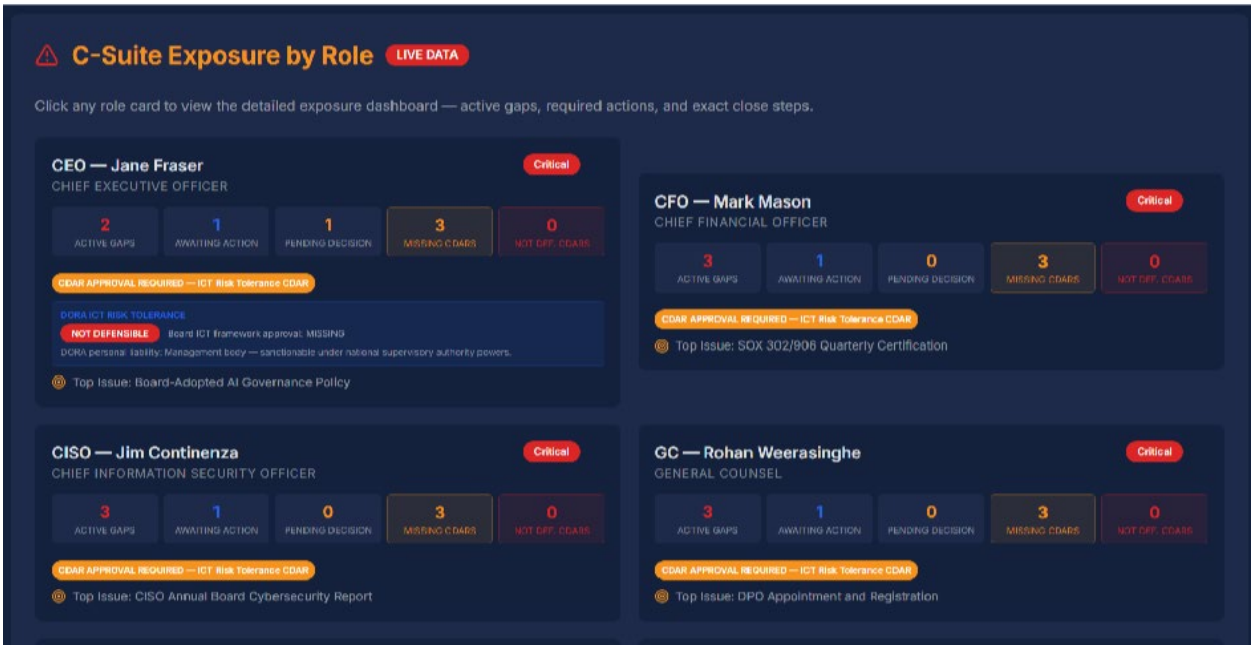


## The Threshold Nobody Set: Why Legal Defensibility Requires Infrastructure

**Geoff Hancock's** May 2026 **article on CISO accountability** is the best description of the problem I have read. I want to start there, because I think he is right — and because what I am going to add is not a correction. It is the part that comes next.

### Hancock is right about the operating model

Hancock's diagnosis is precise: the modern **CISO** is being **held accountable** like an **executive** but **positioned** like **a technician**. The **title implies ownership**. The **org chart**, the **budget authority**, and the real **decision rights rarely deliver it**. Then something goes wrong, and the **accountability turns** out to be entirely **real** — and entirely **personal**. The **authority** that should have come with it **was never there**. That is a broken operating model, and he names it cleanly.

His prescription is the right one. **Define** precisely **what security owns**. **Formalize** how **risk decisions** actually **get made**. Assign execution accountability where it truly sits, not where the title suggests it should. Put issues in front of the board in business language rather than technical language. Every CISO should read the piece, and most should act on it.

But there is a **gap** between a better **operating model** and **legal reality** — and that gap is exactly where companies, and the executives who run them, **get into trouble**.

**A better operating model still does not create legal evidence**

Here is the uncomfortable part. You can implement every move Hancock describes and still be undefensible.

Hancock says to formalize risk decisions. He is right that you should. But "formalized" and "legally sufficient" are not the same standard. Regulators and courts do not ask whether a company had a process for making risk decisions. They ask something far more specific: did a **named**, **accountable person decide** — **before** any **harm occurred** — **how much risk** this organization was willing to **impose on other people**, and is there a **contemporaneous**, dated **record** of that **decision** and the **reasoning behind it?**

That is a **different artifact** than anything most **security programs produce**. It is **not** a **risk-appetite** statement. It is **not** a **heat map**. It is **not** a **maturity score** or a policy document. It is a specific, **calculated threshold**: this is the level of risk of harm we are knowingly accepting on behalf of our customers, our partners, our patients, our users — **here is how we reasoned** to it, here is **who signed it**, and here is **when**.

I call that decision the CDAR, the **Calculated Definition of Acceptable Risk**. The term is mine. The legal requirement is not. Under **GDPR Article 32**, **DORA**, **NIS2**, the **EU AI Act**, and **child-safety statutes** such as KOSA, **organizations** are already **required** to **set** and document explicit **thresholds** for the **risk** of harm they **impose** on **others** — per stakeholder type, per regulation. Not in the abstract. Specifically.

Make it concrete. A social platform that admits thirteen-year-olds has a legal obligation to decide, in advance, how much risk of harm to those minors it considers acceptable, and to be able to show the calculation. A bank deploying a new AI model that touches customer money carries the same kind of obligation under a different set of rules. In practice, almost none of these organizations hold a dated record of that specific decision. They have policies. They have controls. They have a risk register. What they do not have is the one artifact a regulator asks for first.

**Most companies** have **never made** that decision — not because they are negligent, but because **no one told them** the **law** now **required it,** and **no process** or software **ever walked them through it**.

**Why this is an infrastructure problem**

This is where I want to extend Hancock rather than argue with him.

An operating model is a set of agreements about who does what. It is necessary. But the threshold decision the law requires cannot be produced reliably with better agreements alone, for three reasons.

First, it **requires applying** a **legal test**, not a business judgment. Courts evaluating reasonableness apply specific standards — the Hand Rule in US negligence law, "state of the art" and adequacy under GDPR, six-factor practicability tests in other jurisdictions. A CISO and a General Counsel improvising in a conference room are not applying those tests in any consistent, repeatable way. They are reasoning by experience and instinct, which is not the same thing, and will not read the same way under examination.

Second, the **decision** has to be **made** per **stakeholder** and per **regulation**, and then kept current as both change. The acceptable risk you may impose on a sophisticated enterprise customer is not the acceptable risk you may impose on a thirteen-year-old, and neither is static. A single annual risk statement cannot carry that weight, and an organization operating across the SEC, a state attorney general, and an EU regulator at the same time needs one defensible standard, not five disconnected ones.

Third — and this is the reason that matters most — the **evidence** has to be **created** at the **moment of decision** and **sealed** so that it cannot be quietly revised afterward. **Reasonableness** is **judged** on what **you knew** and **decided** then, not on what you can assemble later. A record that could have been edged into shape after the fact is not evidence. It is an argument.

No operating model produces that on its own. It requires infrastructure: something that **guides** the threshold **decision**, **applies** the relevant **legal test**, and **captures** the **result** as dated, attributed, **tamper-proof evidence** — **every time** a c**onsequential decision** is **made**, not on the occasions someone remembers to.

**Evidence has to exist before the incident**

Consider what actually happened at SolarWinds. When the SEC brought its case, the question was never whether the company had security controls. It was whether anyone could show why leadership had accepted the risks that had been raised internally. The controls existed. The decision rationale did not — or at least could not be produced in a form that answered the question being asked.

That is the **pattern underneath** the great **majority** of significant **enforcement actions** in cybersecurity and privacy over the past decade. The failure is rarely the absence of controls. It is the absence of a **documented**, **contemporaneous** decision showing that a named executive weighed a known risk and judged it reasonable before anyone was harmed.

You **cannot retrofit** that. The day an incident becomes public is the day it is too late to create the evidence. Either the **threshold was set**, dated, and signed beforehand, **or it was not**. There is **no third option**, and **no** amount of **post-incident effort manufactures** one.

**The operating model gets you the seat. Infrastructure proves the decision.**

So here is how I would put it.

Hancock's work fixes who is in the room and what authority they carry when they get there. That is real, and it is overdue. But sitting in the room with the right authority is not the same as being able to prove, two years later and under examination, that the decision made in that room was reasonable.

The first is an operating-model problem. The second is an infrastructure problem. They are not in competition; a CISO needs both. The operating model gives the CISO a defensible position inside the company. The infrastructure gives the company — and that CISO personally — a defensible position in front of a regulator, a plaintiff's counsel, or a prosecutor.

Read Hancock. Fix the operating model. He has described the seat at the table more clearly than anyone in this field. Then ask the next question, the one his framework does not reach: when it matters most, can you actually prove what was decided at it?

If the answer is no, a better operating model has not solved your problem. It has only made it easier to see.