

HILLSTRONG · THOUGHT LEADERSHIP

The Four Hats of Cyber-Risk Exposure

Where cybersecurity liability comes from for companies that build, service, and operate connected infrastructure.

Chuck Tommey, P.E., CISSP, GICSP

Co-Founder & Principal Consultant,
Hillstrong Group Security, Inc.

Contents

A four-chapter guide to the distinct sources of cyber-risk exposure, and how one program can cover them.

INTRODUCTION One Vulnerability, Four Liabilities

CHAPTER 1 The Enterprise Hat

CHAPTER 2 The Product Hat

CHAPTER 3 The Operator Hat

CHAPTER 4 The Critical Infrastructure Hat

CONCLUSION Scope to the Convergence

INTRODUCTION

One Vulnerability, Four Liabilities

THE FOUR HATS OF CYBER-RISK EXPOSURE

A company that builds connected hardware, services those products under contract, and may also operate them inside critical infrastructure does not have one cybersecurity problem. It has four, and they can originate from the product, remote access, customer SLAs, and/or their own enterprise.

I wrote this after one too many assessment engagements that quickly spiraled into existential risk management exercises. This happens when a capable security team, a real budget, and a mature program meet an operational contract and/or regulatory blind spot the size of the business. The program covered the corporate network well. It had almost nothing to say about the firmware in the field, the service platform reaching into customer operations, or the reporting clock that starts when a grid-connected asset goes down. Each of those is a distinct source of cyber-risk exposure with its own regulators, its own threat actors, and its own failure mode. A program scoped to one of them often leaves one or more of the other three uncovered, and the company does not find out which one was exposed until an incident identifies it.

The hat you skip is the one an adversary will choose.

The model in this book separates those exposures into four hats which cannot be juggled but must be worn simultaneously. A bit absurd, but memorable when you think about how a manufacturer that also services and operates its products accepts all four risk vectors at once:

- The Enterprise. The company as a public business, carrying reputation, corporate IP, and a new wave of board-level governance and disclosure duties.
- The Product. The connected device itself, carrying secure-by-design market expectations and, under the EU Cyber Resilience Act, fleet-wide vulnerability and support obligations.
- The Operator. The service relationship, carrying the privileged remote access that can turn one enterprise compromise into a many-customer breach, plus the contractual and flow-down obligations that come with servicing or operating someone else's equipment.

- **Critical Infrastructure.** The asset's role in the grid or another essential service, carrying sector regulators, nation-state adversaries, and the fastest reporting clocks in cyber.

The reason this matters is risk convergence. The product you design and sell is the asset you service and operate, and it is a node on the grid. One firmware vulnerability can land as a product-liability event under the Cyber Resilience Act, a contractual breach under a customer's service agreement, and a grid-reliability incident with a one-hour reporting clock, with additional board and C-level implications, all from the same root cause. A company that treats those as separate problems handled by separate teams will respond to each one late or incomplete, because no single team owns the whole picture.

This book is for the people who own that picture: the CISO, the head of engineering or operations, the general counsel, and the board members who now carry personal accountability for cyber under regimes like NIS2. It is written for companies in energy, water, and the other sectors where connected products have moved into essential operations and the regulatory environment has moved with them.

The four middle chapters take one hat each. Each chapter names where the exposure comes from, what triggers it, what the current rules require, with the citations to back it, how it fails, and what to do about it. The regulatory references are current as of mid-2026. Reporting clocks and effective dates change, so treat the citations as a starting point and confirm the specifics against primary sources before you rely on them in a filing or a contract.

The closing chapter returns to convergence and answers the question the four hats raise: if the exposures overlap on the same systems, can one program cover all four? It can. The final chapter shows how to scope it, and gives you a short self-assessment to find which hat your current program has left exposed.

Start wherever your risk is sharpest. If you are a public company under new disclosure rules, start with the enterprise hat. If your product just came into CRA scope, start with the product. But read all four before you decide your program is complete, because the hat you skip is the one an adversary will choose.

CHAPTER 1

The Enterprise Hat: When Cyber Risk Reaches the Boardroom

THE FOUR HATS OF CYBER-RISK EXPOSURE



Most security programs start where the budget already lives: the corporate network, the email tenant, the endpoints, the SOC. That is the enterprise attack surface, and it matters. It is also the hat that companies understand best, which is exactly why it gets treated as the whole picture. For a manufacturer that ships connected hardware and services it in the field, the enterprise hat is one of four. This chapter covers what that hat carries, because the answer changed over the last three years and most boards have not caught up.

Where the exposure comes from

The enterprise hat is the company as a legal and financial entity. The exposure here has three sources. The first is reputation and market trust, the softest to measure and the fastest to lose. The second is the corporate IT and intellectual-property surface: the engineering designs, the controls firmware, the pricing, the merger data, all the material a competitor or a nation-state actor would pay to steal. The third source is the one that grew teeth recently, regulation aimed at the company's governance rather than its technology.

Cyber moved from a control the security team implements to a duty the board discharges.

That third source is where most leadership teams underestimate their position. Cyber regulation used to read like an IT problem: patch this, log that, report incidents on this timeline. The current wave of rules points somewhere else. It points at the board.

What the rules now require

Three regimes show the shift clearly.

In the United States, the SEC requires a public company to file a Form 8-K under Item 1.05 disclosing a material cybersecurity incident within four business days of deciding the incident is material. The clock does not start at discovery. It starts at the materiality determination, which puts the pressure on how fast and how defensibly you can make that judgment. Regulation S-K Item 106 adds an annual disclosure of how the company manages cyber risk and how the board oversees it. These rules took effect at the end of 2023, and the first full year of filings showed companies struggling with the same question: who decides materiality, and on what evidence.

The European Union's NIS2 Directive goes further into the boardroom than any prior cyber rule. Under Article 20, the management body has to approve the company's cyber risk-management measures and oversee their implementation. Directors also have to complete cybersecurity training. This is not a delegation to the CISO with a dotted line to the audit committee. The directive names the management body as the accountable party. Article 32 lets authorities, for essential entities, seek a temporary ban preventing a named senior individual from holding management functions, and that ban can follow the person across other essential entities, not only the one where the failure happened. Article 34 sets the fines: for essential entities, up to at least ten million euros or two percent of total worldwide annual turnover, whichever is higher, and for important entities, seven million or 1.4 percent. Turnover means the consolidated group's global revenue, even when a single EU operation triggered the breach.

Australia's Security of Critical Infrastructure Act reaches the board through a different door. A responsible entity with a governing body has to have that body approve the annual Critical Infrastructure Risk Management Program report before it goes to the regulator, within ninety days of the financial-year end. A board that fails to approve or oversee that program can find itself exposed under the general director duty-of-care provisions of the Corporations Act. The financial penalties are fixed civil amounts rather than turnover-based, so they land lighter than NIS2, but the accountability structure is the same shape: the board signs, the board owns it.

Read those three together and a pattern shows up. Cyber moved from a control the security team implements to a duty the board discharges. Personal liability, mandatory director training, and turnover-scale fines are governance instruments, not technical ones.

The failure mode

The failure mode on the enterprise hat is not a breach. A breach is the trigger. The failure is what happens after, when the company cannot show that its board understood the risk, approved a real program, and oversaw it. That gap turns a security incident into a securities-disclosure problem, a regulatory-enforcement problem, and in the NIS2 case a personal problem for named executives.

The hats also compound. A material NIS2 fine is the kind of event that can itself meet the SEC's materiality bar and trigger an 8-K. One operational failure in one country can become a fine, a disclosure obligation, and a management ban in the same quarter. The enterprise hat is where those threads land, because it is the hat that speaks for the whole company.

What to do about it

Three moves separate the companies that carry this hat well from the ones that get surprised.

Fix the materiality decision before you need it. Write down who makes the call, what evidence they use, and how fast the process runs, then rehearse it. The four-business-day clock is survivable only if the determination process is already built.

Put cyber governance in the board's own record. Approvals, oversight, and director training belong in board minutes, not in a slide the CISO presents once a year. NIS2 and SOCI both assume the governing body can produce evidence that it did its job. If it is not written down, it did not happen.

Map your entity classification early. Whether you are an essential or an important entity under NIS2, and whether you are a responsible entity under SOCI, sets your fine ceiling, your supervisory regime, and your personal-liability exposure. For a multinational, a non-EU headquarters does not remove NIS2, because the directive reaches the group through its EU operations and scales the fine on worldwide turnover. That determination is a legal question worth answering before an incident forces the issue.

The enterprise hat is the one every company already wears. The mistake is assuming it is the only one. The next chapter adds a harder hat on top of this one: the product itself, the connected device you shipped, and the liability that rides with every unit in the field.

In the next chapter, The Product Hat: secure-by-design, the update channel, and fleet-wide liability under the EU Cyber Resilience Act.

CHAPTER 2

The Product Hat: Every Unit You Shipped Is Still Your Problem

THE FOUR HATS OF CYBER-RISK EXPOSURE



The enterprise hat covers the company. The product hat covers the thing the company sells, and for a manufacturer of connected hardware that is a different and harder exposure. A software vendor ships a bug and pushes a patch. A hardware manufacturer ships a controller with updatable firmware into a substation, a factory, or a battery site, and that unit runs for a decade or more under someone else's roof. The vulnerability you shipped in year one is still yours in year eight. This chapter is about that hat.

Where the exposure comes from

Two things make the product hat distinct from ordinary software liability.

The vulnerability you shipped in year one is still yours in year eight.

The first is that the product is a cyber-physical system. When it fails, the consequence is not only lost data. A manipulated battery management system can drive a cell into thermal runaway. A spoofed setpoint can trip a process. The failure crosses from the digital domain into the physical one, which pulls product-liability law and safety regulation into a conversation that most software teams have never had to hold.

The second is the update channel. The same remote-update mechanism that lets you patch a fleet is a single point of control over every unit that trusts it. Compromise the channel and you do not have one vulnerable device, you have the entire installed base. The feature that makes a connected product supportable is also the feature that makes it a fleet-wide risk. That is the tradeoff at the center of the product hat, and it does not go away with better patch hygiene. It has to be designed for.

What the rules now require

The European Union's Cyber Resilience Act, Regulation (EU) 2024/2847, is the first horizontal law that treats a connected product's security as a condition of sale. It changes three things for a manufacturer.

It creates a reporting duty on the product itself. Under Article 14, when a manufacturer learns of an actively exploited vulnerability or a severe incident affecting one of its products, it has to file an early warning within twenty-four hours, a fuller notification within seventy-two hours, and a final report within fourteen days of a corrective measure being available. Reports go to the relevant national CSIRT and to ENISA through the CRA's Single Reporting Platform. The clocks resemble the incident-reporting regimes on the other hats, but the trigger is different. It fires on a vulnerability in your product, discovered anywhere in the world, exploited in anyone's deployment.

It requires a software bill of materials. The manufacturer has to produce an SBOM in a common machine-readable format, covering at least the product's top-level dependencies, held in the technical documentation and available to market-surveillance authorities on request. The CRA does not force you to publish it, but it does force you to know what is inside your own product, which many manufacturers cannot currently answer for firmware assembled from years of third-party components.

It sets a support obligation. Security updates have to be provided and vulnerabilities handled across the product's expected lifetime, with a support period generally of five years. That converts every shipped unit into a standing commitment. You cannot sell the device and walk away, and you cannot quietly end support because the product line was discontinued.

The timeline is close. The CRA's reporting obligations apply from 11 September 2026, and the main obligations apply from 11 December 2027. A manufacturer that has not built the vulnerability-reporting process and the SBOM discipline by then is selling into the EU market out of compliance.

The failure mode

The failure mode on the product hat scales with the size of the fleet. A single defensible vulnerability in a single product, exploited across the installed base, becomes a fleet-wide liability event, a reporting obligation on a twenty-four-hour clock, and a market-access problem if the CRA conformity case falls apart. For a manufacturer whose products sit inside customers' critical operations, that same defect also lands on the other hats, which later chapters cover. The point here is narrower. The product you shipped is a liability you still hold, and the law now measures how well you hold it.

What to do about it

Three moves matter most.

Build the SBOM before the regulator asks for it. Start with the products still in production and work backward through the fielded base. You cannot report on, or patch, dependencies you cannot see. The SBOM is the foundation the rest of the product-security case stands on.

Treat the update channel as the crown jewel. Sign your firmware, protect the signing keys as carefully as you protect anything in the company, and design the update path so a compromise of one customer's deployment cannot reach the fleet. The blast radius of the update mechanism is the single largest technical risk on this hat.

Stand up the twenty-four-hour vulnerability-reporting process now, not in 2026. The CRA clock is short, and the reporting decision, like the SEC materiality decision on the enterprise hat, is only survivable if the process already exists. Know who receives a vulnerability report, who decides it is actively exploited, and who files to the national CSIRT and ENISA.

The product hat is where a manufacturer's engineering choices become legal exposure. But a company that also services what it sells wears a third hat, and it is the one that turns a single incident into a many-customer event. That is the subject of Chapter 3: the operator.

In the next chapter, The Operator Hat: maintenance, monitoring, privileged remote access, and the blast radius of a service platform.

CHAPTER 3

The Operator Hat: The Service You Sell Is the Attack Path You Own

THE FOUR HATS OF CYBER-RISK EXPOSURE



Selling the product is rarely the end of the relationship. Customers want the manufacturer to maintain the equipment, patch it on a schedule, respond when it breaks, and increasingly to watch it around the clock. That service business is good revenue and a strong moat. It is also a third hat, and it carries a risk that the enterprise and product hats do not: to deliver the service, the manufacturer holds standing privileged access into other companies' operations. This chapter is about what that access does to your risk position.

Where the exposure comes from

The operator hat has three layers, and they stack.

Compromise the service platform, or the remote-access channel, and the attacker does not reach one customer.

The first layer is preventive maintenance. Time-based work, coordinated to scheduled downtime: feature updates, operating-system upgrades, bug and security fixes. Routine, planned, low-drama.

The second layer is break/fix. Ticketed requests from the customer, plus alerts surfaced by the manufacturer's own monitoring. This is where the service organization starts reaching into the customer environment on demand, often with elevated rights, often quickly.

The third layer is the one growing fastest: continuous monitoring and response, the manufacturer acting like a managed security service provider for its own installed base. Watching the fleet, triaging alerts, responding to incidents inside the customer's operational network.

Each layer deepens one thing: the privileged remote-access pathway from the manufacturer into the customer. That pathway is the actual exposure. To service many customers efficiently, the service platform reaches into many customer environments, frequently through shared infrastructure and standing credentials. The convenience that makes the service scalable is the same property that makes it dangerous.

The blast radius

Here is the sentence that should keep a service-business leader up at night. Compromise the service platform, or the remote-access channel, and the attacker does not reach one customer. They reach every customer that platform touches.

A single-tenant breach is a bad day. A service-platform breach is a many-customer incident with lateral movement between tenants, and it turns the manufacturer into the delivery mechanism for attacks on its own customers. The 2020 and 2021 supply-chain intrusions that moved through trusted software-update and management channels showed the pattern at scale. A service platform with privileged reach into customer operations is the same pattern in a different wrapper.

The exposure runs the other direction too. Because the manufacturer operates the customer's equipment, the customer's regulatory obligations flow back up the contract. If the customer is an essential entity under NIS2 or a responsible entity under a critical-infrastructure regime, the security duties the customer owes get pushed onto the supplier that runs the gear. NIS2 makes this explicit: Article 21 requires entities to manage supply-chain security, including the security of their direct suppliers and service providers. The customer will translate that obligation into contract terms, and those terms land on you.

What the rules and contracts require

The operator hat is governed less by a single statute and more by the contract, which is where the sharpest teeth usually sit.

Service-level agreements (SLAs) carry penalties for missed availability and missed response times, and a security incident that takes the service offline can trigger them directly. Customer audit rights let the customer, or an independent assessor, inspect the manufacturer's security posture, which turns your internal control gaps into someone else's findings. Indemnification clauses assign the cost of a breach, and a manufacturer that caused a customer incident through its own service platform is the party those clauses point at. Flow-down provisions push the customer's own regulatory duties onto the supplier, so the manufacturer inherits obligations it did not sign up for directly.

NIS2 Article 21 sits behind all of this as the regulatory driver that makes customers demand it. When the customer has to manage its supplier risk by law, the manufacturer's security posture stops being a private matter and becomes a contract condition, an audit item, and a renewal risk.

The failure mode

The failure mode on the operator hat is the multi-customer breach. One compromised access path, many affected tenants, penalty clauses firing across the customer base at once, audit findings the manufacturer has to remediate under deadline, and the reputational damage of having been the vector. In the worst version, a customer with its own regulatory reporting obligation reports an incident that originated in the manufacturer's service platform, and the manufacturer learns about its own breach from its customer's regulator.

What to do about it

Three moves reduce this exposure the most.

Segment the service platform by tenant. The default architecture of many service organizations is flat: one platform, shared credentials, broad reach. Redesign it so a compromise of one customer's access cannot pivot to another's. Per-tenant isolation is the single most valuable control on this hat, and it is usually the one that was skipped for speed early in the service business.

Put multi-factor authentication and just-in-time access on every privileged path into a customer. Standing always-on credentials into customer operations are the liability. Short-lived, strongly authenticated, logged access that expires when the work is done shrinks the window an attacker can use and the credential set they can steal.

Read your own service contracts as an attacker would read your defenses. Know which SLAs a security incident triggers, what your audit exposure is, and which of your customers'

regulatory obligations have flowed down to you. The contract is where the operator-hat risk is priced, and most service organizations have never mapped it.

The operator hat turns the manufacturer into part of its customers' attack surface. The fourth hat goes one level higher, to the moment the equipment is not just a customer's asset but a piece of a nation's critical infrastructure. That is Chapter 4.

In the next chapter, The Critical Infrastructure Hat: grid stability, sector regulators, nation-state threat actors, and the fastest reporting clocks in cyber.

CHAPTER 4

The Critical Infrastructure Hat: When Your Asset Belongs to the Grid

THE FOUR HATS OF CYBER-RISK EXPOSURE



The first three hats attach to the company, the product, and the service relationship. The fourth attaches to something the manufacturer does not fully control: the role the equipment plays in a nation's essential systems. A battery site that stabilizes a regional grid, a controller inside a water treatment plant, a substation automation package. Once the asset matters to supply, a new set of regulators and a different class of adversary show up, and they do not care about the corporate org chart. This chapter is about the hat that gets assigned to you by what your equipment does, not by who you are.

Where the exposure comes from

The critical-infrastructure hat has a different origin from the others. The enterprise hat comes from being a company. The product hat comes from selling a device. This hat comes from the asset's function in the grid or in another essential service. That function draws two forces at once.

They are not trying to monetize access, they are trying to be able to turn something off.

The first is a set of sector regulators whose mandate is continuity of supply, not corporate compliance. They regulate the reliability of the electricity system, the safety of the water, the availability of the network. Their obligations attach to the operational asset and the entity responsible for it, and they carry the fastest reporting clocks in all of cyber.

The second is the threat actors who target critical infrastructure specifically. The enterprise and product hats mostly attract financially motivated actors and intellectual-property thieves. The critical-infrastructure hat attracts nation-state operators and hacktivists whose objective is disruption of supply, not data or money. Their intent changes the threat model. They are not trying to monetize access, they are trying to be able to turn something off.

What the rules require

Three regimes show how fast and how seriously this hat is regulated.

In the United States, NERC CIP-008-6 governs incident reporting for the bulk electric system. A responsible entity has to report a Reportable Cyber Security Incident, and attempts to compromise the systems that protect the electronic security perimeter, to the Electricity Information Sharing and Analysis Center and to CISA within one hour of determining that the incident qualifies. One hour. That is the tightest clock in this book, and it exists because the thing being protected is the reliability of the grid itself.

In the European Union, NIS2 places energy squarely in the essential-entity category, which means the Article 23 reporting regime applies in full: an early warning within twenty-four hours of becoming aware of a significant incident, a fuller notification within seventy-two hours, and a final report within one month. Energy operators do not get to argue their way out of essential-entity status, and the supervisory regime for essential entities is proactive, with audits and inspections rather than after-the-fact review.

In Australia, the Security of Critical Infrastructure Act sets a two-tier clock: a cyber incident with a significant impact on a critical-infrastructure asset has to be reported to the Australian Cyber Security Centre within twelve hours, and one with a relevant impact within seventy-two hours. The same Act layers on a risk-management program obligation and, for the most important assets, enhanced obligations.

Put the clocks side by side and the message is unmistakable. Twenty-four hours under NIS2, twelve under SOCI, one under NERC CIP. As the asset's importance to supply rises, the time you are given to report falls. The regulator is buying situational awareness of the grid, and it is buying it fast.

The failure mode

The failure mode on this hat is not measured in euros or dollars first. It is measured in supply. A disruption to a grid-stabilizing asset is a reliability event before it is a compliance event, and the regulatory enforcement, the license risk, and the fines follow the physical consequence. For a manufacturer or operator, the reputational and regulatory damage of being the cause of a supply disruption outweighs the direct penalty. Missing a one-hour or twelve-hour reporting clock on top of the disruption compounds it, because the regulator now has two failures to act on: the incident, and the silence.

What to do about it

Three moves matter most on the critical-infrastructure hat.

Build the reporting capability to the tightest clock you face, not the loosest. If any asset in your fleet sits under a one-hour or twelve-hour regime, the detection, triage, and notification pipeline has to run at that speed. A process built for a seventy-two-hour clock will miss a one-hour one every time. Rehearse it, because a reporting process that has never been exercised will not perform on the day it matters.

Segment the operational network so a compromise cannot reach the control function. The nation-state actor on this hat is aiming at the ability to disrupt supply. Defense in depth between the enterprise, the service-access paths, and the control system is what denies that objective. This is where the operator hat and the critical-infrastructure hat meet: the same privileged remote-access path that services the asset is the path an adversary would use to reach the control function.

Know your obligations per jurisdiction and per asset. A fleet that spans the US bulk electric system, EU energy, and Australian critical infrastructure carries three reporting regimes with three clocks. Map which asset falls under which, and who is the responsible entity for each, before an incident forces the question under time pressure.

Four hats, one company, and they all sit on the same equipment. The final chapter is about the thing that makes this hard and also makes it solvable: the four hats converge on the same systems, which means one well-designed program can address all four, and a program scoped to only one will leave three exposed. That is where we go next.

In the conclusion: the convergence point, and how to build one program that covers all four hats.

CONCLUSION

Scope to the Convergence

THE FOUR HATS OF CYBER-RISK EXPOSURE

The four hats look like four programs. They are not. They are four views of the same systems, and that is the insight the whole model turns on.

The product you sell is the asset you service is the node on the grid. The firmware is one artifact. The remote-access path that services the fleet is the same path an adversary would use to reach a control function. The board that answers to the SEC is the same board that answers to NIS2. When the exposures share the underlying systems, a defect in one place radiates into all four hats at once. That is the risk. It is also the opportunity, because a control placed at the convergence point protects against several hats together.

The product you sell is the asset you service is the node on the grid.

Two forces run through every hat

Before the program design, two cross-cutting realities have to be named, because they refuse to stay inside a single hat.

The first is supply chain. The company inherits the vulnerabilities of the components it buys, and it is itself a supply-chain link for every customer it sells to. Risk flows upstream and downstream at the same time. The Cyber Resilience Act's SBOM requirement, NIS2's Article 21 supplier-security duty, and the flow-down clauses in customer contracts are three expressions of the same fact: no company in this market is assessed on its own perimeter alone anymore.

The second is that the threat actors differ by hat, and the defense has to differ with them. Financially motivated actors and IP thieves come for the enterprise and the product. Nation-state operators and hacktivists come for the grid. A program tuned only for ransomware will be calibrated wrong for an adversary whose goal is to disrupt supply and who is willing to wait. Knowing which actor wants which hat tells you where to spend.

One program, four hats

A security program that covers all four hats shares a spine and branches at the edges. The spine is the set of controls that pay off across every hat, and the convergence point is where you build it.

Identity and access is the clearest example. Strong authentication and just-in-time, least-privilege access protect the enterprise hat against account takeover, the operator hat against the multi-customer blast radius, and the critical-infrastructure hat against the path an adversary would take to the control function. One control, three hats. That is what scoping to the convergence means in practice.

Asset and software inventory is another. You cannot report a vulnerability under the Cyber Resilience Act, satisfy an SBOM request, segment a service platform, or defend a control network without knowing what you have. The inventory serves the product, operator, and critical-infrastructure hats at once.

Incident detection and reporting is the third. Build the pipeline to the tightest clock any of your assets face, which for a grid operator is the one-hour NERC CIP clock, and the same pipeline satisfies the slower NIS2, SOCI, SEC, and CRA obligations underneath it. A process built to the loosest clock will fail the tightest one. A process built to the tightest clock covers them all.

Around that spine, each hat gets what only it needs: board governance evidence for the enterprise hat, secure-development and update-channel protection for the product hat, tenant segmentation for the operator hat, and control-network defense in depth for the critical-infrastructure hat.

A four-hats self-assessment

Run your current program against these questions. The hat where you cannot answer cleanly is the hat you have left exposed.

Enterprise. Can you name who decides materiality for a cyber incident, and produce board minutes showing the board approved and oversaw the cyber program? Do you know whether you are an essential or important entity under NIS2?

Product. Can you produce an SBOM for your top-selling connected product today? Do you have a process that files an actively-exploited-vulnerability report within twenty-four hours? Is your firmware update channel protected against a compromise that would reach the whole fleet?

Operator. If one customer's service access were compromised, could the attacker reach a second customer? Do your privileged paths into customer environments use strong authentication and expire when the work is done? Do you know which customer regulatory obligations have flowed down to you by contract?

Critical infrastructure. Do you know which of your assets fall under a one-hour, twelve-hour, or twenty-four-hour reporting clock? Have you rehearsed the fastest one? Can an adversary who lands on your enterprise or service network reach a control function, or does segmentation stop them?

If those questions are uncomfortable, that is the point. Most programs answer the enterprise questions well and the other three poorly, which is exactly the imbalance that leaves a well-defended company exposed on the systems that matter most.

Where to start

Start with the spine. Identity, inventory, and incident reporting return value across all four hats, so they are the first place to invest. Then close the single weakest hat your self-assessment exposed, because that is the one an adversary will find and exploit. The goal is not four programs. It is one program, deliberately scoped to the point where the four hats meet, so that a single vulnerability stops being four liabilities.

Hillstrong Group Security, Inc. helps companies in energy, manufacturing, and other critical sectors build OT and IoT security programs that hold up across all four hats. The assessment starts on the plant floor, not on a remote scan.

About the Author



Chuck Tommey, P.E., CISSP, GICSP

Co-Founder & Principal Consultant, Hillstrong Group Security, Inc.

Chuck Tommey is Co-Founder and Principal Consultant at Hillstrong Group Security, Inc., an operational-technology cybersecurity firm serving industrial and critical-infrastructure operators. He brings more than 30 years of experience in operational technology, working across engineering, information security, and OT to help asset owners and connected-product manufacturers build security programs that hold up against their regulatory, product, service, and grid obligations.

His credentials span those disciplines: a licensed Professional Engineer (P.E.), a Certified Information Systems Security Professional (CISSP), and a GIAC Global Industrial Cyber Security Professional (GICSP). That combination of engineering, information security, and industrial-control-systems practice is the same convergence this book argues security programs now have to cover.

chuck.tommey@hillstrongsecurity.com · hillstrongsecurity.com



About Hillstrong Group Security, Inc.

Hillstrong Group Security, Inc. is an operational-technology cybersecurity consultancy helping industrial and critical-infrastructure operators build security programs that hold up across their regulatory, product, service, and grid obligations. Blending OT, IoT, and ICS expertise, Hillstrong partners with companies in energy, manufacturing, water, chemicals, food and beverage, automotive, and other essential sectors to achieve operational resilience. Every engagement starts on the plant floor, hands-on and on-site, not on a remote scan.

Our Core Services

- OT Security Maturity Assessments
- OT Risk & Vulnerability Analysis
- OT Program Design & Governance
- Expert-Level OT Security Engineering
- Fractional CISO & Advisory Support
- OT Governance, Risk & Compliance (GRC) as a Service
- OT Staff Augmentation & Contract Specialists
- Resilion.io, AI Execution Engine for OT Cybersecurity

Resilion.io is Hillstrong's AI-powered OT cybersecurity platform, delivering continuous visibility, risk quantification, and program execution (build, execute, manage), so assessment findings translate into measurable resilience.

www.hillstrongsecurity.com · Based in Atlanta, GA · Solving cybersecurity challenges globally