Aera TECHNOLOGY



The Decision-Centric Enterprise

Why the next era of enterprise AI will be measured in decisions, not models, and what that asks of the C-suite

Table of Contents

Executive Summary 3

What Has Changed 4

From Well Informed to Well Decided 5

What Decision-Centric Means 6

The Tech Stack Collapses into Two Tiers 8

Connecting Decisions Across the Value Chain 11

The New Executive Operating Model 13

The Economics of Decisions at Scale 15

Where to Start 17

Conclusion 18

Executive Summary

For two decades, enterprise technology investment has rested on a single premise: better information produces better outcomes. Companies built the warehouses, lakes, pipelines, and dashboards that premise implied, along with the organizations to run them. The premise was half right. Information became abundant. Outcomes did not improve at anything like the same rate, because the enterprise never fixed the step that actually creates value: the decision itself.

Over the past eighteen months, that gap has stopped being an abstraction and become a budget problem. Boards that once asked whether their company had an AI strategy now ask which decisions are being run with AI and what those decisions returned. The question is harder than it sounds, because most enterprises have never measured a decision — not its cycle time, not its quality, and certainly not its cost.

This paper argues that the organizing unit of enterprise performance is shifting from data to decisions, and that this shift changes four things at once: how the enterprise technology stack is architected, how functions are connected, how executives lead, and how the economics of AI are governed. It lays out what a decision-centric operating model looks like in practice, what the leading enterprises are doing that the rest are not, and how to begin without a multi-year replatforming program.

An enterprise is not a collection of data. It's essentially a machine that makes several million decisions a year. Almost none of them are measured.

What Has Changed

Four developments, none of them individually decisive, have combined to move enterprise AI out of the innovation budget and into the operating plan.

Capability Became Table Stakes; Accountability Did Not

The competitive question is no longer what a model can do. Frontier capability is available to every company at a price that keeps falling, which means it confers no advantage on anyone. What remains scarce is the context that makes a model useful inside a specific business — the constraints, the trade-offs, the commercial commitments, the institutional judgment — and the governance that lets a system act on it. Differentiation has moved up the stack, away from the model and toward the decision.

Agents Acquired Hands

The generation of agents deployed in 2024 produced recommendations. The generation deployed now produces transactions: it places the replenishment order, reallocates the inventory, adjusts the price, releases the payment. That is a categorical change, not an incremental one. It converts the governance conversation from a question of accuracy — is the answer right? — into a question of authority: what is this system permitted to do, up to what value, under whose mandate, and how is it reversed? Most enterprises have a well-developed answer to the first question and no answer at all to the second.

Scale Exposed the Real Bottleneck

Enterprises spent the better part of a decade modernizing the data estate, and the modernization largely worked. What it revealed is that insight was never the constraint. The constraint has been the distance between signal and action: the hours or days in which a human reads a dashboard, convenes a call, builds a spreadsheet, and sends an email. The data was fast. The decision was not.

Ownership Moved Out of IT

The clearest signal that a technology has crossed over is that the people who fund it change. Enterprise AI is now discussed by chief operating officers and chief financial officers with a profit-and-loss statement in front of them, not by technology leaders with a roadmap. That has consequences the rest of this paper is concerned with, because operating executives ask different questions: not what is possible, but what it costs, what it returns, and who is accountable when it's wrong.

From Well Informed to Well Decided

It's worth being precise about the limits of the data-driven model, because the criticism is often wholesale, while the investment was not wasted.

Data-driven programs optimized the supply side of information. They made data available, consistent, and queryable, and that was a necessary achievement. But availability is not the same as application. A dashboard is an offer, not an outcome; it presumes a human will interpret it, weigh it against competing objectives, and act. That human is typically one of several, in different functions, with different incentives, working on different cycles.

The result is an enterprise that's extremely well informed — but still slow. Three patterns hold it back:

- **Latency.** Information is produced at machine speed and consumed at meeting speed. The signal arrives on Tuesday and is acted on at the Thursday planning call, by which point the conditions that produced it have changed.
- **Local optimization.** Every function optimizes its own metric correctly. Supply chain protects service levels, finance protects working capital, commercial protects revenue. Each is doing its job, but the enterprise still loses, because nobody owns the trade-off between them.
- **Evaporation.** The reasoning behind a decision lives in the head of the person who made it. When that person leaves, thirty years of judgment leaves with them, so the enterprise has to relearn what it already knew.

None of these is a data problem. All three are *decision* problems, and no additional investment in the data estate resolves any of them.

What Decision-Centric Means

A decision-centric operating model treats the decision — not the report, the dataset, or the model — as the object the enterprise designs, measures, and improves. In practice that means a decision is modeled explicitly, with five things made clear:

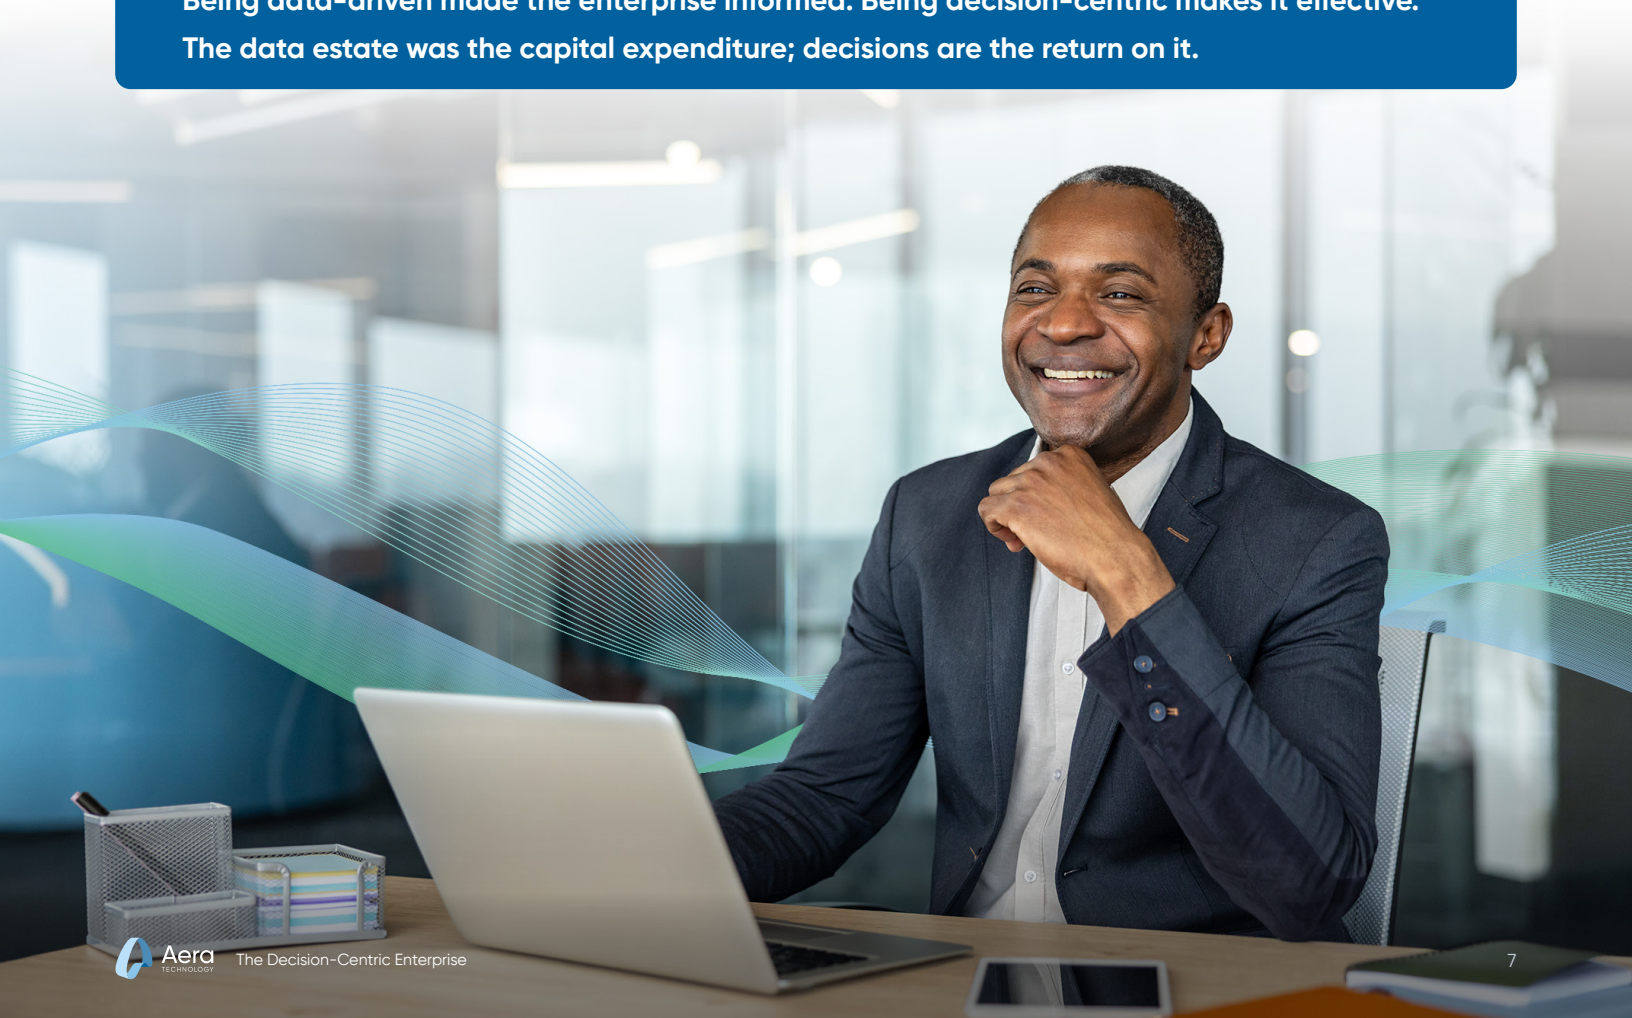
1. **Objective**: what the decision is optimizing for, stated in business terms rather than implied by a metric on a chart.
2. **Constraints**: what it must not violate — service commitments, credit limits, regulatory requirements, capacity.
3. **Trade-offs**: what is being given up in exchange for what, quantified. A point of service level has a working-capital price, and in most enterprises that price has never been written down.
4. **Authority**: who or what is permitted to make this call, at what value, and how it escalates or reverses.
5. **Outcome**: what actually happened, captured and fed back into the system, so the next instance of this decision is better than the last.

The last point is the one that compounds. When decisions are modeled and their outcomes captured, an organization accumulates decision memory: a durable, queryable record of what it chose, why, and what it got. Institutional judgment stops being a property of individuals and becomes a property of the enterprise. Gartner has projected that "By 2030, explicitly modeled business decisions will be five times more trusted and 80% faster than ungoverned decisions, enabled by decision intelligence platform adoption."[1] That's a useful framing, because it locates the advantage in the *modeling*, not the automation.

[1]Gartner®, Magic Quadrant™ for Decision Intelligence Platforms, David Pidsley, Carlie Idoine, Gareth Herschel, Kevin Quinn, Kjell Carlsson, 26 January 2026.

	Data-driven	Decision-centric
Unit of value	The dataset and the dashboard	The decision and its outcome
Primary artifact	A report someone must act on	An action already taken, with its rationale attached
Cadence	Weekly and monthly planning cycles	Continuous, triggered by conditions
Measure of success	Adoption, data quality, query volume	Cycle time, automation rate, value per decision
Where judgment lives	In the heads of experienced people	In modeled policy and decision memory
Characteristic failure	Well informed and slow	Poorly governed autonomy — which is why governance is designed in

Being data-driven made the enterprise informed. Being decision-centric makes it effective. The data estate was the capital expenditure; decisions are the return on it.



The Tech Stack Collapses into Two Tiers

Becoming decision-centric changes more than how an enterprise operates. It changes how its technology is structured. Today's stack was designed around a simple division of labor: software informs, and people decide. As systems begin to decide and act, the stack takes a new shape.

How the Stack Was Built

Most enterprises run on three broad layers. At the base are the systems of record: the ERP and transactional core that books orders, moves stock, and closes the ledger. Above them is a wide middle layer of purpose-built tools. Each one is designed to help a person solve one narrow problem: a planning application, a forecasting tool, a control tower, a dashboard. Integration and workflow software ties them together. On top sits a growing layer of analytics, copilots, and agents.

Each of these layers was a sound answer to a real need. The middle layer has one essential job. It carries information to a person who decides, then carries that person's decision back into a transactional system to be executed. Every hand-off adds a seam, and every seam adds time. This is the distance between signal and action described in the first section, expressed in software.

Where the Middle Layer's Work Goes

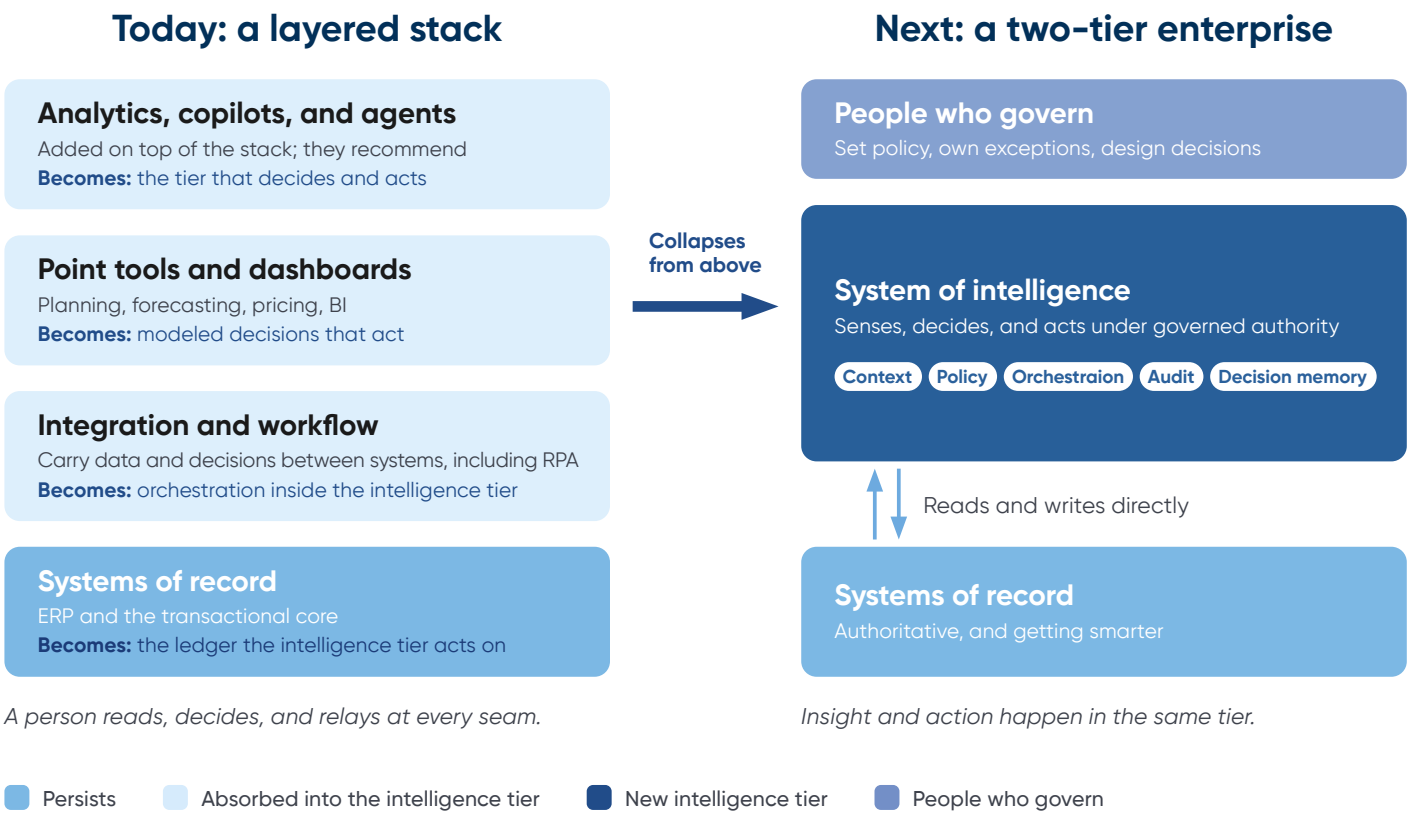
Once a system can both sense and act, reading the data and writing the transaction happen in the same place. The work the middle layer was built for, helping people decide and relaying what they decided, moves into the intelligence itself. Adding AI to each tool improves that tool, but the hand-offs remain.

Dumb software with a layer of AI on top is still dumb software. The opportunity is a simpler architecture, not a smarter version of the old one.

The Two-Tier Destination

The stack that emerges has but two tiers:

- **The system of record** stays authoritative and control-oriented, and it gets smarter over time. It continues to book, post, and reconcile. It becomes the ledger the intelligence tier reads from and writes into.
- **The system of intelligence** senses conditions across the value chain, models decisions explicitly, and acts under governed authority. It can hold a purchase order, rebalance inventory, or adjust a production line. Context, trade-off policy, orchestration, execution authority, audit, decision memory, and unit cost all live here. Its foundation is a shared decision data model (described in the following section). That model takes over the work of both the reporting semantic layer and much of the point-to-point integration beneath the old middle tier.



Collapsing from Above

This shift doesn't require a replatforming program. The intelligence tier sits on top of the systems the enterprise already runs and works directly against them. As each decision moves onto it, the tools that supported that decision are used progressively less, then ultimately not at all. They retire on the enterprise's own schedule. The stack simplifies from above, one decision at a time, while the transactional core keeps running. When a case falls outside policy, it escalates to a person with full context attached, the same path exceptions follow today.

Why the Intelligence Tier Sits Across Systems

Application vendors are adding intelligence to their own products, and those products will improve as a result. But the decisions that matter most *cross systems*. A pricing decision touches margin, inventory, and cash, and the data for each lives somewhere different. Intelligence that lives inside one application sees the enterprise through that application's window. The tier that collapses the stack has to sit across the systems of record, with one view of context, policy, and audit. That same position is what keeps a growing population of agents coherent, as the next section describes.



Connecting Decisions Across the Value Chain

The enterprises making real progress share an architectural conviction: decisions form a network, not a stack of departmental applications. A pricing decision is a margin decision is an inventory decision is a cash decision. Treating those as four systems owned by four functions is what produces the quarterly meeting in which the best arguer wins.

Four practices distinguish them.

A Shared Decision Data Model

The decision data model is not another data lake; it's a shared meaning. ERP, planning, logistics, and commercial systems each carry their own definitions of an order, a cost, a customer, a week. Until those are harmonized into a model built for decisions rather than for reporting, cross-functional decisions degenerate into arguments about whose number is correct. Leading organizations do this work once, deliberately, and treat it as infrastructure.

Trade-Offs Encoded, Not Debated

The exchange rate between competing objectives is the single most valuable undocumented asset in most enterprises. What is a point of fill rate worth in working capital? What is an expedited freight dollar worth in retained revenue? Leaders write these down, attach owners to them, and let systems apply them consistently. That has the secondary benefit of surfacing the cases where the company has never actually agreed.

Authority Assigned by Value and Reversibility

The useful axis is not confidence but consequence. Decisions that are small and reversible should execute autonomously; decisions that are large and permanent should escalate with a full recommendation and an audit trail. Most enterprises get this backwards, applying uniform human review to everything or, worse, nothing. Segmenting a decision portfolio this way is usually the fastest route to measurable value, because it concentrates automation where the risk is lowest and the volume is highest.

Decisions Measured Like a Production Line

What gets instrumented gets improved. The leaders track cycle time, automation rate, exception rate, and realized value by decision type, and they review those numbers the way a plant manager reviews throughput. This is unglamorous work, but it separates programs that scale from those that stall.

The Failure Mode to Avoid

The characteristic 2026 failure is agent sprawl: dozens of agents from a dozen vendors, each locally competent, none sharing context, collectively incoherent. It's shadow IT with an execution budget and a spend curve. But the remedy is not fewer agents; it's a common decision layer underneath them — shared context, shared policies, and a single audit trail — so that those agents work together instead of colliding. That common layer is the intelligence tier described in the previous section.

The pattern shows up clearly under stress. When a tariff or input-cost shock lands, a connected enterprise reprices, re-sources, and re-plans inventory as a single decision with one set of trade-offs. A disconnected one runs three sequential functional projects and finishes a quarter late, by which time the shock has moved. Aera customers describe exactly this contrast: The Estée Lauder Companies has spoken publicly about moving beyond dashboards to systems that autonomously optimize competing objectives in the supply chain, and Kerry Group has described automating thousands of inventory decisions with savings in the millions while reducing waste and risk.

The New Executive Operating Model

When decisions become a designed system, the job of leading changes shape. The most useful analogy is the trading floor: a firm does not scale by hiring a better trader, it scales by building a desk with a mandate, position limits, and risk controls. Executives are making the same transition, from making the call to designing the system that makes ten thousand of them.

Role	From	To
CEO	Allocating capital annually; stating strategy in narrative terms	Allocating decision authority continuously; specifying strategy precisely enough that a system can optimize for it
CFO	Reporting outcomes and approving project budgets	Owning the trade-off policy and instrumenting cost and return per decision rather than per project
CIO / CTO	Delivering systems of record and applications	Running a two-tier architecture: a system of record that stays authoritative and a system of intelligence with context, orchestration, and a control plane for agents (identity, permissions, audit, unit cost); retiring point tools as decisions migrate
Functional Leaders	Making the calls, two hundred a week	Setting policy, owning exceptions, and supervising twenty thousand — judgment encoded rather than judgment in the moment

Two consequences deserve to be said plainly. The first is that an autonomous system will faithfully maximize whatever it is told to maximize, which forces a precision about strategy that most organizations have never had to supply. Vague strategy survives in a slide deck; it does not survive contact with an optimizer.

The second is organizational. The layers of management whose principal function is relaying information upward will thin out, because the information now moves without them. The parts of the role that grow are the parts that were always the hardest: deciding what the enterprise should optimize for, handling the exceptions that policy cannot anticipate, and developing the people who will do both. Companies that treat this purely as a headcount exercise will capture the cost reduction but lose the capability.

The technology organization changes shape in the same way. Teams built around individual applications, reports, and integrations consolidate around the intelligence tier. Their work moves from keeping tools running to designing and governing decisions.

You can automate a decision. You cannot automate the accountability for it. That's what keeps this an executive agenda rather than a technology one.



The Economics of Decisions at Scale

The last shift is the one most likely to determine which programs survive the next budget cycle, and it is the least discussed.

Pilots hide economics. At a hundred decisions a month, compute cost is a rounding error and the only question that matters is whether the output is any good. At ten million decisions a year, cost becomes a line item, and the governing question changes: what does this decision cost to make, and what is the margin on it?

This is the first category of enterprise software in which cost scales with usage rather than with seats. Success is now a cost event: the better the system works, the more it's used, and the more it consumes. That breaks the budgeting model most technology organizations have relied on for thirty years, and it is why AI spend has migrated from the technology budget into capital planning, where it is compared against other uses of the same money.

Two Numbers

A decision-centric enterprise governs AI with two figures, tracked by decision type: value per decision and cost per decision. If value exceeds cost and the volume is meaningful, the decision is scaling without further debate. If the organization cannot compute both, then it's not ready to scale — and this, more than model quality, explains why so much AI investment has produced no attributable return.

The arithmetic is unforgiving in both directions. Ten million decisions a year at ten cents each is a million dollars; the same volume at two cents is two hundred thousand. That gap is a function of architecture, not of vendor discount. And if each of those decisions saves a dollar, the board conversation is about how fast to scale, not where to cut.

Consolidation Funds the Shift

As decisions move onto the intelligence tier, the licenses, integration upkeep, and report maintenance tied to the tools they replace can be retired. That adds a second line to the business case: alongside value per decision, the spend released from the old middle layer. Over time, the intelligence tier is increasingly funded by the tools it makes unnecessary.

Route Compute by Decision Value

Not every decision deserves a frontier model. A fifty-dollar replenishment call should run on a cheap deterministic path; a five-million-dollar sourcing shift earns expensive reasoning and a human in the loop. Most cost overruns come from applying uniform compute to wildly non-uniform decisions — a portfolio problem that is solved in design, not in procurement.

Governance Is an Economic Feature

The largest hidden cost is not tokens. It's rework, reconciliation, and human review of output that nobody trusts. A system requiring a person to verify every answer has no economics at all, however cheap its inference. Explainability, audit trails, and bounded authority are therefore not a compliance tax levied against the business case — they're the mechanism that makes the business case real, because they're what allow a human to stop checking.

Where to Start

The approach that works is narrow, instrumented, and owned by the business rather than by a center of excellence. A practical sequence:



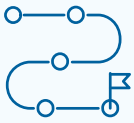
1. Pick one decision that has a P&L owner.
Not a model, and not a use case. A decision with real volume, a clear owner, and a measurable outcome — inventory deployment, payment-term exceptions, promotional allocation. Volume is what makes the economics legible.



2. Baseline it before changing anything.
Cycle time, cost, and outcome quality. Most organizations discover here that they cannot answer basic questions about a decision they make ten thousand times a year, which is itself the finding.



3. Model the decision explicitly.
Objective, constraints, trade-offs, authority. Insist on the trade-off exchange rate in writing; the argument it provokes is the valuable part.



4. Automate by consequence, not by confidence.
Start with the small and reversible cases, and keep the escalation path and audit trail live from day one. Trust is built by the reversals that worked, not by the accuracy statistics.



5. Connect the second decision before scaling the first.
A second and third decision in an adjacent function, connected through the same decision data model, is where cross-functional value appears — and where the difference between a platform and a point solution becomes obvious.



6. Retire as you migrate.
When a decision runs on the intelligence tier, retire the tools that supported it rather than running both. Weigh every new request for a point solution against one question: will this decision live on the intelligence tier? Each retirement simplifies the stack and funds the next decision.

Two quarters of this produces something more useful than a strategy document: a decision that runs, a number attached to it, and an organization that has learned what governed autonomy actually feels like.

Conclusion

The enterprises that will compound an advantage from AI are not the ones with privileged access to models. They're the ones that have made their decisions explicit: modeled, connected across functions, bounded by governed authority, measured in unit economics, and run on a simpler, two-tier stack. That work is unglamorous next to the model announcements, and it is the part that does not commoditize.

The shift from data-driven to decision-centric is not a repudiation of the last twenty years of investment. It is the claim on its return. The question a board should be asking is no longer whether the company has an AI strategy. Rather, it's a simpler and much harder one: which decisions are we running with AI, what did they cost, and what did they earn?

About Aera Technology

Aera Technology is a leader in agentic decision intelligence. Aera connects data across enterprise systems, models decisions explicitly, recommends and executes them under governed authority, and captures outcomes as decision memory so the enterprise improves with every decision it makes. Aera was named a Leader in the 2026 Gartner® Magic Quadrant™ for Decision Intelligence Platforms.

Aera TECHNOLOGY

aeratechnology.com

linkedin.com/company/aera-technology

Are you ready to lead the future of decision-making?

The way your organization makes and executes decisions is a powerful lever for competitive differentiation. Aera equips your business with the intelligence and automation needed to surface untapped value, eliminate waste, and accelerate sustainable, long-term growth.

Take the first step today:

See how Aera can revolutionize your operations with real-time, AI-driven decisions.

Request a Demo

Discover how decision intelligence can deliver measurable results tailored to your business.

Speak with an Expert