resillion

The 48-hour OT vulnerability decision playbook

From OT vulnerability to reportable incident: a decision workflow for the first 48 hours after a new OT vulnerability is identified, designed to work alongside regulatory reporting timelines.

Part 1: Why this matters (and why timelines are tightening)

The hardest part of OT vulnerability management isn't finding vulnerabilities. It's deciding what to do next.

Severity scores like CVSS are designed to describe technical risk.
They don't describe operational reality.

They can't tell you:

- whether a vulnerability is likely to be exploited in your environment
- whether exploitation would disrupt operations or safety
- or whether immediate action would reduce risk – or make things worse

That uncertainty drives two familiar responses:

Patch panic, where everything is treated as urgent
Patch paralysis, where nothing moves at all

In OT environments, both lead to unnecessary risk.
What teams need isn't more alerts. They need a clear, defensible decision process.

Why 48 hours matters now

Across Europe, multiple regulations can impose formal reporting and decision timelines when vulnerabilities emerge or serious incident thresholds are met.

These timelines may range from as little as 4 hours to the more common 24–72 hour notification, depending on the circumstances.

So even when you start in 'vulnerability mode', you may quickly need to operate in 'regulatory clock mode' if the situation meets a reporting trigger.

In practice, that means within the first 4–24 hours you need to be able to produce (and evidence) a few minimum decision outputs:

- a defensible view of **scope** and **potential impact** (your 0–4 hour output.)
- a clear, proportionate **response choice** under uncertainty - monitor, reduce exposure, fix, contain (often a useful 4–24 hour output).
- a contemporaneous **decision record:** what you knew, what you assumed, what you did, and when you'll update (good practice within 24–48 hours).

How to adapt this playbook for different regulatory timelines

Different regimes use different triggers and start times. Don't solve that by building separate processes. Keep one operational workflow and add a simple layer to track multiple reporting clocks as the situation becomes clearer.

Use these rules of thumb:

- **Set 'Awareness time' (T0)** once and record it. Use it as your internal anchor even if some regimes key off "classification".
- **Track two things in parallel:** the decision workflow stage (0–4 / 4–24 / 24–48) and which regulatory triggers might apply based on what you know right now.
- **Default to the shortest credible deadline** until you can rule a regime out. This protects you when classification changes mid-stream.
- **Treat early reporting as iterative:** notify what you know, state assumptions, and commit to updates. Don't wait for a perfect root cause.
- **Separate owners:** one accountable for the reporting clock (legal/compliance/DSO) and one for the technical workflow (security/engineering).

Regulatory clocks (quick reference) *

When vulnerability or incident handling may transition into regulatory reporting

Regime	Trigger (simplified)	Key reporting timeline (from awareness / trigger)
NIS2	Significant incident (essential/important entities)	Early warning: within 24h of awareness Notification: within 72h Final: no later than 1 month after the 72h notification
DORA	Major ICT-related incident (financial sector)	Initial: within 4h of classification as major (and no later than 24h after awareness) Intermediate: within 72h of initial Final: no later than 1 month after intermediate
GDPR	Personal data breach with risk to individuals	Authority: without undue delay, where feasible within 72h of awareness Data subjects: without undue delay if high risk
Cyber Resilience Act (CRA)	Actively exploited vulnerability or severe incident (manufacturers)	Early warning: within 24h of awareness Notification: within 72h Final: within 14 days (vulnerability once patch available) or 1 month (severe incident)
EU AI Act	Serious incident (high-risk AI systems)	Standard: immediately and no later than 15 days Death/ health harm: no later than 2 days Widespread infringement/critical infrastructure: no later than 10 days

*These timelines do not apply to all vulnerabilities. They become relevant only where the legal trigger conditions are met (for example, based on exploitation, impact, or how the event is classified). Always verify the applicable trigger, start time, and deadline for your organisation and incident.

Practically: the 0–4 hour stage can help you quickly assess whether you might be in scope for an early-warning deadline (e.g., NIS2/CRA) or a fast 'classification then notify' obligation (e.g., DORA).

The 4–24 hour stage can help you build a decision record that supports timely escalation and, where needed, a notification draft for legal/compliance review (e.g., NIS2/GDPR/CRA) and an initial DORA report if applicable.

The 24–48 hour stage should focus on stabilising risk and locking the narrative: what you did, why you did it, what you don't yet know, and when you will update.

From analysis to decision

Regulatory timelines add urgency, but they do not change the nature of the decision that needs to be made. Whether a vulnerability ultimately qualifies for regulatory reporting or not, effective handling still depends on answering three key questions:

- What would the impact be if this was exploited here?
- How likely is that exploitation in reality?
- What does a plausible attack path look like in our environment?

Answering those questions consistently turns vulnerability management from debate into discipline.

Part 2: The 48-hour decision workflow

This playbook is designed for **real operational environments,** where patching is constrained and decisions must balance security and uptime. (For the regulatory context and timeline differences, see Part 1.)

It is not a maturity model or a compliance checklist.
It's a **decision workflow** teams can apply consistently when a new vulnerability appears - producing a record that stands up to operational, executive, and (where applicable) regulatory scrutiny.

0–4 hours: Intake and scoping

OBJECTIVE: Confirm relevance and exposure, not fix everything immediately

Focus on establishing technical facts:

- Register the vulnerability in the risk register or tracking system
- Confirm whether the affected product, component, or firmware exists in your environment
- Identify where affected systems sit in the architecture:
    - IT / OT boundary
    - Safety critical zones
    - Control plane components (gateways, jump hosts, historians, remote access)
- Assess basic exposure:
    - network reachability
    - authentication requirements
    - existing segmentation and monitoring

KEY QUESTION:

Is this vulnerability theoretically relevant, or practically reachable in our environment?

DELIVERABLE:

A short **impact and exposure statement,** with assumptions clearly recorded.

4–24 hours: Risk determination and attack path analysis

OBJECTIVE: Decide whether this is a monitoring issue, an exposure reduction issue, or a fix now issue

At this stage, teams combine multiple inputs:

Severity and criticality

- CVSS or vendor severity as a starting point
- Operational criticality of the affected system
- Consequences of compromise (loss of view, loss of control, loss of recovery)
  Answering those questions consistently turns vulnerability management from debate into discipline.

Exploitation signals

- Evidence of exploitation in the wild
- Public advisories or alerts indicating active abuse
- Whether similar vulnerabilities are commonly weaponised

Credible attack path Rather than abstract risk, teams map a plausible sequence:

- How initial access would occur
- What privileges the vulnerability provides
- How an attacker could move laterally
- What systems could realistically be impacted next

This does not require deep threat modelling. A simple "if this, then that" chain is sufficient.

Detection implications

- What signals would indicate exploitation?
- Are relevant logs already collected?
- Would alerts be distinguishable from normal operations?

DELIVERABLE:

A decision ready risk entry, including:

- urgency level
- rationale
- recommended response class

24–48 hours: Mitigation and execution

OBJECTIVE: Reduce risk immediately without introducing instability
When immediate patching is not safe or feasible, teams prioritise risk suppression:

Technical mitigations

- Restrict or disable exposed management interfaces
- Reduce reachable network paths through segmentation
- Tighten authentication and access control scopes
- Disable unused services or functions
- Apply rate limiting or protocol restrictions where possible

Detection improvements

- Add or tune alerts based on the identified attack path
- Validate that detection actually triggers
- Ensure logs are retained and reviewable

Patching strategy

- Plan patching with testing and rollback
- Align with maintenance windows
- Treat patching as a controlled change, not an emergency reflex

Communication and governance

- Record the decision and mitigations implemented
- Communicate clearly with operations and engineering
- Set a defined review point

DELIVERABLE:

Implemented controls, updated risk record and a documented next step.

Common failure patterns to avoid

- **Treating severity as urgency**
  High severity does not automatically justify disruptive action.
- **Waiting for "perfect certainty"**
  Delayed decisions often create more risk than imperfect ones.
- **Relying on a single data source**
  Scores and feeds are inputs — not decisions.
- **Skipping documentation**
  Undocumented decisions don't survive audits, incidents, or staff changes.

How Resillion helps

We help OT and hybrid environments apply this playbook in practice by:

- tailoring the 48-hour workflow to your architecture
- identifying high risk attack paths specific to your environment
- defining compensating controls where patching is constrained
- helping teams document decisions in a defensible way

Regulatory note:

This document provides general information and an example operational workflow. It is not legal advice. Regulatory obligations, triggers, and reporting timelines vary by jurisdiction, sector, entity classification, and competent authority guidance, and may change over time. Always confirm applicability and deadlines with your legal/compliance function and relevant regulator guidance for the specific incident.

Any support services described are intended to help operationalise security decision-making and documentation. They do not, on their own, guarantee regulatory compliance, and should be used alongside your organisation's legal/compliance review and incident reporting processes

Get in touch

Request an OT vulnerability decision workshop
Get a 48-hour playbook tailored to your environment
✉ | Email us at: hello@resillion.com