INNOVATION

Human Error To Architectural Failure: Redefining Cyber Accountability

By Federico Simonetti, Forbes Councils Member.

Xiid's CTO, Federico Simonetti, is a security expert with a background in ethical hacking, law enforcement and academia.

After every major breach, the same phrase shows up in the postmortem: human error. Someone clicked the wrong link. An admin misconfigured a setting. A contractor reused a password.

On the surface, this sounds like accountability. In reality, it's a tell. It signals that the organization is still treating cybersecurity as a behavior problem rather than a design problem.

Boards no longer have that luxury. Cyber risk now sits alongside financial risk and regulatory exposure as a core governance responsibility. Regulators and investors expect directors to demonstrate real oversight of technology risk, not just sign off on budgets and hear an annual briefing.

If your security strategy still depends on people consistently doing the right thing in complex, high-pressure environments, the problem is not your people. It's your architecture.

Human Error Is Predictable; Architecture Is A Choice

Analyses routinely show that the majority of breaches include a human factor. Think of misdirected emails, credential reuse, misconfigurations or a convincing phish. But those behaviors are inevitable. Employees choose speed when deadlines loom. Developers use tools that help them ship faster. Executives approve exceptions that keep revenue flowing.

If you know a behavior is inevitable and still design your defenses around preventing it, that's not "user error." That's a design decision. Most programs are built on probabilistic controls that help but depend on human attention, memory and judgment, including:

- Awareness campaigns
- Simulated phishing
- Policy reminders
- Anomaly monitoring

Deterministic controls work differently. The risky action simply cannot happen. The system doesn't present a directly reachable path to it. The critical application is not on a routable network segment. The privileged access path does not rely on long-lived credentials that can be reused if leaked.

From a governance standpoint, that distinction matters.

Three Places Architectural Risk Hides

For most organizations, architectural risk accumulates in three places: reachability, standing access and segmentation.

1. Reachable Systems

If your crown-jewel systems are directly reachable from the public internet—via open ports, static endpoints or broad VPN access—it only takes one stolen credential, one successful phish or a zero-day exploit for an attacker to walk in. CISA's analyses of the most exploited vulnerabilities highlight the pattern: Internet-exposed services with known weaknesses become a systemic risk, not just IT cleanup.

Boards should ask:

• Which revenue-critical or safety-critical systems can be probed or scanned from outside the organization?

• How many third parties provide indirect reachability into those systems?

If the answer is “most of them” or “we’re not sure,” training isn’t your main problem. Exposure is.

2. Standing Access And Static Secrets

Long-lived credentials and broad admin rights turn minor mistakes into major incidents. If a developer accidentally commits an access token or pastes a credential into an external AI tool, an attacker can replay those secrets at scale. The governance question should be: Have we designed our systems so that a leaked credential has very little value or doesn’t exist in the first place?

Short-lived, narrowly scoped access and credential-less authentication patterns dramatically shrink the blast radius of inevitable slips.

3. Flat Or Coarse-Grained Segmentation

Even when initial access is small, flat networks and coarse segmentation let attackers roam. If the compromise of a single workload gives access to an entire subnet or environment, you’ve traded one user mistake for an organizational incident.

Boards should look for evidence of segmentation at the process or service level, what some teams call pico segmentation, rather than just VLANs and high-level zones.

Three Board-Level Questions That Change The Conversation

Directors need sharper questions that pull the discussion up from “awareness” and down into “how we actually built this.”

1. Where do we still depend on humans never making a mistake?

Ask your CISO to identify the top scenarios where a single action, like clicking a link, misclassifying a system, reusing a password or approving an exception, could still result in material impact.

Then ask: What deterministic controls could make those actions non-catastrophic?

Examples include moving critical systems off addressable networks and into non-routable, out-of-range exposure patterns or replacing static credentials with ephemeral, policy-driven access that cannot be reused indefinitely.

2. How much of our attack surface is still reachable today, and is that number shrinking?

Dashboards tend to emphasize vulnerability counts and patch status. Boards should also analyze reachability: how many services tied to critical processes can be touched from outside, and how that trend is changing over time.

At minimum, request these recurring metrics:

• Count of internet-reachable services associated with crown-jewel systems

• Number of partners with network-level access into core environments

If those numbers aren't heading down quarter over quarter, you're not truly de-risking the architecture.

3. If an adversary harvested our traffic today, what would still matter five years from now?

Concerns about "harvest now, decrypt later" quantum scenarios have pushed many organizations to revisit their cryptography, especially for long-lived data.

Boards should extend that view to architecture:

• Where is sensitive data decrypted within our environment?

• Where do intermediaries terminate and re-encrypt traffic, creating extra exposure?

The goal is not just stronger algorithms; it's fewer places where high-value data exists in the clear and fewer entities that can access it, even if today's protections weaken over time.

Redefining Accountability From The Top Down

Blaming human error gives boards an easy out. If the problem is people, the solution is more reminders.

Training and culture both matter, but neither is a substitute for an architecture that's designed to bring human error to nuisance level instead of headline level.

In a world where AI accelerates both attack and defense and regulators are asking tougher questions about cyber governance, accountability has to move upstream, from the last person who clicked to the decisions about what was reachable. The next time human error appears in a post-incident briefing, treat it as a starting point, not a conclusion.

Ask: What in our architecture made that mistake matter so much?

Boards that get comfortable asking and acting on that question will reduce breach impact and set a new cyber accountability standard that's grounded in design, not blame.

By Federico Simonetti, COUNCIL POST | Membership (fee-based). Xiid's CTO, Federico Simonetti, is a security expert with a background in ethical hacking, law enforcement and academia. Read Federico Simonetti's full executive profile here.
Find Federico Simonetti on LinkedIn. Visit Federico's website.