguration than a nation-state zero-day. Design accordingly.

Suppliers. Not adversaries exactly, but players whose decisions create your exposure. Your MSSP. Your software vendors. Your payroll provider. Your managed file transfer system — see MOVEit in May 2023, approximately 3,000 US and 8,000 global downstream victims from a single Progress Software vulnerability. The cloud provider whose signing key was compromised and whose incident report you found unsatisfying — see Storm-0558 in 2023, and the Cyber Safety Review Board's April 2024 conclusion.

Now consider what this decomposition implies. You are not defending against one player. You are defending in a multi-game tournament, where different opponents at different tables have different strategies and different payoffs. The defences that work against one may be useless against another. MFA kills credential stuffing against ransomware affiliates and does nothing against a supply-chain compromise. Network segmentation makes state-disruptor pre-positioning harder and does nothing against an insider with legitimate access. Cyber insurance transfers financial risk from ransomware and does nothing to protect against strategic disruption that has no insurable payout equivalent.

The first question of a strategic CISO is never *what is the threat?* It is always *which threat, to which asset, via which path, by which adversary?* The answer determines which game you are playing, and which game you are playing determines which moves make sense.

2 The Three Levers

You have three fundamental ways to change the game. Not fifty. Not the long list of controls in the CIS Top Eighteen. Three. Every useful security action is, at its core, an exercise of one or more of these three levers. Understanding this is what separates a CISO who reacts from a CISO who plays.

2.1 Information Control — The Fog of War Goes Both Ways

The attacker has advantages you do not have. They see your external surface. They can scan at their leisure. They can buy credentials on underground markets. They can read your LinkedIn pages and your SEC filings and your vendor press releases.

You have an advantage they do not have. You see your internal network. You know what your

normal looks like. You control the ground they are trying to traverse. The question is whether you use it.

Information control has two faces: what you reveal and what you conceal.

Conceal what hurts you. Your internal architecture. The specific tools your SOC uses. The location of your crown-jewel assets. Your incident response playbook. The vendors with access to your most sensitive systems. These should not be on your external website. They should not be in job postings that list specific product names. They should not be casually mentioned at conferences. This is not paranoia; it is basic operational security, and most organisations practise it badly.

Reveal what deters. That you investigate every alert. That you have an active threat hunting programme. That you prosecute. That you cooperate with law enforcement. That your defensive posture is unpredictable. Deployed Stackelberg security systems — Milind Tambe's ARMOR at LAX, IRIS for the Federal Air Marshals, PROTECT at US ports — work precisely because they commit publicly to randomised, payoff-weighted defence. The attacker cannot optimise against a policy they cannot predict. You cannot use the academic Stackelberg solver for your network, but you can use the principle: vary your inspection schedule, vary your detection cadence, make yourself unpredictable.

Deceive intelligently. Honeypots and honey tokens have a long, mixed reputation. Against a sophisticated adversary running a careful operation, they are probably not going to fool the human operator for long. But that is not the game. The game is to make the adversary *uncertain*, to raise their cost of distinguishing real from fake, to slow them down, and to give yourself a high-fidelity alert when someone touches the tripwire. Thinkst Canary and similar deception platforms do this well because they optimise for low false-positive rates, not for fooling the sophisticated. A single canary file in your file share that no legitimate user should ever touch is worth more than ten additional SIEM rules, because the base rate of false positives is nearly zero. If it fires, something is wrong.

Tell your board the truth. This is the information-control move that CISOs most often get wrong, and it is the one with the highest return. The temptation — especially after a near-miss — is to present the board with a sanitised summary. "We have a robust security programme." "We are aligned with NIST CSF." "Incidents were contained." This is information control against your own organisation, and it is a strategic mistake. If the board does not know your near-misses, they underfund you. If they underfund you, the near-misses become hits. If hits happen, you are blamed for not telling them about the near-misses. The information asymmetry that served you in the short run eats you in the long run.

The rule is: reveal ugly truth upward, reveal nothing of operational value outward, and make the external observer uncertain about what you know and what you are doing. Most organisations do this exactly backwards — they broadcast their architecture and conceal their failures.

2.2 Commitment — The Strategic Value of Burning Bridges

One of the most counterintuitive results in game theory is that reducing your own options can make you stronger. Thomas Schelling demonstrated this in *The Strategy of Conflict* in 1960: if two trucks are approaching each other on a narrow bridge, the driver who visibly throws his steering wheel out the window wins the game. He has committed. He cannot swerve. The other driver must.

Cybersecurity is full of places where commitment changes outcomes. The CISOs who understand this get their budgets. The ones who do not spend their careers negotiating.

Commit to the regulator. This is the single most powerful commitment device available to a critical-infrastructure CISO, and it is dramatically underused. You are not asking your CFO for money. You are informing your CFO of a binding legal constraint. "We cannot legally operate the SCADA estate without this control. NIS2 Article 21 requires it. The consequence of non-compliance is a fine of up to 2% of global turnover and personal liability for named members of management." The CFO is no longer negotiating with you; they are negotiating with the European Commission, and they will lose. This is not a rhetorical trick. It is the explicit design of NIS2, DORA, NYDFS Part 500, CIRCIA, the SEC cyber disclosure rule, and the CMMC programme. These regulations exist, in significant part, because legislators recognised that internal negotiation between CISOs and CFOs produces suboptimal security outcomes, and external commitment was needed to fix the coordination failure. Use the tool that was built for you.

A caveat: commitment via regulation only works if the commitment is *credible* and *costly to violate*. NYDFS's 500.17(b) certification requires a named CEO and CISO to personally attest; the April 15 deadline is real; enforcement actions against First American, Residential Mortgage, and PayPal demonstrate that fines actually happen. SEC Form 8-K Item 1.05 requires a four-business-day disclosure of material cyber incidents; the *in terrorem* effect of the SolarWinds Tim Brown case has already changed disclosure behaviour across industries even though the SEC dismissed remaining charges on 20 November 2025. DORA's TLPT and third-party oversight regime, effective from 17 January 2025, has already produced concrete compliance spending — 38% of EU CISOs and 47% of UK CISOs surveyed report DORA compliance costs above one million euros. These are credible commitments. Contrast them with PCI Level 2 self-assessment, or internal policy documents that nobody enforces: those are soft commitments, which adversaries and employees alike learn to ignore.

Commit to automation. When you configure your EDR to automatically isolate a host on high-confidence malicious activity, you are making a commitment. A human who can be social-engineered cannot override that rule in the moment. The attacker who calls your helpdesk at 2 a.m. and begs them to let him back on the VPN is negotiating with a machine, not a person. The machine will not negotiate. This is the MGM case in reverse — in September 2023, a ten-minute social-engineering phone call to the MGM IT helpdesk enabled Scattered Spider and ALPHV to compromise over a hundred ESXi hypervisors and cost MGM over $100 million in Q3.

A helpdesk that could not reset credentials without out-of-band verification could not have been social-engineered that way. The cost of automation is rigidity. The benefit is unnegotiability.

Commit to the CEO. The hardest commitment is the one you make to yourself and your CEO: what you will do, and what you will not do, when something goes wrong. The CISOs who survive the current liability environment are the ones who have established, in writing, before incidents, what the CEO's decision authority is, what the CISO's decision authority is, who gets notified and when, what the disclosure timeline is, and what happens if there is disagreement. Joe Sullivan's conviction in 2022 and the Ninth Circuit's 2025 affirmation were not primarily about the breach. They were about the cover-up — about disclosure choices made under pressure without a prearranged framework. The lesson is not "always disclose immediately." The lesson is "decide in advance, in writing, with legal counsel, what your disclosure framework is, so that you do not make the decision in the panic of a 2 a.m. incident call."

Commitment is not bravery. Commitment is discipline. You remove your own future optionality so that you cannot be bullied, negotiated with, or confused in the moment. Every CISO we have seen survive a major incident has made these commitments in advance. Every CISO we have seen fail has made them in real time.

2.3 Mechanism Design — Changing the Rules of the Game

The third lever is the most elegant and the most underused. Instead of playing harder inside the current rules, you change the rules so that playing well is the rational choice for every player.

Consider the MFA problem. A CISO who treats MFA as a policy — "you must use MFA" — is playing a losing game against human nature. Users will find workarounds. They will share tokens. They will fatigue-spam themselves into approving attacker push notifications. They will resent the CISO. The CISO will spend political capital on enforcement. The programme will slowly decay.

A CISO who treats MFA as a mechanism-design problem asks a different question: how do I make the secure path the *easy* path? The answer is passwordless single sign-on with phishing-resistant FIDO2 authenticators. The user does not fight MFA because there is no password to fight. Compliance becomes the path of least resistance. The CISO has changed the game so that the user's rational choice — minimise effort — aligns with the security outcome. The control is invisible; the security is strong; the political cost is zero. This is not a fantasy. This is how Google, after a phishing incident in 2017, eliminated account compromise across 85,000 employees: mandatory hardware security keys, implemented as the easiest way to log in. The mechanism solved the problem that the policy could not.

Or consider software supply chain. A CISO who reads the SBOM and chases every CVE is playing a losing game of whack-a-mole. A CISO who requires their critical vendors to carry cyber insurance with contractual security minimums, who uses DORA's Critical Third-Party Provider regime to delegate oversight of the biggest vendors to regulators, and who contractually caps supplier data retention — has changed the game. The supplier now has skin in the game.

The insurer now audits the supplier. The CISO has outsourced part of the work to players whose incentives are aligned with the outcome. This is mechanism design.

Or consider insider risk. A CISO who runs phishing simulations and sends frustrated emails to repeat clickers is treating the symptom. A CISO who redesigns the expense-reporting process so that it does not require clicking suspicious-looking PDFs, who implements least privilege so that a single user's compromise does not cascade, and who creates a blame-free reporting culture so that employees report their own clicks rather than hiding them — has changed the game. The user is no longer the adversary; the process is. And processes can be redesigned.

The core principle of mechanism design, stated most clearly by Ross Anderson in his work on the economics of security, is: *align incentives with outcomes.* Wherever a security outcome depends on someone doing something costly for reasons that do not benefit them, you have a mechanism-design problem, not a compliance problem. The solution is not louder enforcement. The solution is to redesign the mechanism so that the aligned choice is the easy choice.

Two caveats. First, mechanism design is slower than policy. You cannot redesign a process in a sprint. It takes months, sometimes years, and requires cross-functional support. Second, it is expensive in the short run and cheap in the long run. Budget and communicate accordingly.

One more thing. Mechanism design is how you solve the weakest-link problem. Hal Varian's taxonomy — weakest-link, sum-of-efforts, best-shot — tells you which mechanism matters when. Passwords are weakest-link: the worst password on the network is the one that matters. MFA raises the floor for everyone. Patch management is sum-of-efforts: the overall posture depends on everyone patching. Detection is best-shot: one excellent threat hunter adds more than ten mediocre ones. Design each mechanism for the game it is actually playing. Using weakest-link tools on sum-of-efforts problems is why MFA and Zero Trust did not prevent SolarWinds.

3 The Sectors

The levers are universal. The details are not. What follows is the shape of the game in each of the three sectors this paper addresses. Read your own sector closely. Read the other two carefully — the patterns that are obvious in one are often obscured in another, and the cross-sector reader learns faster.

3.1 Energy — The Kinetic Sector

Energy is the sector where cyber becomes physical. This is both the reason it is interesting and the reason it is terrifying. In finance, a breach costs money. In defence, a breach costs secrets. In energy, a breach can cost lives.

The canonical case history begins on 23 December 2015, when Sandworm — Russia's GRU military intelligence unit — used the BlackEnergy malware to cause the disconnection of roughly 30 substations in western Ukraine, interrupting electricity for approximately 225,000 customers for between one and six hours. It was the first confirmed cyberattack to cause a power outage. The attack was remarkable less for its technical sophistication than for its operational discipline:

the attackers had been inside the utilities for months, they coordinated their attack across three regional distribution companies simultaneously, and they disabled the utilities' ability to remotely restore service, forcing manual field restoration.

One year later, on 17 December 2016, the same actor group deployed Industroyer — also known as CrashOverride — against a transmission substation in Kyiv. Unlike BlackEnergy, Industroyer was purpose-built grid malware. It spoke IEC-101, IEC-104, IEC-61850, and OPC DA — the native protocols of electrical substations. It was the first malware designed from the ground up to attack the physics of electricity delivery, not the IT systems around it. The outage was shorter — about an hour — but the message was longer: a state actor had demonstrated a repeatable, scalable capability against grid infrastructure.

In June 2017, the game changed again. XENOTIME — attributed to Russia's Central Scientific Research Institute of Chemistry and Mechanics — deployed Triton, also called Trisis, against a petrochemical facility at Petro Rabigh in Saudi Arabia. Triton targeted Schneider Electric Triconex safety instrumented systems — the systems whose entire purpose is to prevent explosions and release of toxic materials. It is the first malware known to be designed to potentially kill people. The attack failed because of a bug in the malware itself, not because the defences worked.

Then came Colonial Pipeline on 7 May 2021. A DarkSide ransomware affiliate used a single compromised VPN account without MFA to gain access to Colonial Pipeline Company's IT network. The company proactively shut down its 5,500-mile pipeline — one of the largest fuel pipelines in North America — to prevent lateral spread to OT. It paid a $4.4 million ransom, of which $2.3 million was later recovered by the FBI. The pipeline was down for six days. Gasoline shortages spread across the US East Coast. The President of the United States addressed the nation. The attack was not technically sophisticated. It was boring. That is the point: a weakest-link MFA gap in an IT VPN produced a national-security-relevant physical outcome.

On 24 February 2022, the day of Russia's invasion of Ukraine, Viasat's KA-SAT satellite network was disabled by the AcidRain wiper. Approximately 40,000 to 45,000 modems across Europe were bricked. A collateral consequence was that roughly 5,800 Enercon wind turbines in Germany lost their remote control capability. The outage affected a NATO country in a manner the attacker almost certainly did not specifically intend, demonstrating that in interconnected infrastructure, externalities are the rule, not the exception.

In January 2024, a Russia-linked group deployed FrostyGoop against the Lviv district heating utility, using the first malware known to directly abuse Modbus TCP port 502. Roughly 600 apartment buildings lost heat for nearly two days in sub-zero conditions. The attack's elegance was in its minimalism: Modbus TCP is one of the most widely deployed industrial protocols in the world, with approximately 46,000 internet-exposed devices globally, and FrostyGoop demonstrated that the protocol can be weaponised without an implant — legitimate Modbus commands, sent with malicious intent, cause legitimate devices to do malicious things.

And then there is Volt Typhoon. CISA Advisory AA24-038A of February 2024, co-signed by NSA, FBI, DOE, EPA, TSA, and the Five Eyes partners, concluded that PRC-sponsored actors

had maintained footholds in US critical infrastructure for at least five years with no exfiltration and no ransom, for the purpose of disruption in a future Taiwan contingency. This is the game-theoretically most important development of the decade. It breaks the expected-value model of the cyber-criminal game. There is no short-term payoff. The payoff is strategic optionality. The defender cannot calculate an expected loss because the payoff function is not probabilistic; it is conditional on a geopolitical trigger that may or may not fire.

If you are an energy sector CISO, your game has four players you must take seriously: ransomware affiliates, who treat your IT network as a payday; state disruptors, who treat your OT network as strategic leverage; hacktivists, who treat your brand as a target of opportunity; and regulators, who are watching all three and will hold you accountable for all three.

Your defensive strategy must therefore be multi-game:

- **Against ransomware affiliates**, the weakest-link model applies. Raise the baseline: MFA everywhere, EDR on every endpoint, ruthless patch cadence on externally-exposed systems, network segmentation between IT and OT, immutable offline backups tested quarterly. Colonial Pipeline would not have been Colonial Pipeline with MFA on the VPN. This is straightforward economics.
- **Against state disruptors**, raise the cost of discovery and the cost of persistence. Volt Typhoon relies on living-off-the-land techniques; they are harder to detect than malware-based intrusions precisely because they use legitimate administrative tools. Your response is behavioural detection, baseline monitoring of what normal looks like for privileged accounts, and aggressive alerting on any deviation. You also need OT-aware monitoring; Dragos's 2025 Year in Review reported that 17% of industrial organisations still run shared IT/OT domains, which is a gift to any attacker who achieves IT access. Segment aggressively.
- **Against hacktivists**, your response is resilience and communications. They will not usually cause lasting damage. They will damage your reputation if you are caught unprepared. Have an incident communications plan. Know who speaks. Know who does not.
- **Against regulators**, commit early. NERC CIP in North America, NIS2 in the EU, CAF in the UK, SOCI in Australia. Use the regulator as your commitment device. Map every control in your programme to a specific regulatory requirement. When the CFO asks why, you are not defending a judgement; you are reporting a legal constraint.

One last point specific to energy. The asset lifecycles in OT are 20 to 40 years. Some of the controllers in your substations were installed before anyone had heard of cybersecurity, and replacement cycles are measured in decades. This means many of your defensive *actions*, in the PAPI-T sense, are unavailable — you cannot patch what the vendor no longer supports, and you cannot replace what costs ten million dollars per unit and takes three years to procure. Accept this. Work around it with compensating controls, network segmentation, strict access control, and monitoring. A mature energy CISO spends more time on compensating controls than on primary controls, because the primary controls often do not exist.

3.2 Finance — The Repeated Game

Finance is the sector where the game-theoretic frame fits best. This is not a coincidence. Finance invented much of modern risk management. It has the tightest regulatory apparatus, the most mature information-sharing infrastructure, and the clearest dollar-denominated payoff functions. It is also where the repeated-game dynamic is most visible: you will be attacked, again and again, by the same kinds of adversaries, using mostly the same techniques, for the same reason.

The canonical case is the Bangladesh Bank heist of 4 and 5 February 2016. Lazarus Group — DPRK-attributed, Reconnaissance General Bureau — issued 35 fraudulent SWIFT instructions from Bangladesh Bank's account at the Federal Reserve Bank of New York, attempting to move $951 million. Of that, $81 million was moved into Rizal Commercial Banking Corporation in the Philippines, laundered through Manila casinos, and largely lost — only about $15 million was recovered. The operation exploited three weaknesses: a flat internal network at Bangladesh Bank that allowed lateral movement from a phishing foothold to the SWIFT terminal; custom malware that disabled the printers that would have printed confirmation receipts; and a carefully chosen three-time-zone weekend that coincided with Chinese New Year, delaying the manual verification that would have caught the fraud.

SWIFT's response is the best example in the paper of a sector-wide mechanism-design fix. The Customer Security Programme, launched in 2016 and mandatory from 2017, imposed a set of baseline controls on every participating institution. The logic was not that every institution would fail without the controls; it was that the *network* was only as strong as the weakest participant, and the system needed a floor below which no participant could drop. This is Varian's weakest-link model applied at sector scale, with a regulator-equivalent — SWIFT, a member-owned cooperative with coercive power over its members — enforcing the floor.

Then came the Equifax pattern, which is really a series of patterns: a missed patch, a blind spot in monitoring, a brittle response, and an enormous downstream cost. 147.9 million US records, another 15.2 million in the UK, and so on. The financial sector does not own Equifax's credit reporting role, but the pattern repeats in finance constantly. Capital One in 2019: 106 million records, a misconfigured AWS WAF, a former AWS employee's SSRF attack against the metadata service, $80 million OCC fine, $190 million class action settlement. The canonical cloud shared-responsibility failure.

MOVEit in May 2023 is the weakest-link problem rotated ninety degrees. Cl0p, the TA505-affiliated ransomware group, exploited CVE-2023-34362 in Progress Software's MOVEit Transfer managed file transfer product. Roughly 3,000 US and 8,000 global organisations were affected, including 1st Source, First National Bankers Bank, Putnam, TIAA, Prudential, and countless pension and payroll providers whose customers had never heard of MOVEit. No customer's MFA, Zero Trust, or EDR posture prevented compromise. The vulnerability was in a shared upstream product. This is sum-of-efforts, not weakest-link; the sector-wide defensive posture depended on every organisation's supply chain hygiene, and most organisations did not have visibility into their fourth-party risk. DORA's Critical Third-Party Provider regime, which came into force on 17 January 2025, exists specifically to address this class of problem.

Then there is ICBC Financial Services on 8 November 2023. LockBit 3.0, exploiting Citrix Bleed — CVE-2023-4966 — disrupted approximately $9 billion in Treasury-backed trades, pushed the day's failed-trade volume to $62.2 billion, and forced ICBC to settle trades manually via a USB stick carried by messenger in Manhattan. A single mid-sized broker-dealer, for a day, briefly dislocated the world's risk-free rate. This is systemic cyber risk as a macroprudential event, and it is precisely what Basel's operational resilience framework and the FSB's cyber incident reporting regime are designed to detect.

Finance's regulatory architecture is the most developed in the world. DORA imposes 24-hour initial incident reporting, 72-hour follow-up, Threat-Led Penetration Testing aligned to TIBER-EU, and Critical Third-Party Provider oversight by the European Supervisory Authorities. NYDFS Part 500's Second Amendment, effective from 1 November 2023, added Class A audits, April 15 CEO and CISO certification, 24-hour ransom payment notification, mandatory penetration testing from May 2025, universal MFA from November 2025, and board-level cyber literacy requirements. The SEC's Form 8-K Item 1.05 requires four-business-day disclosure of material cyber incidents. Basel, PCI-DSS v4.0, the UK FCA and PRA, the Monetary Authority of Singapore, the Hong Kong Monetary Authority — every major jurisdiction has a detailed cyber regulatory regime, and they are increasingly converging on common principles of operational resilience, third-party risk management, and incident reporting.

Finance also has FS-ISAC, which is the best counterexample to the view that information-sharing is a doomed prisoner's dilemma. With approximately 5,000 member institutions across 75 countries, representing roughly $100 trillion in assets, FS-ISAC demonstrates that under the right conditions — legal safe harbour under CISA 2015, trusted peer reputation, aggressive anonymisation, and a threat environment that disproportionately punishes free riders — sharing can become the equilibrium strategy rather than the free-rider strategy. The sector's willingness to share is, along with the regulatory apparatus, the reason finance is relatively well-defended despite being the most obviously attractive target in the world.

For the finance CISO, the game has a particular shape. Your adversaries are mostly economic: ransomware affiliates, financial-crime-focused state actors like Lazarus, initial access brokers selling to whoever pays, and insider threats. Your regulator is aggressive and numerous. Your peers are organised and, within limits, cooperative. Your systemic exposure is real — both as a target of sector-wide events and as a contributor to them.

Your defensive strategy should therefore be:

- **Commit early and visibly to the regulator.** This is the most powerful lever in finance. DORA, NYDFS, the SEC rule, and the Basel framework are not obstacles; they are your commitment devices. Map every control investment to a regulatory requirement. The CFO is not arguing with you; they are arguing with the ECB.
- **Participate actively in FS-ISAC and equivalent sector bodies.** The sharing game works in finance. Use it. Intelligence about a campaign aimed at your sector is worth more than ten additional tools. Your participation is also part of what keeps the sharing equilibrium stable.

- **Obsess over third parties.** The MOVEit lesson is that your defences do not protect you from your supply chain's defences. Know your critical third parties. Know their critical third parties. Require contractual security minimums. Require incident notification. Use cyber insurance as an enforcement mechanism — require vendors to carry it, and let their underwriters do the enforcement for you.
- **Design for systemic resilience, not just your own survival.** A finance CISO who thinks only about their own organisation misses the point. If your failure cascades, you have failed twice: your organisation and the system. Operational resilience — the ability to keep delivering critical services even under partial failure — is what the regulators are increasingly requiring and what the sector needs.

3.3 Defence — The Long Game

Defence is the sector where the frame is most stretched, and where the stakes are highest. The adversary is patient, state-funded, and uninterested in your insurance coverage. The time horizons are measured in decades, not quarters. The payoff functions are political, not financial. And the supply chain — the defence industrial base, some 300,000 contractors in the US alone — is both the primary target and the primary vulnerability.

Start with APT1. In February 2013, Mandiant published a report attributing hundreds of intrusions across at least 141 victim organisations to People's Liberation Army Unit 61398, headquartered in Pudong, Shanghai. In May 2014, the US Department of Justice indicted five named PLA officers for economic espionage against Westinghouse, US Steel, Alcoa, ATI, SolarWorld, and the United Steelworkers union. The indictment was largely symbolic — the officers were never extradited — but it established a new norm: cyber-enabled economic espionage attributable to a named state actor would be publicly attributed and legally consequenced, even without prosecution. This was an information-control move at the nation-state level.

Then APT10, also known as Stone Panda or menuPass — an MSS Tianjin operation with front companies including Huaying Haitai. Operation Cloud Hopper, documented by PwC and BAE Systems in 2017, compromised approximately a hundred downstream organisations via managed service providers, including Hewlett Packard Enterprise and IBM. The game-theoretic lesson: if the attacker cannot efficiently compromise every target, they compromise the shared service that serves every target. This is the defence-sector analogue of MOVEit, five years earlier.

The 2011 RSA SecurID breach is the watershed. In March 2011, RSA — then owned by EMC — was penetrated via an Adobe Flash zero-day delivered in a spear-phishing email with a Microsoft Excel attachment titled, in the most famous filename in cybersecurity history, "2011 Recruitment plan.xls". The attackers exfiltrated information related to RSA's SecurID two-factor authentication tokens. Two months later, that information was used in an attempted attack against Lockheed Martin, which was detected and contained by Lockheed's internal security team using an internally developed framework they called the Cyber Kill Chain. Northrop Grumman and L-3 were also targeted. EMC spent $66.3 million on the response and eventually replaced approximately 30,000 customers' tokens. The attack is the defence-sector supply-chain

case par excellence: compromise the security vendor to compromise the defence primes.

The OPM breach, discovered in April 2015 but with the first intrusion in March 2014 via compromised KeyPoint contractor credentials, exfiltrated approximately 22.1 million SF-86 and SF-85 records — the background investigation forms for US security clearances — including 5.6 million fingerprints. The attack was attributed to MSS Jiangsu State Security. The payoff was not dollars. The payoff was a database of every cleared American employee, their family members, their foreign contacts, their financial histories, and their biometrics — intelligence of a kind that will be useful for decades.

The SolarWinds Orion compromise, disclosed in December 2020, was APT29 — Russia's SVR. Some 18,000 customers downloaded the trojanised Orion update; approximately 100 were actively exploited, including nine US federal agencies. The technical sophistication was extraordinary; the strategic sophistication was that SVR had compromised a software vendor used by nearly every major enterprise and government agency, and had done so with an exquisitely disciplined operational security that kept them undetected until a FireEye analyst noticed an anomaly in their own environment.

Then Microsoft Exchange ProxyLogon in March 2021 — HAFNIUM, with four zero-days affecting approximately 250,000 servers globally and 30,000 US organisations. And Storm-0558 in May and June 2023 — forged Azure AD tokens using a stolen 2016 MSA consumer signing key, used to access 25 organisations including the US State Department, exfiltrating 60,000 State Department emails. The Cyber Safety Review Board's April 2024 report rejected Microsoft's crash-dump hypothesis for how the key was stolen and concluded that Microsoft's security culture was inadequate for a company of its importance. The cloud era has made the defence-sector supply-chain problem worse, not better.

And, of course, Stuxnet. Operation Olympic Games, attributed to the US and Israel, using four zero-days and a USB-based air-gap crossing to destroy approximately 1,000 IR-1 centrifuges at Natanz in 2010. It remains the canonical offensive case study, and its lesson for defensive CISOs is not complicated: air-gaps are only as strong as the humans who carry media across them.

For the defence-sector CISO, the game has specific characteristics. Your primary adversaries are state intelligence services — APT28 and 29 from Russia, APT10 and 41 from China, and the Lazarus family from North Korea. Your regulatory environment is the most prescriptive and punitive — CMMC 2.0 finalised in October and September 2024, effective from December 2024 and November 2025 respectively, with phased implementation through 2028; DFARS 252.204-7012; NIST SP 800-171 and 800-172; ITAR; the NISPOM. Your supply chain is enormous, heterogeneous, and persistently under-resourced — as of October 2025, only approximately 431 organisations had achieved CMMC Level 2 final certification out of roughly 80,000 DIB firms that will require it.

Your defensive strategy must therefore be:

- **Assume breach.** You will not keep APT29 out indefinitely if you are on their list. The game is not prevention; the game is detection speed, containment speed, and damage limitation. Mandiant's 2025 M-Trends report noted that 44% of zero-day exploitation in 2024 pivoted to

edge devices — VPNs, firewalls, gateways — precisely because endpoint EDR coverage has improved. The attackers move where you are not looking. You must look everywhere.

- **Segment ruthlessly.** Air gaps where possible. Strict segmentation where not. Defence-in-depth is not a buzzword; it is a function of assuming any single control will eventually fail and requiring the attacker to defeat multiple independent controls in sequence.
- **Supply chain is the attack surface.** The same principle as in finance, but more intense. Know your subcontractors. Know their subcontractors. Use CMMC not as a compliance checkbox but as a floor — require higher. Where possible, segregate CUI-bearing environments from everything else.
- **Counter-intelligence, not just cybersecurity.** Defence-sector attackers are patient spies, not criminals. They recruit insiders. They monitor your personnel moves. They watch your conference attendance. Your defensive posture must include traditional counterintelligence practices — access control, compartmentalisation, personnel security, exit procedures — in addition to technical controls.
- **Coalition defence.** Use JCDC, the DIB-ISAC, the NSA Cybersecurity Collaboration Center, and the Hunt Forward Operations programme. You are not defending alone, and you should not pretend you are. The sharing equilibrium in the defence sector is subsidised by government and reinforced by classified threat briefings to cleared personnel. Use both.

4 Playing the Game Every Day

The levers and the sectors give you a strategic frame. They do not tell you what to do on Tuesday. The rest of this paper is about Tuesday.

4.1 The Board Game

Your single most consequential ongoing strategic interaction is not with the attacker. It is with your board.

This will sound wrong to most technical CISOs. The attacker is the enemy. The board is the employer. Surely the board is on your side?

The board is not on your side and not against you. The board has its own payoff function, which is enterprise value, measured quarterly, against a reference class of peer organisations. Security affects enterprise value in both directions: too little security and a breach destroys value; too much security and operating costs destroy value. The board's rational strategy is to approve the level of security investment that maximises expected enterprise value, which is almost never the level you think is optimal.

Your job is not to persuade the board to adopt your utility function. Your job is to demonstrate that your proposals are maximising *their* utility function. This is a translation problem.

Speak economics. "We have a $50,000 proposed investment that reduces the probability of a $2 million loss scenario by 70%." This is a sentence a board can evaluate. "We need to implement CIS Control 6 sub-control 6.4" is not. The Factor Analysis of Information Risk — FAIR — framework, Jack Jones's quantitative approach to risk expressed in loss event frequency times loss magnitude, is the best-established vocabulary for this translation. You do not need to produce Monte Carlo simulations of every control decision. You do need to produce defensible expected-loss estimates for the big ones, and you need to present the board with a portfolio that looks like an investment portfolio — a collection of risk-reducing investments ranked by expected return.

Tell them about near-misses. The single highest-leverage information-control move a CISO makes is to report near-misses to the board truthfully. If you report only successful defences, the board concludes you are either wildly effective or unneeded; the first is flattering and the second is fatal. If you report the successful defences *and* the near-misses — the phish that would have worked but for the automated quarantine, the misconfiguration that was caught in staging but would have been catastrophic in production, the vendor disclosure that arrived 48 hours before the exploit went public — the board sees the operational risk they are actually running, and they fund you accordingly. This is not scaremongering; it is the opposite. It is replacing their intuitive model of cyber risk, which is based on headlines, with your informed model, based on your actual environment.

Make regulation your ally. When NYDFS Part 500 requires CEO and CISO certification, you have gained a powerful internal ally: the CEO, who now has personal legal exposure to your programme quality. When SEC Form 8-K Item 1.05 requires four-business-day disclosure, you have gained access to the audit committee, whose chair now has skin in the game on your incident response capability. When DORA requires board-level approval of your ICT risk management framework, the board is no longer approving your budget as a favour to you; they are approving their own compliance with a binding regulation. Use these levers.

Do not hide behind regulation. If your only justification for a control is a regulation, the board will eventually conclude that security is a compliance exercise, and they will delegate it to the compliance function. Regulation should be one of several justifications, not the only one. Always also have the economic case, the operational case, and the threat-landscape case.

Have a standing board metric. Pick three to five indicators that tell the true story of your programme — not all fifty. Report them every quarter, the same way, so the board can see trends. Candidates: mean time to detect and mean time to respond for critical-severity incidents; patch compliance on internet-exposed systems; percentage of privileged access requiring phishing-resistant MFA; third-party risk register coverage; tabletop exercise frequency and outcomes. The metrics matter less than the consistency.

Prepare for the incident before the incident. When — not if — you have a material incident, the board will be looking for two things: that you are in control, and that you told

them in advance this was a real risk. If both are true, you will be supported. If either is false, you will not. The investment in board education before incidents is what buys you the support you need during them.

4.2 The Daily Play — A CISO's Operational Rhythm

Strategy lives in daily practice. A strategic frame that does not change how you spend Tuesday afternoon is a strategic frame that does not matter. Here is what strategic play looks like in daily operations.

Every day, ask the PAPI-T questions of your highest-priority decision. Not every decision. The big one. The one where you are about to sign off on a control choice, a vendor selection, a policy change, or an incident response action. Five questions: who are the players? What actions are available? What are the payoffs for each player? What does each player know? What is the time structure? You will be surprised how often this five-minute exercise changes the answer.

Every week, review your lever balance. Are you playing too much information control and not enough commitment? Too much mechanism design and not enough signalling? Most CISOs over-rely on whichever lever they are most comfortable with. Forcing yourself to audit the balance produces a more robust game.

Every month, run a small red-team exercise. Not a full penetration test. A single, narrow, specific exercise: what would it take for an attacker to achieve *this one outcome* against *this one system*? A well-designed monthly exercise surfaces more issues than a quarterly full pen test, because the cycle time is short enough to adapt and the scope is narrow enough to actually fix what you find.

Every quarter, review your adversary model. Who is actually attacking your sector? What techniques are they using? What has changed in the last three months? This is not a strategy retreat. This is a 90-minute meeting with your threat intelligence function and your senior SOC staff. Update your model. Make it specific. If your model says "nation-state and cybercrime" and does not name specific groups or techniques, it is not a model; it is a poster.

Every year, redesign one mechanism. Not all of them. One. Pick the security-user-experience friction point that generates the most tickets or complaints, and redesign the mechanism so that the aligned choice is the easy choice. One mechanism per year, compounded over five years, is how you build a programme where security is the path of least resistance.

Continuously, invest in your people. This is obvious and therefore often skipped. Your SOC analysts' mean tenure is the single most underrated metric in your programme. An analyst who has been with you for three years has internalised your normal, your architecture, your incident history, and your organisation's politics in a way no external hire can. Pay them. Promote them. Give them training. Give them interesting work. The single biggest mistake

CISOs make with their teams is treating them as fungible — and the adversary, meanwhile, runs stable teams with multi-year tenure.

Never, ever, take credit for a clean week. The week was not clean because you did good work. The week was clean because of a thousand small failures that did not compound into a big one, plus some amount of luck. The moment you start taking credit for clean weeks is the moment you stop preparing for dirty ones.

4.3 The Coordination Problems You Did Not Know You Had

Some of the hardest problems in your programme are not between you and the attacker. They are between you and the players who are nominally on your side.

Your MSSP. You are paying them for effort. Effort is hard to observe. They have an incentive to appear to be doing a lot while actually doing the minimum. This is classic principal-agent moral hazard. Solve it by requiring outcome-based SLAs: mean time to detect by severity, percentage of incidents escalated within contractual windows, percentage of false positives versus true positives. Measure what you want, not what is easy. And cultivate the capability to spot-check their work — run a tabletop where you inject a synthetic incident and measure whether they catch it.

Your software vendors. They are externalising security costs onto you. Every unpatched CVE in their product is a cost they are passing to their customers. The answer is not to chase every CVE; the answer is to make vendor security posture part of your procurement decision, contractually bind them to disclosure and patch timelines, and — where you are a big enough customer — use that leverage to improve their upstream practices. Ross Anderson has been writing about this for 25 years and it remains the biggest unsolved mechanism-design problem in cybersecurity.

Your employees. They are not your adversaries. They are rational agents optimising for a different utility function — productivity, convenience, promotion, survival of the next quarter. When they bypass your controls, it is almost always because the controls make their jobs harder for unclear benefit. Your response is not more enforcement; it is better mechanism design. Cormac Herley's NSPW 2009 paper on user rationality is required reading. Your users are not stupid. Your mechanisms might be.

Your peers at other organisations. You are in a coalition with them against a shared adversary, whether you acknowledge it or not. The sharing equilibrium — you tell me about your incident, I tell you about mine, we both defend better — is fragile but worth cultivating. FS-ISAC works. The E-ISAC works. The Defense Industrial Base Cybersecurity Program works. JCDC works, imperfectly. Participation is not a cost; it is an investment with a compounding return. Free riding feels cheap in the short run and ruinously expensive when your peer's unreported incident turns out to be a campaign that reaches you three weeks later.

Your insurer. This is the newest player and, in some ways, the most consequential. Cyber insurance underwriting has become, de facto, a commercial regulator. The mandatory MFA, EDR, and backup requirements your underwriter imposes are enforceable in a way your internal policies are not, because failure to comply voids coverage. Use this. Make the underwriter your ally in conversations with the CFO: "If we do not implement this control, our cyber insurance premium will rise and our coverage limits will drop." The Merck v. ACE American litigation — settled in January 2024 after the Appellate Division of the New Jersey Superior Court affirmed coverage for NotPetya in May 2023 — reshaped the market by establishing that war exclusions as then-written did not cover nation-state cyber incidents, leading to the Lloyd's LMA Y5381 bulletin of August 2022 mandating state-sponsored cyber exclusions from 31 March 2023. Your coverage is narrower now. Your control requirements are stricter. This is both a cost and a mechanism-design tool.

4.4 On Losing

You will lose. Not every game, but some. The question is not how to prevent losing; it is how to lose in a way that preserves the ability to keep playing.

Norsk Hydro lost on 19 March 2019, when LockerGoga encrypted a significant portion of their IT environment. Here is what they did: they refused to pay the ransom. They called Microsoft's Detection and Response Team. They went fully transparent, holding daily press conferences, publishing status updates, and inviting scrutiny. They spent somewhere between $52 and $75 million on recovery. Their share price went up. Not down — up. The market rewarded their transparency and their resilience. This is what a good loss looks like.

Maersk lost on 27 June 2017, when NotPetya — almost certainly GRU-originated, aimed at Ukraine but spread far beyond — destroyed 4,000 servers, 45,000 PCs, and required rebuilding 2,500 applications. The recovery cost Maersk $250 to $300 million; NotPetya's total cost across all victims, including Merck and FedEx's TNT subsidiary, was approximately $10 billion. Maersk's recovery depended on a single Ghana domain controller that survived because it had been offline during a local power outage. Maersk rebuilt in ten days. They have since become one of the most open organisations in the industry about what happened and what they learned. This is also what a good loss looks like.

MGM lost on 11 September 2023, when Scattered Spider and ALPHV compromised them via a ten-minute social-engineering phone call to the IT helpdesk. They refused to pay. They lost over $100 million in Q3 alone. They operated manually for approximately ten days. Caesars Entertainment, attacked by the same group in roughly the same period, paid $15 million of a $30 million demand. MGM's refusal likely shaped the game — the group was arrested months later, in part because its business model depended on payment rates that MGM's stance eroded. This, too, is what a good loss looks like.

What do good losses have in common?

- *Preparation before the incident.* Incident response plans, tested in tabletops. Board communications frameworks, agreed in advance. Legal counsel identified. External IR retainers in

place. Crisis communications playbooks drafted. Regulatory notification processes mapped. Identity, backup, and recovery capabilities actually tested, not just described in documents.

- *Honesty during the incident.* Telling the board everything you know, as you know it. Telling regulators within their statutory windows. Telling customers when you must and often when you should even if you do not have to. Telling employees what is happening and what you need from them. The organisations that cover up are the organisations whose losses become catastrophes; see Sullivan, Uber, 2016.
- *Ruthless clarity about priorities during recovery.* Not everything can come back at once. Decide what matters. Make the hard calls. Do not restore systems before they are cleaned. Do not pay a ransom you cannot verify. Do not promise timelines you cannot meet.
- *Learning after the incident.* Not a blameless post-mortem that surfaces nothing. A genuine post-incident review that asks hard questions, documents honest answers, and produces specific, trackable improvements. The organisations that emerge stronger from incidents are the ones that use incidents as forcing functions for changes that should have happened earlier.

One final thing about losing: you will be blamed. Some of the blame will be fair. Some will not. The CISO who survives is not the one who avoids blame; it is the one who has prepared their documentation, their decisions, their communications, and their compliance posture in advance such that the blame that does arrive is proportionate to actual failures and survivable in the sense of professional continuity. Document your risk acceptance decisions. Get signatures. Keep records. When the lawyer asks you in depositions what you knew and when, you want to have a clear, written, dated answer.

5 The Long View

5.1 What the Game Does Over Time

The game is not stable. It evolves. What worked five years ago does not work now, and what works now will not work five years from now. A CISO who optimises only for the current configuration is fighting the last war.

Some trends are visible and will continue.

Zero-day exploitation has shifted to edge devices. Mandiant's M-Trends 2025 reported 44% of observed zero-day exploitation in 2024 targeted VPNs, firewalls, and gateway devices — up sharply from prior years. The reason is evolutionary: endpoint EDR coverage has improved, and the attackers have moved where defenders are looking less closely. If you have not inventoried, patched, and monitored your edge devices with the same rigour as your endpoints, you are defending the wrong terrain.

Ransomware has diversified. The pure encryption model has given way to double extortion — encrypt and threaten to leak — and triple extortion — encrypt, threaten to leak, and attack customers and partners — and increasingly to pure exfiltration extortion with no encryption at

all, the Cl0p model perfected on MOVEit. The shift reflects both improved backup practices, which reduce the leverage of pure encryption, and the economics of data-market resale. Expect further diversification.

State pre-positioning is now a stated strategy. Volt Typhoon was not an accident. It was the visible surface of a deliberate PRC doctrine of establishing footholds in adversary critical infrastructure for use in a future contingency. The US Cyber Command's Persistent Engagement doctrine — codified in the 2018 DoD Cyber Strategy and articulated by General Nakasone's April 2018 Command Vision — is the American analogue, operating on adversary networks rather than on domestic ones. The combined effect is that every critical infrastructure CISO is, whether they know it or not, operating in a contested domain where both offence and defence have given up on deterrence by retaliation in favour of continuous contact.

AI is the wildcard. Large language models are making phishing cheaper, vulnerability research faster, and social engineering more scalable. They are also making threat detection, log analysis, and SOC tier-one work more automatable. The net effect on the offence-defence balance is unclear and probably sector-dependent. What is clear is that any CISO who treats AI as a hype cycle they can ignore will be surprised in ways they do not want to be.

Regulation is tightening, not loosening. NIS2, DORA, SEC, NYDFS, CMMC, CIRCIA, and the equivalent frameworks in Australia, Japan, Singapore, the UK, and India are not going away. They are converging on a small number of core principles: mandatory reporting, personal liability, board accountability, third-party oversight, and operational resilience. CISOs should plan their programmes against this regulatory baseline, not against current enforcement practice.

Liability is individualising. Joe Sullivan's conviction, the Tim Brown SEC case, NYDFS's personal certification requirement, DORA's management accountability — these are not isolated developments. They are a coherent trend toward holding named individuals personally responsible for cybersecurity outcomes. The implication is that every CISO should, in the next twelve months if they have not already: confirm their D&O insurance coverage and its cyber-specific terms; establish clear written decision authority with their CEO; ensure their incident response framework has pre-approved disclosure paths; retain personal legal counsel with cyber experience; and document their risk acceptance decisions meticulously. IANS's 2025 compensation benchmark reported that D&O coverage for CISOs rose from approximately 40% to over 50% in a single year. If you are in the 50% without it, you are out of step with the market.

5.2 What This All Adds Up To

If you take nothing else from this paper, take this.

Cybersecurity is not a technical problem. It is a strategic problem with technical components. The technology is necessary and not sufficient. What makes a programme succeed or fail is the interaction between people with different goals, different information, different resources, and

different payoffs — and the structural question of whether those interactions produce defensible outcomes or not.

You are playing this game whether you know it or not. Playing it well is a discipline, not a talent. It requires the vocabulary of players and actions and payoffs; the three levers of information control, commitment, and mechanism design; and the humility to know that no single lever is sufficient and that the game will keep going regardless of what you do.

You cannot win in the sense of achieving security. You can play well, which means raising the attacker's cost above the attacker's benefit, aligning your internal players' incentives with security outcomes, preserving optionality for yourself while removing it for your adversary, and losing — when you lose — in ways that keep you playing tomorrow.

Security is a verb. It is something you do, continuously, not something you have. It is a game you are playing right now, whether you chose to or not. The question that separates the CISOs who last from the ones who do not is whether they are playing deliberately, or being played.

Play deliberately. The game will not wait.

A A Working Vocabulary

Action.
A move available to a player. The set of actions available to each player defines the game's strategic space.

Bayesian game.
A game in which at least one player has private information about their own type — motivation, capability, payoff. The other players hold probabilistic beliefs about that type. Real cyber is always a Bayesian game, because you rarely know who exactly is attacking you or why.

Commitment device.
A mechanism that removes a player's future optionality in a way that changes other players' behaviour. Regulations, contracts, automated responses, and public commitments are all commitment devices.

FAIR — Factor Analysis of Information Risk.
Jack Jones's quantitative risk framework expressing risk as loss event frequency multiplied by loss magnitude. The most widely adopted vocabulary for translating security risk into financial terms.

FlipIt.
Ron Rivest, Ari Juels, Alina Oprea, and Marten van Dijk's formal model of continuous-time stealth-takeover games. The canonical game-theoretic model of advanced persistent threats.

Information asymmetry.
A condition where one player knows something another does not. Nearly every aspect of security involves information asymmetry; strategic play consists, in part, of managing which asymmetries are in your favour and which are not.

Mechanism design.

The branch of game theory concerned with designing the rules of a game so that rational players, pursuing their own objectives, produce outcomes the designer wants. Much of regulation is mechanism design. So is good user experience in security.

Nash equilibrium.

A state in which no player can improve their payoff by unilaterally changing their strategy, holding others' strategies fixed. Useful as a concept; rarely the right solution concept for real cybersecurity, because multiple equilibria almost always exist and selection is unstable.

Payoff.

The value a player assigns to an outcome. Payoffs are not dollars; they are whatever the player cares about. Different players have different payoff functions.

Principal-agent problem.

A coordination problem in which one party — the principal — delegates to another — the agent — whose effort is hard to observe and whose interests may diverge. Your MSSP, your software vendors, and in some framings your employees all involve principal-agent dynamics.

Quantal Response Equilibrium — QRE.

Richard McKelvey and Thomas Palfrey's refinement of Nash equilibrium that assumes players make noisy, approximately-rational choices rather than perfectly-rational ones. More realistic than Nash for most real interactions.

Repeated game.

A game played multiple times in sequence, allowing for reputation, learning, and conditional strategies. Nearly all cyber is a repeated game.

Signalling game.

A game in which one player's action conveys information about their type to other players. Honeypots, incident disclosures, threat intelligence publications, and regulatory announcements are all signalling moves.

Stackelberg game.

A sequential game in which one player — the leader — commits to a strategy first, and the other — the follower — responds optimally. The defender-commits-first model underlies most deployed security games, including Milind Tambe's ARMOR and PROTECT systems.

Varian's taxonomy.

Hal Varian's classification of collective security problems into weakest-link — system security depends on the least-protected component — sum-of-efforts — system security is the sum of all participants' efforts — and best-shot — system security depends on the most-protected component. Different problem classes require different mechanisms.

B Further Reading

For the CISO who wishes to go deeper, the following is a curated list of works that have shaped this paper and that reward careful reading.

Foundational game theory applied to security. Milind Tambe, *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned* (Cambridge, 2011). Tansu Alpcan and Tamer Başar, *Network Security: A Decision and Game-Theoretic Approach* (Cambridge, 2010). Jeffrey Pawlick and Quanyan Zhu, *Game Theory for Cyber Deception* (Birkhäuser, 2021). Marten van Dijk, Ari Juels, Alina Oprea, and Ronald Rivest, "FlipIt: The Game of Stealthy Takeover," *Journal of Cryptology* 26(4), 2013.

Economics of security. Ross Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*, 3rd edition (Wiley, 2020), particularly chapters on economics and psychology of security. Anderson's foundational 2001 paper, "Why Information Security Is Hard — An Economic Perspective," remains essential. The annual WEIS — Workshop on the Economics of Information Security — proceedings are the definitive venue for the field.

Behavioural and bounded rationality. Colin Camerer, *Behavioral Game Theory: Experiments in Strategic Interaction* (Russell Sage / Princeton, 2003). Cormac Herley, "So Long, and No Thanks for the Externalities: The Rational Rejection of Security Advice by Users" (NSPW 2009).

Strategic context. Thomas Schelling, *The Strategy of Conflict* (Harvard, 1960) and *Arms and Influence* (Yale, 1966). Rebecca Slayton, "What Is the Cyber Offense-Defense Balance? Conceptions, Causes, and Assessment," *International Security* 41(3), Winter 2016/17. Ben Buchanan, *The Hacker and the State* (Harvard, 2020). Joseph Nye, "Deterrence and Dissuasion in Cyberspace," *International Security* 41(3), 2016/17.

Threat landscape. Andy Greenberg, *Sandworm* (Doubleday, 2019). Nicole Perlroth, *This Is How They Tell Me the World Ends* (Bloomsbury, 2021). Kim Zetter, *Countdown to Zero Day* (Crown, 2014). David Sanger, *The Perfect Weapon* (Crown, 2018).

Practitioner frameworks. Adam Shostack, *Threat Modeling: Designing for Security* (Wiley, 2014). Jack Jones and Jack Freund, *Measuring and Managing Information Risk: A FAIR Approach*, 2nd edition (Butterworth-Heinemann). Bonney, Hayslip, and Stamper, *CISO Desk Reference Guide*.

Commentary. Bruce Schneier's essays, collected at `schneier.com` and in several volumes. Dan Geer's talks, particularly the 2014 Black Hat keynote on cyber realpolitik and the 2003 report "CyberInsecurity: The Cost of Monopoly." Tyler Moore and Rainer Böhme's ongoing work on security economics.

This field manual does not end here. The game does not end. Keep playing.