HILLSTRONG GROUP SECURITY

HILLSTRONG · THOUGHT LEADERSHIP

Operationalizing OT Vulnerability Management

How Manufacturing Leaders Can Prioritize Cyber Risk Without Relying on Patching Alone

Roger Hill | CEO, Hillstrong Group Security, Inc.

Contents

CHAPTER ONE	Why Patching is Not a Panacea in OT Security
CHAPTER TWO	How to Recalculate CVSS Scores in an OT Context
CHAPTER THREE	The RISE Model: Prioritizing OT Vulnerabilities Where It Matters Most
CHAPTER FOUR	Mapping OT Business Risk Using the QUICK Framework
CHAPTER FIVE	How to Build a Sustainable OT Vulnerability Management Program
	About the Author
	About Hillstrong Group Security, Inc.

HILLSTRONG GROUP SECURITY

OT VULNERABILITY MANAGEMENT · CHAPTER ONE

Why Patching is Not a Panacea in OT Security

Debunking the Dangerous Illusion of Patch-Centric Protection

The notion that patching is the foundation of all good cybersecurity hygiene has become almost doctrine. It's a logical and practical approach in the IT world: patches are regularly issued, environments are flexible, and downtime can be managed with limited disruption. But the industrial world doesn't operate with the same assumptions. When this doctrine is lifted and dropped into operational technology (OT) environments without translation, it sets security programs up to fail. Not because patches are inherently evil but because their role in OT is fundamentally different.

I've seen this pattern repeatedly over the years: CISOs and CIOs with enterprise IT backgrounds walk into manufacturing and critical infrastructure environments and assume that patching is a first-order defense strategy. And it's understandable. It's what they know. But that assumption quickly collides with reality: fragile legacy systems, certification requirements, tightly coupled production lines, and unforgiving uptime expectations. All of this demands a recalibrated mindset around risk, and the blunt instrument of patching no longer makes sense as the default answer.

The Difference Between Can't and Won't

When a system isn't patched in IT, it's often seen as negligence or poor hygiene. In OT, unpatched systems are usually the result of informed, risk-based decisions. You don't patch a SCADA system controlling a multi-million-dollar extrusion line without understanding what a reboot could do. You don't update PLC firmware mid-shift without accounting for how it affects a plant's output targets or how it might violate FDA validation.

I've sat across from plant operations teams that have been blamed for being “obstructionists” when they push back on patching requests. In truth, they're often the most risk-aware people in the room: they're just evaluating risk through physical safety, production loss, and regulatory exposure.

This is why we must stop treating patching as a virtue in and of itself. It is a tool (one of many) and sometimes, it's wrong.

Why the Patch-Centric Model Breaks in OT

Let's start with cadence. Patching in IT environments can happen monthly, weekly, or even daily. Vulnerability management systems and patch orchestration tools are designed to work in continuous cycles. OT environments (especially manufacturing and energy) don't operate on those same clocks.

Patch windows in OT can range from quarterly to annually. In some cases, especially with embedded devices or legacy control systems, patching may not even be technically possible without replacing the

device altogether. Even when patches are available, applying them often requires shutting down operations, validating that the update won't disrupt custom application code, and, in some cases, re-certifying the system altogether.

This makes patching in OT not just a security issue. It's a business risk decision. That's a different paradigm. Security leaders must understand that the “availability” in the CIA triad is not just about uptime but revenue continuity, worker safety, and operational trust.

And now we're facing another wrinkle. In early 2025, the Common Vulnerabilities and Exposures (CVE) program, run by MITRE and used as the backbone for tracking known vulnerabilities, nearly lost federal funding. CISA had to step in to prevent a disruption to the program. The implications of that moment weren't just bureaucratic. They were a stark reminder: if your entire vulnerability management program hinges on third-party databases or a single stream of patching intelligence, you are building on sand. Even institutional support for foundational cybersecurity programs is not guaranteed.

This reinforces what those of us in the OT trenches already knew. Patching is only one mechanism, and not the most dependable one. The risk decisions we make must outlast funding cycles and agency politics.

If Not Patching, Then What? is the real question.

We need a credible alternative if we know that patching can't be the dominant strategy. The answer is layered: a combination of mitigation, segmentation, detection, and risk-informed prioritization.

Some of the most resilient OT environments I've assessed didn't have the highest patch compliance. Still, they had robust network segmentation, tight control over engineering workstations, clear asset visibility, and disciplined change control processes. These made it very hard for a vulnerability to turn into an exploit.

Mitigation strategies include:

- Deploying industrial firewalls and unidirectional gateways to isolate high-risk assets
- Leveraging read-only patching registries for validation before broad deployment
- Creating clear operational playbooks for compensating controls when patching isn't viable

None of these are new ideas. But doing the basics well (and consistently) in OT is far more valuable than chasing a theoretical patch compliance number.

Real-World Example: Choosing Safety Over Compliance

A few years ago, I worked with a pharmaceutical manufacturer who discovered a critical vulnerability in a component of their environmental control system. The vendor issued a patch within days. Their IT security team flagged it as a top priority and pushed hard for immediate deployment. But the site operations team pushed back.

Why? Because the control system had undergone FDA validation. Applying the patch would require not just downtime but a full completion of the clean room environment, a process that would take weeks, cost hundreds of thousands in compliance testing, and delay production runs of a high-margin oncology drug.

Instead of patching, the site implemented firewall rules to block potentially exploitable ports, restricted remote access, and increased monitoring. They reviewed and logged every interaction with the device and restricted physical access to engineering laptops.

Security purists might balk, but the decision balanced risk, business value, and compliance requirements, and it worked.

This isn't an exception. It's the norm for mature OT environments.

The Metrics Trap

One of the most damaging things I see is organizations trying to force IT-centric metrics into OT. Patch compliance is a perfect example. When CISOs are told to report the "percentage of patched assets," they often realize how poor that signal is in an OT context. The number tells you very little about the actual risk.

A factory that has patched 90% of its assets but hasn't mapped its network or doesn't know which Engineering Workstations (EWS) are Internet-facing might look good on paper but is operating blind. Meanwhile, a plant with only 40% patch compliance but strong segmentation and access controls is likely in a far better security posture.

Security metrics in OT need to evolve. We should be asking questions like:

- What percentage of OT assets have defined risk mitigation plans?
- How many unpatched vulnerabilities are mitigated through compensating controls?
- What's the time to containment when a new OT exploit is identified?

These are more actionable, and more aligned with operational realities.

Toward a Risk-Based, Context-Aware Future

Patching has a role. But in OT, it must be part of a broader ecosystem of risk reduction, not the North Star.

If we want to secure industrial environments, we need to meet them where they are. That means building programs that prioritize the business context, respect the constraints of safety and uptime, and measure success not in theoretical compliance but in actual defensibility.

The future of OT vulnerability management isn't about how fast we patch. It's about how well we protect what matters in the way that works best for the business.

How to Recalculate CVSS Scores in an OT Context

In the IT world, a critical vulnerability with a CVSS score of 9.8 triggers alarms, tickets, executive notifications, and a countdown to remediation. That number carries urgency. It drives board-level discussions, internal SLAs, and compliance audits. But in OT, that same 9.8 might be functionally irrelevant, or, at the very least, deeply misleading.

The Common Vulnerability Scoring System (CVSS) has become the de facto standard for evaluating software vulnerabilities, and for good reason. It offers a consistent, structured way to assess technical severity. However, its original design never accounted for the unique constraints, architectures, and business dependencies that define operational technology environments.

If you're leading cybersecurity for a power utility, a food manufacturing plant, or a pharmaceutical operation, you've likely encountered this disconnect. A centralized scanner flags dozens of critical devices. Security leadership wants a resolution. But your plant floor teams are quick to point out that many of those devices aren't exposed, accessible, or, in some cases, even patchable.

So where does that leave us? CVSS is not wrong. It's just incomplete. If you're going to use it in OT, you need to recalibrate it through the lens of context, compensating controls, and business impact.

Why Raw CVSS Doesn't Work in OT

Let's be clear: CVSS Base Scores were never designed to capture the realities of physical operations. They focus on inherent exploitability (network access, required privileges, user interaction) and potential technical impacts like confidentiality, integrity, and availability.

However, in OT, exploitability is often contained by design. Assets are physically separated, network traffic is tightly scoped, and changes require specialized access. You can have a critical vulnerability on a PLC, but if that device sits behind two layers of segmentation, requires proprietary software to interact with, and has no internet exposure, the real-world risk is nowhere near the base score.

Conversely, a low-severity vulnerability on a misconfigured historian server with persistent access to Level 2 networks might present far more risk than its CVSS number suggests. This inversion is common and dangerous if misunderstood.

The Missing Variable: Environment

CVSS includes an "Environmental Score" component, but few organizations use it. In OT, this piece becomes essential. It's where we introduce the true variables of risk:

- Network segmentation and firewall rules
- Physical access controls
- Device role and functional criticality
- Dependency mapping (what breaks if this breaks?)
- Monitoring and detection coverage

Applying the environmental score means asking the right questions:

- Can this vulnerability be exploited remotely?
- Is it accessible from any unmanaged asset?
- Does it control or influence a safety function?
- How often is the device used, updated, or accessed?

When you apply this context, CVSS 9.8 might drop to a real-world risk score of 4.2. Or it might stay high. But at least you're seeing the truth, not just the math.

A Practical Example

Take a Level 1 device: a servo controller in a packaging line. The scanner reports a CVSS 10 vulnerability related to remote code execution. Sounds dire. But let's dig deeper.

- The device communicates via a proprietary protocol that is not routable beyond the cabinet.
- It sits behind an unmanaged switch, meaning no external visibility unless intercepted through an engineering workstation.
- The production line is segmented from the corporate network with a firewall and dual NIC jump host architecture.
- The vulnerability requires a specific firmware version that this device doesn't run.

When you factor all this in, the real-world risk is negligible. Yet the vulnerability still triggers a critical alert in a dashboard. If your program is driven solely by base CVSS scores, you'll waste time, budget, and credibility chasing ghosts.

Creating Your OT Risk Lens

It's time to stop accepting raw CVSS scores as gospel. Instead, use them as inputs: starting points for a more mature, OT-aware evaluation process.

Here's a simplified structure I often recommend:

Start with CVSS Base Score: It provides a standardized measure of inherent technical severity.

Apply Asset Role Weighting: Is the asset a PLC controlling critical batch processes or a maintenance laptop?

Assess Network Context: Is the asset routable from business networks, remote access systems, or the cloud?

Evaluate Compensating Controls: Are there isolation measures, jump hosts, NAC policies, or read-only access enforcement?

Account for Business Process Impact: If the asset goes down, what will be the operational cost, compliance exposure, or safety risk?

The output is a context-calibrated risk score (call it "OTCVSS" if you like). What matters is that your prioritization is rooted in reality, not in theory.

Why This Matters for Leadership

CISOs and CIOs must be able to explain to the board why a flood of critical vulnerabilities doesn't automatically equate to a crisis. Without that translation, security teams look like they're asleep at the wheel, or worse, ignoring real threats.

When you build a context-aware scoring model, you create alignment. You can show why one site's 30 unpatched criticals are lower risk than another site's 3 that reside on high-impact assets with exposed interfaces. You shift the narrative from "Why haven't you patched this" to "How are we reducing risk in the smartest way possible."

This approach also builds credibility with plant leadership. Instead of bringing them problems, you're showing them a path forward, one that respects production constraints and regulatory realities.

CVSS Isn't Going Away, But It Must Evolve

We're not advocating for the abandonment of CVSS. It's valuable. But in OT, it must evolve. And that evolution starts with you.

If you're leading vulnerability management for industrial environments, your responsibility is not just to identify issues. It's to interpret them accurately and prioritize response based on real operational impact.

That means pulling in context from engineering, operations, and compliance. It means looking past the red flags on dashboards and asking: What's the actual exposure? What does this asset do? Who can touch it? What happens if it fails?

These questions aren't outside the scope of cybersecurity. They are cybersecurity in OT.

The RISE Model: Prioritizing OT Vulnerabilities Where It Matters Most

Focus Limited Resources Where They'll Deliver Maximum Impact

If every vulnerability were equal, our jobs would be a lot simpler. We could line them up, knock them down, and call it a day. However, the reality inside OT environments tells a different story. One where not all vulnerabilities are worth the same effort, and not all fixes yield the same return.

When you're operating in a world where patching isn't always feasible, and resources are stretched thin, prioritization isn't just a helpful tactic. It becomes the foundation for effective risk management. That's why we built the RISE model. It's not a theory or a vendor framework. It's born out of years of watching teams chase the wrong vulnerabilities while real risks quietly simmered just out of view.

RISE stands for **Risk**, **Impact**, **Scoring**, **Improvement**, and **Effort**. It's a pragmatic lens designed to help OT cybersecurity leaders answer a deceptively complex question: Where should we act first?

The Problem with Traditional Prioritization

Too often, external forces dictate prioritization in OT environments: scanners flagging CVSS 9s and 10s, compliance audits calling for remediations, or dashboards that paint risk in shades of red. These inputs are useful, but they're incomplete. They fail to capture the nuance of physical systems, safety-critical dependencies, and the operational pain associated with even minor changes.

We've already covered how raw CVSS scores can be misleading. But beyond scoring, the entire structure of traditional prioritization assumes the cost of action is low, and the benefit is high. That's rarely the case in OT.

You need a model that respects the tradeoffs. That's what **RISE** is built to do.

1. Risk. Given the asset's exposure and function, this is the potential for exploitation. It includes things like:

- Is the vulnerability reachable from an unmanaged or remote system?
- Is the device involved in production or safety-critical processes?
- Could this be a pivot point to higher-value targets?

Risk isn't a score pulled from a database. It's context-specific. OT must account for both cyber pathways and physical process exposure.

2. Impact. What's at stake if this vulnerability is exploited? Not just technically, but operationally:

- Could it halt a production line?

- Would it violate a safety constraint or environmental regulation?
- Does it tie back to compliance reporting or customer deliverables?

You're not just protecting assets. You're protecting outcomes. Impact elevates that reality.

3. Scoring Improvement. This is the most overlooked factor in OT vulnerability conversations. It's about the delta. How much risk do we eliminate if we remediate or mitigate this issue?

Sometimes, fixing a vulnerability cuts off a major attack path. Other times, the fix offers only marginal improvement because other exposures still exist. We must stop spending 80% of our effort for 5% of the risk reduction.

4. Effort. The cost of mitigation isn't just technical. It's political, operational, and economic.

- Does the fix require downtime?
- Will it impact validated systems or require requalification?
- Is there a playbook, or is this a one-off engineering lift?

In OT, effort is a gating factor. Even a high-risk issue might need to be deferred if the mitigation introduces unacceptable disruption. And that's not weakness. That's leadership.

Applying RISE in the Field

Let's say you have five vulnerabilities across different devices:

- A Level 1 drive with a denial-of-service CVE
- A remote access service on an engineering laptop
- A firmware issue on a PLC used in packaging
- An out-of-date OS on a historian
- An unpatched browser vulnerability on a QA station

Traditionally, the ones with the highest CVSS scores would rise to the top. But through the RISE lens, your ranking might look very different.

The engineering laptop may have lower CVSS but presents a high risk, high impact, and moderate scoring improvement, with low effort to isolate or remediate. That becomes your top priority.

The historian OS might have a high score but sit behind a series of segmentation boundaries with no real exposure and patching it might require 12 hours of requalification. That drops.

It's not about ignoring severity. It's about putting it in context.

Building a RISE Dashboard

Applying the RISE model doesn't require buying new tools or launching a major project. It requires alignment between teams. OT security must coordinate with engineering, operations, and safety to understand real-world variables.

Start small. Build a spreadsheet. Score your top 10 open vulnerabilities using RISE on a 1–5 scale for each category. You'll immediately see where intuition and dashboards diverge, and that's where your leadership comes in.

Over time, you can operationalize RISE into your governance framework. Assign ownership, create workflows for exceptions, and build reporting that tells the story of decisions, not just alerts.

The Leadership Imperative

Boards don't want noise. They want decisions. They want to know that the OT cybersecurity team spends its time and budget on what matters most. RISE gives you a language to explain that.

When a regulator asks why a critical vulnerability hasn't been patched, RISE allows you to respond confidently: "Yes, it's high severity. But based on our RISE model, it represents low real-world risk, minimal scoring improvement, and high operational cost. We've applied compensating controls and are focusing our efforts on higher-leverage issues."

That's not deflection. That's risk management.

A Model Grounded in Reality

The RISE model isn't about chasing perfection. It's about optimizing decisions. It helps you focus on how intervention changes outcomes. That's the kind of leadership OT environments need right now.

You can't fix everything. You're not supposed to. But with RISE, you can fix the things.

Mapping OT Business Risk Using the QUICK Framework

Why Understanding What's "Critical" To The Business Is More Important Than Scoring Systems

One of the hardest things to explain to people outside the plant is that not all assets carry the same weight. A vulnerability on one controller might be a nuisance, while the same issue on another could cost millions. That difference often has nothing to do with network placement or software version. It has everything to do with the asset's function in the business.

Cybersecurity doesn't happen in a vacuum, especially in operational technology environments. It happens during uptime demands, product quality targets, safety rules, compliance constraints, and customer commitments. And if your vulnerability management strategy doesn't account for that reality, it will always miss the mark.

That's why we use the QUICK framework. It classifies OT assets not by their technical specs but by their role in delivering business outcomes.

QUICK stands for:

Quality, **Uptime**, **Information Integrity**, **Compliance**, and **Key Deliverables**.

It's not a replacement for technical analysis. It's a decision filter. One that puts context around the systems you're defending, and helps you triage vulnerabilities based on what's really at stake.

The Cost of Uniformity

Most enterprise vulnerability management programs apply the same logic to all assets: score, prioritize, patch. But in OT, uniformity creates risk.

You can't treat a QA lab workstation the same way you would a control system that regulates temperature for a fermentation process. One runs analysis and the other governs a multi-million-dollar operation that is subject to strict regulatory oversight.

Yet many security programs do just that, because they lack a model to differentiate assets based on business impact. **QUICK** fills that gap.

What Each QUICK Dimension Tells You

Quality

Does the asset influence product quality? Think vision systems, batching logic, and temperature controllers. A vulnerability that affects product specs can trigger recalls, scrap costs, and brand damage.

Uptime

If this system goes offline, does it impact production? Uptime-critical assets are often deeply embedded and harder to restart safely. Not all downtime is equal: some lines can catch up overnight, and others create ripple effects that last for weeks.

Information Integrity

Does the asset ensure correct, trusted data? Historians, batch records, and lot tracking systems fall into this category. Vulnerabilities here can result in invisible errors that only surface after the product has shipped.

Compliance

Does the system support regulatory, environmental, or safety requirements? Think environmental monitoring, clean room control, or audit logs. A failure or exploit may not just disrupt operations. It can lead to fines or legal exposure.

Key Deliverables

Is the asset tied to a contractual or business-critical obligation? For example, a production line that feeds a tier-one supplier or a machine involved in pilot runs for new product introductions. These assets represent an outsized business risk because they impact external commitments.

From Asset Inventory to Risk Categorization

Let's say you've built a decent asset inventory. That's a good start. But now it's time to overlay meaning.

Using the QUICK model, classify your assets based on their business dependencies. You don't need perfection, just clarity. A simple 1–5 score across each category is enough to begin.

The goal isn't to create new documentation bureaucracy. The goal is to elevate what your controls engineers already know into a shared language that your security program can use.

Your line supervisors, process engineers, and plant managers know which machines can't go down. They know which batch reports are sacred and which are backups only. QUICK gives security a framework to capture that knowledge, and use it.

Why It Works in the Field

A few years ago, I walked into a facility that had flagged a PLC vulnerability as a high priority because of its CVSS score. The remediation plan called for taking the asset offline during a maintenance window. But a process engineer pulled me aside.

"That controller manages flow balance between two pressure-sensitive tanks. If you pull it down, even briefly, you're shutting down the line and introducing product waste. We'll be throwing away $50,000 in raw materials."

That PLC wasn't marked as a high value in the CMDB. But under QUICK, it would've been red-flagged for Uptime and Key Deliverables. We adjusted the plan.

This is what makes QUICK valuable: it captures what the business truly values, not what the database assumes is important.

Integrating QUICK into Governance

Once you've classified your assets, use QUICK scores to influence your vulnerability triage process. When a vulnerability is detected:

- Overlay its RISE score with the asset's QUICK rating
- Identify where business-critical assets intersect with high-risk
- Prioritize actions that defend your most valuable operational dependencies

This isn't about coddling operations or diluting security. It's about protecting what matters most. QUICK doesn't replace technical analysis. It refines it.

Bridging the Business-Security Gap

Too many OT security programs struggle to communicate their priorities to executive leadership. QUICK gives you a narrative.

When questioned why certain vulnerabilities were escalated (or deferred), you can say:

"These assets underpin our FDA compliance and are tied to critical customer deliverables. Remediation risk must be balanced with business continuity. Here's the plan to manage that risk effectively."

That message lands. It shows that security isn't working in isolation. It's operating in service of the business.

Build on What People Already Know

The truth is, you don't have to invent this intelligence. It already exists inside your plants. QUICK is about surfacing it and turning operational insight into a cybersecurity context.

Use it to guide vulnerability prioritization, shape segmentation strategy, and frame discussions with the board.

Because when you know what matters most, you can defend it best.

HILLSTRONG GROUP SECURITY

OT VULNERABILITY MANAGEMENT · CHAPTER FIVE

How to Build a Sustainable OT Vulnerability Management Program

From Project To Practice: Embedding VM In Your OT Security Cycle

There's no shortage of urgency when it comes to OT vulnerability management. The threat landscape keeps expanding. Regulatory pressure keeps mounting. And vendors keep promising solutions that will solve it all with one more sensor, scanner, or platform.

But if experience teaches one thing in this space, it's that no technology can compensate for the absence of a sustainable program. Tools come and go. The real impact is created through governance, alignment, and consistent execution, especially in the high-stakes, high-friction environments that define industrial operations.

Building an OT vulnerability management program that lasts means accepting a few hard truths. First, you can't fix everything. Second, some risks will persist for years. Third, success isn't measured by how many vulnerabilities you close. It's measured by how effectively you reduce risk without disrupting what the business is built to deliver.

From Scanning to Strategy

Most OT vulnerability efforts start with a scanning exercise. That's fine. Visibility matters. However, the gap between asset discovery and actual risk reduction is wide. If you don't define what happens after the scan, you've just added noise.

What does a sustainable program look like? It has clear ownership, defined decision paths, and repeatable processes for triage, mitigation, and monitoring. It integrates with operations, not just IT. And it doesn't fall apart when one SME leaves or one platform changes.

Think of it as an operating rhythm, not an incident response plan. It has to run monthly, quarterly, and annually across a fleet of sites, each with its own constraints.

Define Roles and RACI Early

One of the first cracks in most programs shows up around ownership. Who owns remediation? Who owns exception approval? Who determines whether a compensating control is adequate?

This is where RACI comes in: Responsible, Accountable, Consulted, and Informed. But in OT, the standard IT RACI often misses the mark. You need to account for:

- Plant operations leadership

- Controls engineering teams
- EHS or quality stakeholders
- IT security and GRC functions

Too often, decisions get bottlenecked because the wrong group is driving, or worse because no group is.

Establishing your RACI early prevents drift. It also builds clarity when sites push back on remediation due to production concerns. You're not escalating the issue. You're following the program.

Centralized Governance, Decentralized Execution

Industrial enterprises tend to be federated. That's not going to change. Trying to enforce a monolithic vulnerability program across vastly different sites rarely works.

Instead, your goal should be to build a centralized governance model with site-level flexibility. That means:

- Defining corporate standards and thresholds (e.g., max allowable time-to-mitigate for criticals)
- Providing a shared playbook for triage and exception handling
- Allowing sites to implement controls that match their unique systems and staffing

Your program must account for everything from legacy Windows 7 HMIs in one plant to air-gapped microcontrollers in another. This requires a uniform process and localized adaptation.

Make Risk Reporting Real

Executives don't want vulnerability lists. They want decision context. That's where the work you've done with RISE and QUICK becomes invaluable.

Instead of reporting "We have 374 open vulnerabilities," report:

- The percentage of OT assets with a risk-adjusted mitigation plan
- The number of high-risk vulnerabilities mitigated this quarter
- Sites with the highest risk density by asset function and business impact

You're telling a risk story, not reciting technical inventory. That builds credibility with the board and with plant leadership.

Normalize Exceptions, But Track Them

In IT, the idea of permanent exceptions to patching feels uncomfortable. In OT, it's often necessary. But that doesn't mean it should be unmanaged.

Build an exception process that tracks:

- Why remediation wasn't feasible
- What compensating controls were applied
- When the exception will be reviewed again

A site that has 10 exceptions but full documentation and layered controls is in a much better place than one that claims perfect compliance but can't explain its risk posture.

Invest in Repeatable Cycles

You don't need a perfect program on Day 1. What you need is a cadence: monthly vulnerability review meetings, quarterly scorecard updates, and annual risk reclassification.

Over time, this rhythm does more than close gaps. It changes culture. It makes vulnerability management part of how you operate, not something that spikes before audits.

You'll also start to see the benefits compound. The same asset classification that informs vulnerability prioritization can inform segmentation projects, procurement reviews, and detection rule tuning. Security becomes embedded, not bolted on.

Don't Let the Program Die on Paper

The final trap I've seen is beautiful policies and frameworks that live in SharePoint and never make it to the shop floor. Don't mistake documentation for execution.

Go walk the line. Sit with the control engineers. Ask what works, what doesn't, and what gets ignored. You'll quickly find out where the gaps are, and where the real levers of change live.

A sustainable vulnerability management program is pragmatic, visible, and rooted in how people work. Anything less will collapse under the weight of operational complexity.

The Path Forward

You've now seen the pieces:

- Why patching alone fails in OT
- How to rethink CVSS in context
- How to prioritize using RISE
- How to understand asset value with **QUICK**
- And how to build a program that brings it all together

The path to resilience isn't paved with tools. It's built through alignment, adaptation, and sustained effort. Vulnerability management in OT isn't a sprint or a checkbox. It's a commitment to protecting what matters in the ways that matter most. This is the work, and it's how we move the mission forward.

About the Author



Roger Hill

CEO, Hillstrong Group Security, Inc.

Roger Hill is the founder of Hillstrong Group Security and a recognized leader in industrial cybersecurity. With over 30 years of experience spanning operational technology and industrial automation, Roger brings a sharp, forward-thinking perspective to securing critical infrastructure. At Hillstrong, he pioneers new strategies to bridge cyber and physical risk, helping global organizations build resilient, future-ready systems. His work is rooted in real-world execution, not theory, and his mission is to lead by doing.

He is a vocal advocate for eliminating hype and focusing on a pragmatic approach with practical outcomes that align cybersecurity with business impact. Roger's thought leadership continues to influence how modern manufacturers approach governance, resilience, and the convergence of OT and IT security.

HILLSTRONG GROUP SECURITY

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