HILLSTRONG GROUP SECURITY
OT CYBERSECURITY SERIES

THE OT GOVERNANCE PLAYBOOK FOR MANUFACTURING CISOs

GOVERNANCE, VISIBILITY, AND DECISION RIGHTS



HILLSTRONG GROUP SECURITY

CHAPTER 1

OT Cyber Risk Needs Senior Visibility Above the Plant Level

1

OT Governance for Manufacturing CISOs



Most manufacturers do not lose control of OT cyber risk because they lack another dashboard, control family, or plant security project. They lose control because the risk stays trapped in plant language too long.

A site has a remote access exception for an OEM. Another site has the same exception. A third site inherited the same pattern during a line expansion. Nobody lies about it. Nobody hides it. The problem is worse. Everyone treats it like a local operating fact instead of an enterprise decision. By the time corporate security sees the pattern, the company has already standardized its exposure.

That is the real argument for senior visibility.

OT cyber risk stops being local when it can hit multiple plants, distort capital decisions, expose the same vendor pathway across business units, or force the company into bad tradeoffs between uptime and control. At that point, it is not just a plant problem, an engineering problem, or a technical debt problem. It is enterprise operating risk.

That distinction matters for manufacturing CISOs because OT security programs often get discussed at the wrong altitude. The plant talks about uptime. Engineering talks about feasibility. Security talks

about controls. Procurement talks about vendors. The board hears about OT only after a material incident, a budget request, or a modernization initiative goes sideways. That is late. The company has already allowed local decisions to harden into enterprise exposure.

The evidence base behind this series is clear on one point above all others: material OT cyber risk should be visible above the plant and business-unit level. CISA's Cross-Sector Cybersecurity Performance Goals call for named leadership for cybersecurity overall and named leadership for OT-specific cybersecurity in organizations with OT assets. ASD/ACSC guidance says the board or executive committee should define cyber roles, integrate cybersecurity with business functions, and receive regular briefings. The UK Cyber Governance Code of Practice pushes the same direction — cyber oversight belongs with directors and governance bodies, not buried inside day-to-day operators.

That does not prove one universal reporting line. It does prove something more useful. If a manufacturer cannot show OT cyber risk above the plant, the enterprise is making repeated risk decisions without enterprise visibility.

The plant is where the constraints live. It is not where the full risk picture lives.

Manufacturing leaders are right about one thing. Plants carry realities that corporate teams often flatten or ignore. They carry shutdown windows that move twice a year and never by enough. They carry fragile controllers that support revenue-critical lines but cannot be patched on the timetable a corporate standard assumes. They carry OEM dependencies that were baked into the equipment long before anyone in the current security team had the job. They carry safety implications that make clean IT answers look reckless.

That is all real. It is also incomplete.

The plant sees the local constraint. The enterprise sees the pattern. You need both.

A site may be right that a given HMI cannot be updated before peak production. A line engineer may be right that a segmentation change has to wait for a maintenance window. A plant manager may be right that a remote vendor session is the only way to restore a critical process before the weekend shift. None of that settles whether the same exception exists in nine other sites, whether the same vendor touches half the footprint, whether the enterprise has funded the root fix, or whether the company has normalized a risky operating habit.

That is the gap senior visibility is supposed to close.

OT cyber risk should not be handled as a local technical matter alone. It should be visible above the site and business-unit level because the consequences and patterns reach higher than the site that happens to feel them first.

Repeated local decisions become enterprise exposure fast

This is where OT governance usually gets lost.

One local exception looks manageable. Ten local exceptions with the same root cause are a governance problem. One plant granting remote OEM access during a shutdown looks practical. Twelve plants using the same access method with different sponsors and no shared control model is an unmanaged trust boundary.

You should press hard on that distinction because it changes the whole conversation with leadership.

Translate plant language into enterprise language

What the plant says	What the enterprise needs to hear
"This OEM needs remote access to restore the line."	"We have an unmanaged trust boundary replicated across multiple sites with no enterprise control owner."
"We can't patch during peak season."	"We are carrying accepted risk on revenue-critical assets that the enterprise has not formally chosen."
"It's how the integrator set it up."	"We have inherited an architectural pattern from a third party without enterprise review."
"The HMI is too fragile to touch."	"We have asset-class fragility that constrains response options across the footprint."
"We've always done it that way."	"We have normalized an operating habit that has never been risk-assessed at the enterprise level."

Senior visibility is not a dashboard. It is a decision path.

This is where CISOs often lose executives. They say OT needs more visibility. Leadership hears 'more reporting.' They say OT needs board attention. Leadership hears 'another cyber presentation.' They say the enterprise needs a better view of plant risk. Leadership imagines a bigger dashboard full of controls nobody trusts.

That is the wrong framing.

Senior visibility is not a visual product. It is a decision path. It means OT issues show up in enterprise risk review when they stop being local. It means recurring exception patterns reach executives before they become normalized operating practice. It means major modernization risk gets discussed before the rollout creates a new shared vulnerability across sites.

The board does not need a lecture on Purdue levels to do its job. It needs to know where the company has material OT risk themes, which unresolved issues repeat across the footprint, where a vendor or architecture creates concentration risk, and which investment decisions affect resilience or operating continuity.

For manufacturing CISOs, the practical test is simple: when a serious OT issue appears, can you move it from plant language to business language without stripping out the operational truth? If the answer is no, you do not have senior visibility yet.

What boards should actually see

The board should not see a pile of OT noise. It should not see 80 slides of architecture diagrams, scanner counts, or vague control maturity scores divorced from operational consequence. It should see the issues that change enterprise risk posture. That usually means five things:

Material OT risk themes. Not isolated alarms. Patterns. Repeated remote-access dependencies. Aging unsupported assets in critical process areas. Shared weak points across similar lines or similar business units.

Recurring exception patterns. How many critical exceptions exist, how long they stay open, what classes of assets they affect, and whether the same reason keeps appearing in different sites.

Concentration risk. Which OEMs, integrators, remote access methods, or network patterns create shared exposure across the estate.

Unresolved issues with operating consequences. The board needs to know where the company is carrying known OT risk because of production timing, capital deferral, or unresolved ownership conflict.

Investment choices that shape resilience. Not generic budget requests. Specific asks tied to repeated enterprise exposure, modernization backlog, incident authority, or vendor control weaknesses.

Where most CISOs lose the room

This is not a knowledge problem. It is a framing problem.

You lose the room when you overexplain OT architecture before you explain operational consequence. You lose the room when you open with threat headlines instead of showing how the company's own decisions create repeated exposure. You lose the room when you ask for funding before proving that the issue is systemic rather than local.

The better move is to start with the operational truth and then widen the aperture. Say the line cannot be patched before the quarter ends. Fine. Ask how many other lines depend on the same asset family.

Say the vendor needs remote access to restore the process during an outage. Fine. Ask which sites use the same method and who owns the enterprise control standard.

That is how you earn credibility with both operations and the board.

Attach OT risk to levers the company already responds to

Executive teams do not act on OT risk because OT risk is interesting. They act on it when it attaches to something they already pay attention to. The usable levers:

Cyber insurance renewal. Carriers now ask specific OT questions and write specific OT exclusions. A material gap shows up in pricing or coverage. That is a CFO conversation, not a security conversation.

Regulatory exposure. SEC cyber disclosure, TSA security directives, CIRCIA reporting, NIS2 in Europe. These create board-level reporting obligations the General Counsel already tracks.

Customer security questionnaires. OEMs, automakers, defense primes, and large buyers increasingly require OT-specific attestations. A failed answer affects revenue.

M&A diligence findings. Acquired plants almost always carry inherited OT debt. Surfacing that during integration is the cleanest moment to get capital attention.

Internal Audit findings. An audit finding lands differently than a security memo. Use that channel deliberately.

Peer-company incidents. Norsk Hydro, Clorox, JBS, MKS Instruments, Dole. When a peer takes a public hit, the board's tolerance for 'we'll get to it' drops for sixty to ninety days. That window is leverage.

★ YOUR MOVE

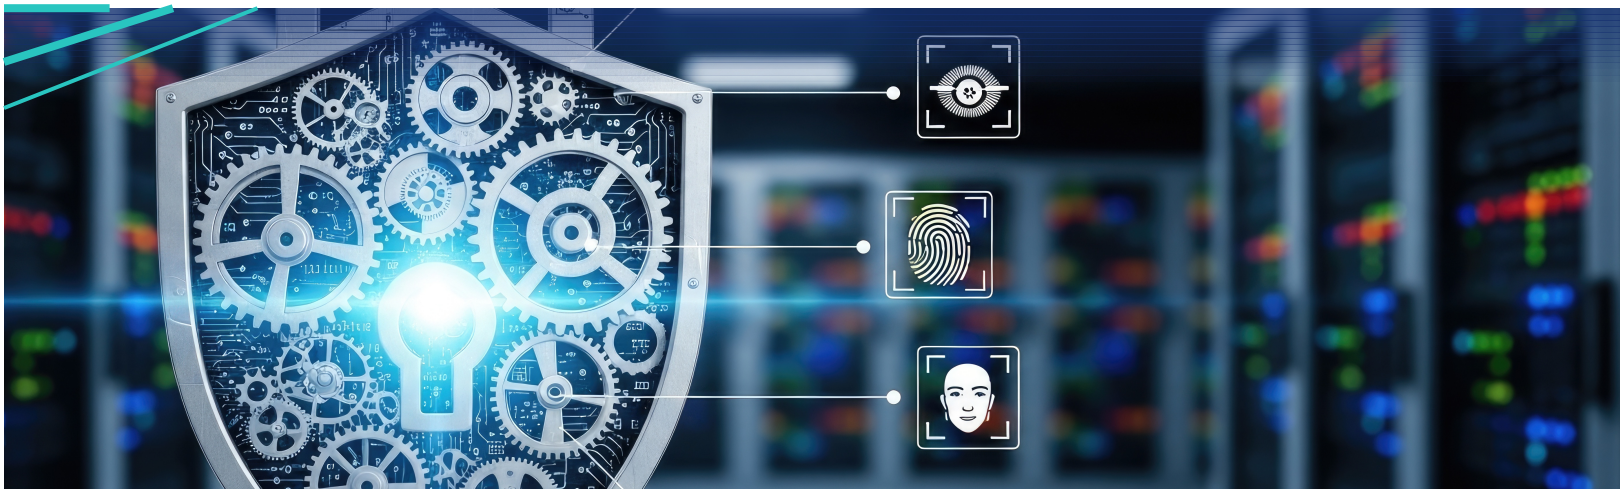
Before next week's operations review, pick a single OT issue that has been raised at more than one site — a remote access pattern, an unsupported asset class, a recurring exception, a vendor dependency. Count the sites it touches. Write one sentence that turns the count into an enterprise statement. Bring that sentence to the next ops meeting and read it out loud. You are not asking for a decision yet. You are establishing that the pattern exists, that you can name it, and that you intend to keep naming it until the enterprise decides what to do about it.

CHAPTER 2

Named Leadership and Named System Ownership Matter More Than the Org Chart

2

OT Governance for Manufacturing CISOs



Most OT ownership problems are not hidden. The company just avoids naming them. That sounds harsh. It is still true.

Ask three questions in a room full of security, engineering, and operations leaders. Who owns enterprise OT cyber direction. Who owns each critical system or site. Who owns the risk when local reality collides with enterprise policy.

In too many manufacturers, the room gets quiet by question three. Some people will answer with titles. Some will answer with functions. Some will answer with committee names. Somebody will say ownership is shared. Somebody else will say the plant owns the asset, so the plant owns the risk. Security will point to standards. Engineering will point to uptime. Operations will point to production targets. Nobody is lying. That is the problem. Everybody owns part of the truth and nobody owns the seam where the hard decision has to get made.

That is why the org chart debate wastes so much time. Should OT sit under the CISO. Should it sit under operations. Should it sit under engineering. Should it run through a hybrid model. Those questions matter less than people think. The stronger conclusion from the source set is narrower and

more useful. The enterprise needs someone accountable for standards, reporting, and program direction. The site or system needs someone accountable for local facts, feasibility, and execution.

Named ownership does.

The org chart is not the operating model

This point gets missed in almost every internal debate. An org chart tells you where people report. It does not tell you who can make a high-cost call when the line is running hot, the vendor needs access, the maintenance window is gone, and the enterprise standard no longer fits the local situation.

That gap matters more in manufacturing than in most enterprise security work. In IT, a central team can often enforce a control with less concern for physical process impact. In OT, the same control can create production loss, safety consequences, or a restart sequence the plant cannot absorb. That means the enterprise needs a governance model that can hold two truths at once. Corporate leadership has to own standards, direction, and risk visibility. Local owners have to own the physical and operational reality that determines whether a control can be applied safely.

The enterprise needs somebody who owns direction. The plant needs somebody who owns local truth. If neither is named, the company does not have an OT operating model. It has a collection of assumptions.

The three ownership questions every manufacturer should answer

Most OT confusion compresses down to three questions:

Who owns enterprise OT cyber direction.

Who owns each critical system or site.

Who owns the risk when local constraints make the enterprise standard hard or impossible to apply.

If a company can answer those three questions cleanly, a lot of noise disappears. If it cannot, the rest of the governance language is decoration.

Who owns enterprise OT cyber direction

This is not the same as who supports the plants. It is the role that owns standards, reporting, enterprise prioritization, and program direction. In many manufacturers, that role belongs with the CISO or equivalent enterprise cyber executive. Enterprise accountability should stay with the CISO or equivalent cyber executive, even when OT operational translation is delegated. That is a practical recommendation rooted in the evidence, not an ideological position.

Somebody has to own the baseline — standards for remote access, identity, segmentation, exception review, incident escalation, and risk reporting. Somebody has to decide when a repeated plant issue is no longer local. Somebody has to represent OT risk in front of the board, the executive committee, and

the budget process. If that role is diffused across plants or left inside engineering without enterprise authority, the company can still run sites. It cannot run an enterprise OT program.

Who owns each critical system or site

This is the part companies skip because it sounds obvious. It is not obvious. Owning the asset in a maintenance database is not the same as owning the system in a governance sense. A real system owner can explain what the system does, what constraints shape it, which shutdown windows exist, which vendors touch it, what exceptions are open, what compensating controls are real, and which changes are safe to sequence without breaking process performance.

A plant without a real system owner creates the same failure pattern every time. Security asks for a control. Engineering says it is risky. Operations says timing is bad. Maintenance says they need a vendor. Nobody can translate the issue into a defined local decision. The meeting ends with 'we need to discuss offline.' The exception survives by inertia. That is not a control gap. That is an ownership gap.

Who owns the risk when local reality conflicts with enterprise policy

This is the hardest question because it exposes whether governance is real. Corporate sets a standard. The plant cannot follow it on the required timetable. The exception should be documented. The local constraint should be described clearly. Compensating controls should be explicit. A named risk owner should decide at the right level of materiality. That is the difference between an operating model and a standoff.

What the enterprise owner actually owns

Manufacturers get into trouble when they say the enterprise owns OT security but never define what that means. The enterprise owner should not be expected to understand every local controller quirk, restart sequence, or process constraint. That is not the job. The enterprise owner does need to own five things:

Standards. Not slogans. Standards for remote access, identity, segmentation, logging, incident escalation, exception handling, and governance expectations.

Program direction. Which OT initiatives matter this year. Which architectural gaps get funded. Which repeated exposure patterns need enterprise action.

Executive reporting. Someone has to translate plant-specific issues into enterprise risk language without flattening them into cyber wallpaper.

Shared governance. Someone has to convene the cross-functional forum where security, operations, engineering, safety, and risk can settle issues that do not belong to one silo.

Systemic risk visibility. The enterprise owner has to know when a local issue is repeating across the estate.

What the local owner actually owns

The local owner does not own enterprise policy. The local owner owns reality. That includes maintenance constraints, outage sequencing, process behavior, local vendor dependencies, restart risk, the truth about asset condition, and the local facts behind any exception request.

A good local owner can say: 'We cannot patch this controller before the next shutdown because the packaging line is in peak run, the OEM requires onsite support for restart, and the line feeds two downstream processes that would need synchronized downtime. We can apply these compensating controls now. We need this long-term fix in the capital plan.' That is useful. A weak local ownership model produces: 'We cannot do it right now. The line is too important.'

Where ownership usually breaks down

Corporate writes a standard the plant cannot execute on schedule.

Engineering owns the asset but has no formal escalation path for enterprise risk.

Operations owns uptime but not cyber exception logic.

A vendor gets sponsored locally with no enterprise access model behind the relationship.

A business-unit leader funds a workaround because the root fix sits in somebody else's budget.

Everybody assumes somebody else owns the repeated issue.

The worst version of this problem is polite. No one fights openly. Meetings stay civil. Everyone agrees OT is important. Months pass. The same exception stays open. Nothing is owned enough to move. Manufacturing CISOs need to call that what it is: unassigned accountability with good manners.

What this means for manufacturing CISOs

If you are the CISO, stop spending the first hour of the argument on where OT should sit. Start with the names. Who owns enterprise direction. Who owns the system. Who owns the risk when local constraints collide with the standard. If those names do not exist, fix that first. If those names exist but nobody knows the decision boundaries, fix that next. You do not need one perfect org chart. You need named people who can answer for standards, systems, and exceptions when the easy answer disappears.

★ YOUR MOVE

Before your next governance review, get three names on paper: who owns enterprise OT cyber direction, who owns your most critical production system, and who would sign the exception if that system could not meet the standard next quarter. If you cannot fill in all three with confidence, that is your immediate governance gap — not the org chart.

CHAPTER 3

OT Governance Has to Cross Security, Engineering, Operations, and Safety

3

OT Governance for Manufacturing CISOs



The real OT problem lives in the seams. That is where standards collide with uptime. That is where safety collides with cyber urgency. That is where a clean enterprise control meets a dirty plant reality and nobody wants to own the tradeoff.

Most manufacturers understand this at a gut level. Few structure around it. They create a reporting line. They assign a program lead. They launch a steering committee. Then they assume the organization has solved OT governance. It has not. The question is never whether a box sits under the CISO, the COO, engineering, or some hybrid model. The question is whether security, engineering, operations, safety, risk, and the enabling functions can make a hard decision together before a production-impacting event makes the decision for them.

That is where OT governance fails.

NIST describes OT cybersecurity work as spanning IT, control engineering, operators, security experts, and enterprise risk management. ASD/ACSC supports formal steering-committee or advisory-board structures to coordinate cyber and business leadership. SANS reinforces the same lesson from incident and tabletop practice: roles and responsibilities have to be worked out across

engineering, operations, safety, and IT/security before an incident tests them.

A clean reporting line does not settle a real OT decision

This is where many OT programs stall after the first governance redesign. The company picks a home for OT cybersecurity. Good. Somebody now has formal accountability. Better than before. Then a real issue arrives.

A site needs a production-season exception on segmentation work. The OEM still needs remote access for break-fix support. The plant manager will not take a line outage without operations coverage. The safety lead wants a hazard review before any network change. Corporate security says the current path violates the standard. Who decides.

The org chart does not answer that. The org chart tells you who gets invited first. It does not tell you how the tradeoff gets settled. That is why OT governance becomes cross-functional the moment the issue stops being theoretical.

No single function owns the full problem

Each function brings something essential — and leaves a gap if absent:

Security brings standards, threat logic, incident structure, and the discipline to push recurring issues into enterprise view.

Engineering brings system behavior, fragility, control logic, restart risk, and the local knowledge that stops a clean security plan from turning into a bad plant decision.

Operations brings uptime, throughput targets, change timing, shift realities, and the business consequences of getting the sequence wrong.

Safety brings consequence thinking — real judgment about what happens when a control change, outage sequence, or degraded mode creates process risk.

Risk and compliance bring escalation language, tolerance boundaries, formal ownership expectations, and the discipline to decide whether the organization is accepting a business risk or merely delaying it.

Procurement and vendor management bring contract leverage, third-party terms, and visibility into how external access models spread across the footprint.

Infrastructure and architecture teams bring shared services, identity models, remote access patterns, and the enterprise consequences of one plant's local workaround becoming standard practice.

What a cross-functional OT governance model actually does

Bad governance models sound busy. Good ones move decisions. A real OT governance mechanism should do five things:

1. Review major exceptions and force them into the open before they become normal operating habit.
2. Resolve recurring cross-site issues — especially the ones no single plant can solve alone.
3. Prioritize shared initiatives where the enterprise has to choose between competing needs, limited shutdown windows, and limited budget.
4. Surface operational conflicts early, while there is still time to change the decision path instead of waiting for an incident or outage.
5. Sponsor exercises, post-incident reviews, and lessons learned that expose governance weakness before the same problem repeats.

What an exception path actually looks like

Take a common one. Plant X requests a ninety-day deferral on Phase 2 segmentation because the new SKU launch consumes every available shutdown window through the end of the quarter. The exception arrives in the forum with five things on the page:

The security standard being deferred and the specific control objective it serves.

The operational driver — what production outcome justifies the deferral.

The compensating control in place during the deferral window.

The date and trigger for re-evaluation.

The named risk owner if a cyber event occurs in the deferral window.

The forum decides. The decision is logged. The risk owner signs. That is the difference between governance and email. The exception still happens. The standard does not get quietly abandoned. The risk does not float without an owner.

Where companies get this wrong

Too much security, not enough operations. Security teams often build the strongest governance language and the weakest operational legitimacy. They know the standard. They do not own the restart sequence. If the plant sees the governance forum as a one-way policy engine, it will route around it.

Leaving safety out until late. A cyber control change may look rational until someone asks what happens if the process enters a degraded state halfway through execution. Safety does not fix bad security design after the fact — safety changes the decision before the design lands.

Inviting plant leaders only after standards are already set. That creates resentment for a reason. If local leaders show up too late, the conversation stops being collaborative. It becomes compliance theater followed by exception handling.

Executive sponsorship with no authority below it. Some programs look strong because the steering committee has senior names attached. Then every real decision gets pushed back down into unresolved middle layers. Executive support matters. It does not replace decision rights close to the work.

Cadence without closure. The meeting happens. Notes get captured. Actions get assigned. The same issue comes back a month later because nobody had authority to settle it, fund it, or elevate it. That is how cross-functional governance becomes overhead instead of leverage.

Exercises expose weak governance faster than policies do

A written policy can hide confusion for months. A tabletop exposes it in thirty minutes. Who can authorize isolation. Who can accept degraded operations. Who can say the plant cannot absorb the change. Who can overrule the local preference when the risk pattern is bigger than one site. Who owns the external vendor call if OEM remote access is part of the recovery path.

A good exercise does not end with 'We need better communication.' It ends with something sharper: We do not know who owns the risk acceptance when the line cannot meet the standard. We do not have a clear sponsor for third-party access in outage conditions. We let safety in too late. Those are governance findings. They matter more than whether the tabletop deck looked polished.

★ YOUR MOVE

Pull the last three major OT exceptions your organization granted. For each, write down who actually decided, what the compensating control was, and whether a re-evaluation date exists. If any of the three are unclear on any of those points, you do not have a forum yet — you have meetings. Map your current forum against the five failure modes in this chapter and score yourself honestly. Pick one upcoming production-impact decision and route it through the forum on purpose — not as a ceremony, as a working session under real constraints.

CHAPTER 4

A Useful OT RACI Starts With Decision Rights, Not Departments

4

OT Governance for Manufacturing CISOs



A RACI that looks clean in PowerPoint usually fails in the first real conflict. That is not because RACI is useless. It is because most companies use it to map functions instead of decisions.

Security gets a box. Engineering gets a box. Operations gets a box. Somebody adds safety, procurement, and risk because the matrix looks mature with more names on it. Then a real OT issue shows up and the model falls apart.

A fragile controller cannot be patched before peak production. A vendor needs remote access during an outage. A segmentation change will improve the standard but could disrupt restart timing. A ransomware event forces a choice between immediate isolation and continued production under uncertainty.

Now the question is no longer who is Responsible, Accountable, Consulted, or Informed in theory. The question is who decides. That is where most OT RACIs fail.

A useful OT RACI should be built around decision points, not generic departments. The strongest decision areas are: who owns enterprise OT standards; who owns the site-level implementation plan; who approves OT exceptions; who has authority during a production-impacting incident; who sponsors

and governs third-party OT access.

Decision one: who owns enterprise OT standards

Many organizations say standards are shared. That usually means the standard will get written by whoever has the strongest process discipline and then fought over by everybody else. Enterprise security leadership should own enterprise OT standards, with OT-specific input and formal consultation from operations and engineering. Standards need a home that can hold consistency across sites, connect to enterprise reporting, and see whether one plant's workaround is becoming another plant's default design. The plant should shape how the standard gets implemented. The plant should not own whether the standard exists.

Decision two: who owns the site-level implementation plan

The enterprise issues a standard. Everyone agrees in principle. Then the site says the sequence will not work. The site or system owner, working with local OT engineering and operations, owns the implementation plan. A corporate team can own the standard for segmentation, logging, remote access, or recovery testing. It should not pretend to own the site-level sequence for how that control lands on a live production system. Corporate owns the standard. The site owns the implementation plan. The site still has to explain the plan in operational terms, not in excuses.

Decision three: who approves OT exceptions

This is where companies reveal whether the RACI is real. Most OT environments carry exceptions — unpatchable assets, temporary connectivity allowances, deferred segmentation changes, unsupported systems with compensating controls. The issue is not whether those exceptions exist. The issue is who can approve them, at what level, based on what evidence, and for how long.

A weak OT RACI says, 'Operations and security will align on exceptions.' That is not a decision path. That is a hope. A better model says the system owner documents the issue, the local operating constraint, the proposed compensating controls, and the business impact. The appropriate risk owner decides whether to accept, reject, time-bound, escalate, or fund the root fix.

Decision four: who has authority during a production-impacting OT incident

This is the hardest decision in the whole model. If the company cannot answer who can isolate, shut down, continue degraded operations, or sequence recovery during a production-impacting incident, the RACI is decorative. The worst possible time to decide who has authority is in that room when the incident is live.

Who can authorize immediate isolation.

Who can authorize continued production under elevated risk.

Who can approve degraded operations.

Who owns recovery sequencing when security urgency and production urgency pull in different directions.

Who speaks for safety consequences.

Decision five: who sponsors and governs third-party OT access

Weak OT governance models break on vendor access faster than anything else. The enterprise needs to own the control model — the access method, approval expectations, identity pattern, monitoring, and review path. The site still needs to sponsor the access because the site knows when it is operationally necessary, who requested it, and what system it touches.

Who owns the enterprise vendor-access standard.

Who sponsors the access locally.

Who approves exception cases.

Who reviews repeated patterns across sites.

Who owns contract and access-model drift over time.

What a useful OT RACI sounds like in practice

Forget the matrix for a second. Listen to how the organization talks. A weak RACI sounds like this:

> "Security owns standards. Engineering owns systems. Operations owns uptime."
>
> True. Useless.

A better RACI sounds different:

> **Decision-Rights Language in Practice**
>
> "Corporate security owns the remote access standard. The site system owner owns the implementation path. If the site cannot meet the standard before the planned outage, the system owner documents the constraint and compensating controls. The plant leader and OT security lead review it. If the exception crosses the materiality threshold, the named enterprise risk owner decides whether to accept, time-bound, or escalate it."

Real governance sounds specific because real decisions are specific. That is the standard manufacturing CISOs should bring into the room.

What boards and executive teams should ask

Who approves a material OT exception.

Who has authority during a production-impacting incident.

Who owns the enterprise standard when plant constraints push back.

Who owns vendor-access governance across sites.

Who decides whether a repeated local workaround has become enterprise risk.

If leadership cannot get crisp answers, the board should assume the OT RACI is not finished. That is not a paperwork problem. It is a risk-management problem.

★ YOUR MOVE

Stop asking whether the RACI has enough names on it. Ask whether it can answer five hard decisions under pressure: who owns the standard, who owns the site plan, who approves the exception, who has incident authority, who governs vendor access. If your current model cannot answer those clearly, fix that before you add more governance machinery. Plants do not need more boxes on a chart — they need clear answers before the line is down and the vendor is on the phone.

CHAPTER 5

OT Exceptions, Vendor Access, and Connectivity Expose Whether Governance Is Real

5

OT Governance for Manufacturing CISOs



Do not judge an OT program by its slide deck. Judge it by what happens when the plant cannot patch, the OEM needs remote access, and the shutdown window is gone. That is where the truth shows up.

Most OT organizations say they have governance. Plenty of them do. On paper. They have standards. They have steering meetings. They have roles. They have architecture diagrams. They have exception forms. They have somebody who can explain the target state in a conference room.

Then a fragile controller reaches production season without a patch path. A vendor needs remote access for a recovery step nobody planned for. A line expansion adds connectivity faster than governance can catch up. A compensating control exists in the ticket and nowhere else. That is when you find out whether the organization has a governance model or a filing system.

OT exceptions should not live in email, local custom, or tribal knowledge. The decision not to update is a business risk decision that should be recorded in the wider risk-management framework. Connectivity decisions should tie to business-case logic, acceptable risk, and a named senior risk owner.

A plant can claim to follow the standard. A steering committee can claim to oversee the standard. A RACI can claim to assign the standard. The exception process tells you whether any of that survives contact with the real operating environment.

OT exceptions are normal

This needs to be said first because too much OT writing still talks as if exceptions are a sign of failure by themselves. They are not. Legacy systems are normal. Fragile assets are normal. Production timing is normal. OEM dependencies are normal. Safety constraints are normal.

A plant that cannot patch a controller during peak run may be making the right local call. The existence of the constraint is not the indictment. The governance question is what happens next. Does the issue get documented. Does someone own the risk. Does the exception have a review path. Do repeated patterns rise above the plant.

Normal does not mean informal

A lot of manufacturers treat OT exceptions as operational facts instead of governed risk decisions. The plant says the patch cannot happen. The OEM says remote access has to stay open for support. Security says the issue is understood. Then the organization moves on. That is not governance. That is a verbal accommodation.

Non-updating is not merely a technical delay. It is a decision with business ownership and traceable consequence. The same logic applies more broadly. If a vendor needs a remote path that violates the normal access model, that is a trust-boundary decision. If a modernization project adds connectivity between systems that used to be isolated, that is an enterprise risk decision. If a plant cannot follow the segmentation timetable, that is a risk acceptance decision waiting for somebody to name it.

What a bad exception process looks like

Local email chains. The exception lives in messages between a plant engineer, a security lead, and a vendor contact. Six months later nobody can say who approved it, which compensating controls were supposed to exist, or when the issue was supposed to come back for review.

Undocumented vendor access. The OEM still connects because operations needs them. Nobody can explain whether the enterprise recognizes the same pattern across the other sites that use the same vendor.

No named risk owner. Everyone is involved. Nobody is accountable. Security knows the issue. Engineering knows the constraint. Operations knows the business impact. No one owns the acceptance of the residual risk.

Compensating controls that exist in the ticket but not in the plant. A document says the local firewall rule was tightened or remote access is limited to specific windows. The configuration never

changed, or the local team assumed someone else handled it.

Repeated exceptions with no enterprise view. Plant 2 has the issue. Plant 5 has the same issue. Plant 9 uses a slightly different explanation for the same issue. Nobody aggregates the pattern. That is not dramatic failure. It is drift.

What a credible exception process looks like

A credible exception process is not glamorous. It is disciplined. Every exception needs:

A **named requestor** — someone who puts the issue on paper with the real constraint, not vague phrases.

A **named risk owner** — not a broad committee. A person or clearly defined role with authority to accept, reject, time-bound, escalate, or require additional controls.

A **business justification** — specific consequence: missed production, unsafe timing, lack of validated restart support, whatever the real reason is.

A **clear constraint description** — 'Plant cannot patch right now' is not a description. 'The line runs continuously through the quarter, the OEM requires onsite validation, and the next safe outage is week 42' is.

Defined **compensating controls** — real controls, not aspirations. If logging increases, say who reviews it. If remote access is limited, say how.

A **review date** — because open-ended exceptions turn into permanent architecture by neglect.

An **escalation threshold** — cross-site repetition, high-hazard process impact, third-party concentration, material business consequence.

Vendor access deserves its own section because it breaks bad governance first

Third-party OT access sits right on the fault line between plant reality and enterprise control. Weak governance lets local urgency win every time. The result: one site sponsors access one way. Another uses a different method for the same OEM. A third has an older arrangement nobody wants to touch. Local teams know the story. The enterprise does not know the pattern.

The enterprise should own the control model. The plant should own the local sponsorship and the business case. The exception process should connect the two. If the same OEM supports five plants, the company should not be learning about the access model one outage at a time.

Connectivity changes should trigger governance, not just design review

A modernization project adds a new data path. A line upgrade needs a remote support function. Everyone talks about architecture. Fewer people talk about ownership. Connectivity should tie back to business-case logic, acceptable risk, and named senior risk ownership. Connectivity is not just a design matter. It is a governance event.

Who asked for the connection.

What business value it serves.

What risk it introduces.

What control assumptions it depends on.

Who owns the residual risk.

How a CISO should use exception data

Exception handling is not only a control mechanism. It is an intelligence feed about how the OT program is actually working.

Repeated exceptions identify repeated control failures. If the same exception reason keeps appearing, the problem is probably architectural, budgetary, or structural.

Exception data exposes underfunded modernization work. If the organization keeps deferring controls because the underlying assets are too fragile, leadership should stop calling it a plant problem and start treating it as a capital and lifecycle problem.

Exception data surfaces vendor concentration risk. If the same vendor access workaround appears across a cluster of critical sites, that is a trust-boundary pattern the enterprise should be governing directly.

Exception data shows where plant constraints are systemic. The plant gets blamed for resisting the standard. In reality the enterprise may have created the same impossible timing condition in multiple places.

The fastest way to tell whether OT governance is real

Look at the exception process, the vendor access model, and the last connectivity decision no one wanted to own. If exception patterns never reach enterprise governance, the board is seeing OT risk after the decisions that create it have already been made. That is too late.

The Full Argument, In One Paragraph

This series has argued five things: senior visibility matters because repeated local exceptions become enterprise risk; named leadership and system ownership matter because someone has to speak for the standard and someone has to speak for the system; cross-functional governance matters because security, engineering, operations, safety, and risk all shape the decision; decision-rights discipline matters because a real RACI has to answer who approves the exception, who controls vendor access, and who owns the risk when the standard breaks; and the exception process is where all of those claims either survive or collapse.

★ YOUR MOVE

Stop treating exceptions as paperwork. Treat them as evidence. They tell you where the standard does not fit the operating reality, where the enterprise keeps making the same compromise, where vendors are shaping your trust boundary more than your own controls are, and where a local plant issue has already become enterprise risk. Push for named risk owners. Push for real compensating controls. Push for cross-site visibility. Push for connectivity governance early, not after the design has hardened. Most of all, push for honesty.

ABOUT THE AUTHOR

Roger Hill

Founder & CEO, Hillstrong Group Security



Roger brings over 30 years of experience in industrial automation and cybersecurity. His extensive background in OT security has positioned him as a trusted advisor to manufacturing CISOs and senior operations leaders worldwide.

His expertise lies in developing and implementing robust OT GRC programs that address the unique challenges faced by global manufacturing operations. Throughout his career, he has been dedicated to helping organizations achieve operational resilience through strategic security initiatives and pragmatic solutions.

Roger's methodologies continue to drive the development of secure and compliant OT environments for leading manufacturing companies across the automotive, food and beverage, chemical, and pharmaceutical sectors.

www.hillstrongsecurity.com | Atlanta, GA

About Hillstrong Group Security

Hillstrong Group Security is an OT and ICS cybersecurity consultancy specializing in hands-on security assessments, governance program development, and operational resilience for global manufacturers. Unlike remote-only security firms, Hillstrong consultants physically assess OT environments on the plant floor — bringing operational credibility to every engagement. Hillstrong serves clients across the automotive, food and beverage, chemical, pharmaceutical, consumer goods, and logistics sectors.

Our Core Services

- Resilion.io OT Cybersecurity ProgramOps
- OT Security Maturity Assessments
- OT Risk and Vulnerability Analysis
- OT Program Design and Security Architecture
- Expert-Level OT Security Engineering
- OT Staff Augmentation and Contract Specialists
- Fractional CISO and Executive Advisory Support
- OT Governance, Risk, and Compliance as a Service

Resilion.io

Resilion.io is Hillstrong's proprietary AI-powered platform supporting continuous OT risk visibility and program management for enterprise manufacturing clients. Resilion transforms OT cybersecurity from a project into an ongoing operational capability — giving CISOs and plant leadership a single system of record for program execution, risk tracking, and governance accountability. Learn more at **resilion.io**.

www.hillstrongsecurity.com | Based in Atlanta — Solving Cybersecurity Challenges Globally