

DECISION MAKING INSIGHT SERIES

The future challenge requires something different — a living system, not a chart.

Article 1 of 3 — Decision Accountability

Decision Architecture I — The Shape: Gravity Over Hierarchy

Why decision making needs to evolve

BeyondRM | 2026 | Article 1 of 3 | Decision Making Insight Series

A. Clements | D. Lengyel | M. Pergler | D. Roberts | S. Crosina

The Problem This Article Addresses

The management architecture that dominates most organisations today — whether expressed as a pyramid hierarchy or its attempted fix, the matrix model — was built to answer one question: who reports to whom? Both structures assume that authority flows downward and that the organisation's primary challenge is the coordination of known tasks together with modest adaptation to evolving conditions. In an environment of known tasks and generally stable conditions, these are reasonable design choices.

But the future challenge is qualitatively different. It requires attracting unknown opportunities, sensing novel threats before they become crises, and preserving decision-making quality at every layer of the organisation — not only at the top. It requires speed at the point of awareness, genuine empowerment of people with the right skills, and a clear, open, and transparent method of communication that does not lose critical information as it passes upward through layers of reporting. It requires opportunity capture and downside risk mitigation be integrated, and that decision making, ongoing risk management, and resilience operate in a tight loop

Summary

We are at an organisational inflection point. For most of the past century, companies have operated inside a period of relative stability — one punctuated by crises, certainly, but structured around the assumption that the future would broadly resemble the past.

Much of this stability has a name: the Great Moderation, the roughly quarter-century stretch from the mid-1980s to the 2007–08 financial crisis during which growth cycles flattened, inflation was tamed, and macroeconomic volatility fell to historic lows across the developed world. It was during this period that the pyramid hierarchy and its matrix variant were refined into their current form — layered approval chains, standardised reporting cycles, and

centralised risk committees were rational, even optimal, responses to an environment in which the future genuinely resembled the past closely enough to plan around. The architecture was not poorly designed; it was well designed for the conditions of its era. What has changed is the era, not the soundness of the original design logic — and we are now well beyond the Great Moderation, with structures still being asked to perform in conditions they were never built to handle.

That assumption no longer holds. The volatility and uncertainty that characterised the early decades of this century are not a temporary disruption to be managed through. They are the new operating environment.

Two frameworks help make this operating environment concrete rather than abstract. The first, VUCA — Volatility, Uncertainty, Complexity, and Ambiguity — originated in U.S. military strategic thinking at the end of the Cold War and has since become shorthand across industry for conditions in which change is rapid and unpredictable, cause and effect are difficult to trace, multiple interdependent variables interact, and the meaning of events is genuinely open to more than one interpretation. The second, the Cynefin framework, offers a sharper diagnostic: it distinguishes Clear, Complicated, Complex, and Chaotic decision domains, each calling for a different mode of sense-making. Pyramid and matrix architectures were built to excel in the Complicated domain, where expert analysis can identify a correct answer through best practice. But a growing share of the decisions organisations now face sit in the Complex domain, where cause and effect are only visible in hindsight and the right response is to probe, sense, and respond rather than analyse and plan. An architecture calibrated for Complicated problems will systematically misjudge Complex ones — not through incompetence, but because it applies the wrong logic to the wrong domain.

The question this raises is not simply how to improve risk management. It is more fundamental: can an organisation that was designed in the nineteenth century — and refined through the twentieth — make good decisions at the speed and complexity that the twenty-first century now demands? The answer, drawn from over 47 interviews with senior aerospace and defence executives across the United States and Europe, from survey data spanning more than 2,000 risk professionals, and from first-principles architectural thinking, is no.

Three structural ideas redefine the architecture

Gravity replaces hierarchy. Value and information flow toward wherever they are most needed. Teams are semi-autonomous but orbit shared gravitational centres — strategic intent, capability pools, and shared intelligence.

Two operating tempos coexist. A stable core (governance, capability stewardship, institutional memory) and a fluid periphery (project teams, opportunity-response units,

futures-sensing nodes) that can form, deliver, and dissolve without destabilising the centre.

Risk and opportunity: Should share a sensing and interpretation infrastructure, while retaining different consequence thresholds, ethical boundaries, and acceptance authorities.

1. The Problem with Existing Risk Management

It's not a problem with risk management as such. It's more that current management architecture is still in the 19th century. It's certainly not fit for purpose for tomorrow.

The pyramid hierarchy is a reporting structure masquerading as a decision structure. The matrix model was an attempt to fix the pyramid but created two hierarchies instead of one. Both assume that authority flows *downward* and that the organisation's primary challenge is *coordination of known tasks*.

The future challenge is different: it's about attracting *unknown opportunities*, sensing *novel threats*, and preserving *decision-making quality* at every layer — not just the top. That requires a fundamentally different architecture. It requires speed at the point of awareness, empowerment from those people with the right skills and a clear, open and transparent method of communication.

The model that emerges from first principles looks more like a **living system** than a chart. Three structural ideas define it:

Gravity replaces hierarchy. Instead of authority flowing down from a peak, *value and information* flow toward wherever they're most needed. Teams are semi-autonomous but orbit shared "gravitational centres" — strategic intent, capability pools, and shared intelligence.

Two operating tempos coexist. A stable core (governance, capability stewardship, institutional memory) and a fluid periphery (project teams, opportunity-response units, futures-sensing nodes) that can form, deliver, and dissolve without destabilising the centre.

Risk and opportunity. Attracting opportunities and managing risks aren't separate departments — they're two outputs of a single strategic sensing capability embedded throughout the organisation, not siloed at the top.


Risk and opportunity can share:

- environmental sensing;
- strategic intelligence;
- scenario development;
- option analysis;

- resource allocation discussions.

But they do not always share the same ethical or decision standards. An opportunity may be declined because it lacks sufficient upside. A safety or human-life risk may be prohibited because it crosses a non-negotiable boundary, regardless of upside.

Here's how that structure could be visualised:



The diagram shows four concentric zones. It is worth being precise about where this model sits relative to what already exists: it is not a wholesale replacement of the existing management structure but a concept operating at its edge — at the boundary where a new architectural logic meets established institutional form. Existing structures are not simply discarded; they are adapted, zone by zone, so that governance bodies, reporting lines, and capability functions that already work can be repositioned within the new logic rather than rebuilt from nothing. Each of the four zones carries a distinct logic:

The strategic core is small and deliberate. It holds purpose, values, long-horizon intent, and the governance architecture. It doesn't manage day-to-day — it sets the gravitational field that everything else orients around.

The capability layer is the organisation's permanent infrastructure — not departments in the traditional sense, but standing pools of human judgment, knowledge, and decision quality. Future Capability/Foresight, Decision Quality/Escalation Architecture, Resilience Base, and Capability Pools are always there, always available to any team that needs them.

The dynamic periphery is where actual work happens. Teams form around opportunities and risks, draw from the capability layer, operate for a defined period, and dissolve or evolve. No permanent functional silos. Each team has a clear mandate, a time horizon, and a direct line to capability resources. The org chart here is temporal — it shows *what exists now*, not what will always exist.

The sensing membrane is the outermost active layer. Rather than waiting for the market to hit the organisation, this layer is continuously translating weak signals from six domains — technology, regulation, stakeholders, competition, talent, and capital — into inputs that shape which teams form in the periphery.

The critical design principle that makes this work is that **risk and opportunity share the same sensing and decision infrastructure**. There is no separate "risk department" doing compliance work while strategy happens elsewhere. Every signal that enters through the membrane is assessed for both its threat dimension and its option-creating dimension simultaneously. However there is the need for coaching of opportunistic management processes, aggregation and simulations to assess actual potential impacts ranges and the loyal opposition aspect to embed thoughts, solutions and decisions.


This is neither pyramid, matrix, nor pure network. It's closer to what biologists call a **living system architecture** — a stable centre with a permeable, adaptive boundary. The pyramid answered the question "who reports to whom." This structure answers the question "how does good judgment get to the right place at the right time."

The Future-Capability Architect role, a future oriented performance enhancement advisor, or **risk/opportunity manager**, sits at the boundary between the strategic core and the capability layer — holding the dual-track distinction structurally, not just by personal commitment. The Futures Intelligence Lead role operates at the boundary between the capability layer and the dynamic periphery — ensuring that foresight informs team formation rather than arriving as a post-hoc report.

This architecture is driven from the perspective of capabilities. However, A capability without an accountable owner is just potential energy; it never converts into action. We therefore need to include a **decision architecture**. This is key as, from a purely risk management perspective, we have seen that decision ownership and responsibility are holding many companies back from attaining their full potential. Specifically, around opportunities and its possibly one reason why many risk management functions focus on risks rather than opportunities.

A sound decision architecture distributes sensing, expertise, and challenge without diffusing accountability. Those closest to the evidence must be empowered to identify, interpret, and contest risk. The authorised decision-maker must then state the governing assumptions, identify any ethical or safety boundaries, balance the relevant trade-space, accept the residual risk explicitly, and remain accountable as conditions change. Authority may be exercised at different organisational levels, but it must always be visible, commensurate with the consequences, and connected to the knowledge on which the decision depends.

If we rebuild this architecture by placing authority and accountability at the centre.



The redesign makes three structural arguments that the previous diagram didn't:

Decision Authority & Accountability belongs at the centre, not at the top. In the pyramid it sits at the apex — distant, slow, and filtered by layers of reporting. Here it sits at the centre of a network that every function speaks to directly and receives direction from. Authority isn't inherited from rank; it's exercised through the function. However, there are multiple layers of authority:

- **Epistemic authority:** who possesses the relevant knowledge.

- **Decision authority:** who is authorised to choose.
- **Risk-acceptance authority:** who may accept the residual consequences.
- **Accountability:** who remains answerable for the outcome.

The architecture should distribute knowledge and challenge while making decision rights unmistakably explicit.

The two-directional spokes are load bearing. The solid arrows show signals and intelligence flowing *inward* to the authority centre. The dashed arrows show mandates, decisions, and risk acceptance flowing *outward* to functions. Neither direction works without the other — and crucially, the outer cycle (Sense → Interpret → Decide → Execute → Govern → Learn → back to Sense) continues even when the centre isn't actively directing. The centre *anchors* accountability; it doesn't have to supervise every step.

Multi-hat roles are structural, not incidental. The three role badges show what we need to have in practise and what is already present in leading organisational architectures: a Chief Engineer simultaneously occupies Sense, Interpret, and Learn; a CRO occupies Interpret, Learn, and Govern; a Program Manager occupies Decide, Execute, and Govern. This isn't role ambiguity — it's the actual architecture. The *function* is the unit of design, not the job title.

For risk managers interested in identifying opportunities for company strategic change and having that discussion with leadership and board teams this presents the following question: "is this person structurally positioned to accept residual risk and be accountable for the decision?" That's a governance design question as much as a capability question, and it's where most organisations are challenged. A risk manager can put this insight onto the table for discussion.

5. Conclusions

Risk management, as currently practised, is not the problem. The problem is the management architecture within which risk management operates — and to which it has adapted itself. An architecture built for stability, coordination, and top-down authority cannot generate the decision speed, the distributed empowerment, or the continuous sensing that an uncertain future requires. Improving the risk function within a fundamentally unsuitable structure is like installing better instruments in an aircraft that was not designed to fly.

What is needed is not a new risk framework, it's a new decision architecture — one that places authority and accountability at the functional centre rather than the bureaucratic apex, that treats risk and opportunity as two faces of the same sensing capability, and that recognises judgment, ethics, dissent, and the human experience of carrying uncertainty as structural elements, not cultural aspirations.

This decision architecture begets its tools and frameworks on which to build a more future oriented organisation that leverages processes, systems and people in a manner different from the focus today. Away from the “me” and “now” and towards the “us” and “tomorrow”.

The architecture proposed in this article is not a blueprint that any organisation can adopt wholesale. It is a design language — a set of principles from which specific architectures can be built for specific contexts. The strategic core will differ between a defence contractor, a financial services firm, and a public institution. The composition of the capability layer will reflect the nature of the work. The sensing membrane will prioritise different signals in different industries. But the underlying logic holds across all of them: gravity over hierarchy, function over title, accountability at the centre.

For organisations weighing where to begin, the starting point is not a reorganisation project but a design audit: mapping where epistemic authority, decision authority, and risk-acceptance authority currently sit, and where they diverge. Three moves tend to matter most in practice. First, name a functional centre — even a small one — with explicit, visible authority over the trade-space between risk and opportunity, rather than leaving that authority to emerge informally between committees. Second, build the sensing membrane before restructuring the periphery: an organisation that can detect weak signals but has not yet reorganised around them still gains more warning time than one that redesigns its org chart without improving what it senses. Third, treat the transition as iterative rather than wholesale — piloting the model within a single capability pool or business unit generates evidence and organisational muscle memory that a full-scale redesign cannot. None of this requires abandoning existing governance overnight; it requires building the new architecture alongside the old one until the new one earns its authority through demonstrated decision quality.

***"The profession has frameworks, registers, heat maps, and committees.
What it lacks is a coherent explanation of how an accountable leader
should interpret imperfect knowledge, protect dissent, recognise ethical
limits, and accept the remaining risk."***

The risk profession stands at a genuine fork in the road. One path continues to refine existing tools, standards, taxonomies, and software within architectures that were not designed to support the decisions those tools are meant to inform. The other path accepts that the profession's most important contribution is not a better register or a more sophisticated model — it is the embedding of decision quality, ethical clarity, and accountable ownership throughout the organisation, at every layer and in every function.

The organisations that will thrive in the decades ahead are not those with the best risk documentation. They are those that have built the capability to make better decisions, faster, at the point where it matters — and to be genuinely, structurally accountable for them.

Key Takeaways

01 - Management architecture is the real constraint — not risk management capability.

The pyramid hierarchy and matrix model were designed for a world of known tasks and stable conditions. Organisations that try to improve their risk and opportunity management without addressing the underlying architecture will find that better tools, frameworks, and processes are consistently absorbed and neutralised by structures that were never designed to support them. The first question is not "how do we improve our risk function?" It is "is our organisational architecture capable of making fast, accountable, empowered decisions in conditions of genuine uncertainty?" In most organisations, that question has not yet been asked.

02 -Decision authority and accountability must be structural, not assumed.

One of the most consistent findings is that risk management consistently fails not because of analytical inadequacy, but because of accountability diffusion. Risks are identified collectively; residual risk is owned by no one. Opportunities are discussed in strategy sessions; the authority to pursue them is never clearly assigned. The future management architecture places decision authority at the functional centre of the organisation — not at the bureaucratic apex — and makes explicit the question that current structures leave unanswered: who is structurally positioned to accept residual risk and be accountable for the consequences? Risk managers who can put that question on the table, with evidence and a structural proposal, are delivering genuine strategic value.

03 - The silent architecture beneath every risk decision determines whether the visible architecture works.

The risk profession has built sophisticated visible structures: frameworks, registers, committees, models, and software. What it has not built is equal sophistication in the human and philosophical substrate beneath them — the risk philosophy of individual leaders, the incentive effects shaping what the risk function actually protects, the ethical boundaries that should be non-negotiable, the psychological burden on those carrying enterprise-level uncertainty, and the temporal dynamics that make yesterday's decision obsolete today. These are not soft issues. They are

the mechanisms through which risk is interpreted, disputed, accepted, and owned. An architecture that does not account for them is incomplete — not in its documentation, but in its actual decision quality.