



PLAYBOOK

The Five Patterns That Make Enterprise AI Actually Work

a **dentsu** company

INTRO

What Makes Enterprise AI Work

Enterprise AI spending is [projected to reach \\$3.3 trillion by 2027](#). But investment alone **isn't producing results**.

[According to Gartner](#), only 28% of AI use cases in infrastructure and operations fully succeed and meet ROI expectations. Leaders are aligned. Budgets are committed. Pilots are succeeding. And yet over 40% of agentic AI projects are [predicted to be canceled](#) by the end of 2027. When it's time to move from proof of concept to scaled deployment, efforts stall.

What's getting in the way?

Part of it is a definition problem. When most executives hear “enterprise AI,” they picture one of a few things: complex, custom AI builds; a company-wide ChatGPT rollout; or a massive contract with a provider like Salesforce, Adobe, Microsoft, or IBM.

Any of these might be part of a smart approach, but mistaking tactics for strategy is where expensive, underwhelming transformations begin.

Enterprise AI is something more specific. It's the AI that organizations deploy internally to change how their teams work, how decisions get made, and how value gets created and delivered. The agent that auto-classifies incoming service requests before a human ever sees them. The system that synthesizes performance data into a morning briefing for a VP making resource allocation decisions.

The workflow that routes customer data through governed infrastructure so a commerce team can act on it in real time. It rarely makes for a flashy demo, but it's where operational advantage lives.

At Merkle, we've had a front-row seat to AI successes. Through hundreds of machine learning projects over the past three years, we've identified five patterns that separate organizations that escape pilot purgatory from those stuck in it.

This playbook explores those patterns and shows how to apply them, building an enterprise that's more efficient, delivers stronger outcomes, and creates better customer experiences.

What is enterprise AI?

The **AI that organizations deploy internally** to change how their teams work, how decisions get made, and how value gets created and delivered.

”

Contents

05	Pattern 01: Follow the data gravity, not the hype
11	Pattern 02: Start with the decision, not the technology
15	Pattern 03: Redesign the work, not just the tools
21	Pattern 04: Earn trust through measurement
26	Pattern 05: Embed intelligence at every altitude
29	Conclusion: Applying the Patterns Together

PATTERN 01

Follow the data gravity, not the hype




THE FIVE PATTERNS THAT MAKE ENTERPRISE AI ACTUALLY WORK

Follow Your Data Gravity

Every enterprise has one or more systems where data naturally concentrates, whether that's customer data in a CDP or transaction data in a data warehouse. These concentrations are your *data gravity*: the places where your data is owned, deep, and closest to the decisions that matter. When organizations ask where to start with AI, the answer is almost always closer than they think.

Your data already clusters around specific systems and processes, and those clusters tell you where AI will generate value fastest.

Every time data moves between systems, complexity follows. AI that lives close to the data reduces that complexity and the organizational burden that comes with data movement.



The advantages:

Real-time decisioning without latency

Agents access data directly, eliminating the sync delays that slow down activation.

Context that travels with the data

Meaning doesn't get lost in transit. AI interprets data with the full semantic context intact.

Always-fresh data for smarter decisions

AI works from live data rather than stale snapshots, so recommendations reflect what's happening now.

Governance that holds across every interaction

Lineage and access controls stay consistent. There's no gap in accountability as data moves between systems.

The future state is coordination between the AI agents using your most concentrated data and any department-specific agents (think finance, marketing, or operations). But you start where data gravity is strongest. That's true whether you're building your own AI capabilities or working with in-platform solutions.

In practice:

An insurance client built an AI close to their customer service data: call transcripts, tickets, and chat logs. That enabled their AI agents to instantly prioritize and route service inquiries, improving service delivery and saving \$1.5M per month in costs.



The Data Readiness Problem

Data gravity points you in the right direction, but it doesn't guarantee the data is ready when you get there. Three signs your data isn't ready:

- Two analysts pull the same metric and get different numbers. Neither can fully explain why.
- Your data schema is full of fields like “order_status_cd” or “loc_type.” The context for what's allowed in those fields and what the values actually mean isn't written down.
- The logic that turns raw tables into reliable answers isn't documented anywhere. It lives with a handful of people: SQL snippets passed around in chat, BI tools that don't talk to each other, or that one analyst everyone routes questions through.

AI won't fix these ambiguities in your data. **It will inherit them.** Most enterprise data was designed for human interpretation. Experienced employees carry the context that makes it usable. They know which system to trust when two disagree and which definitions shift depending on the business unit.


AI doesn't carry that context. It reads exactly what it's given, and when the underlying meaning isn't encoded, it produces confident answers that are wrong.

Closing the Enterprise AI Data Gap

You need two shifts to make data reliable for AI usage: one structural and one cognitive.


Cognitive: Describe the data.

Build a semantic layer: a set of configuration files (typically formatted as YAML or JSON) that live inside a tool like Looker or Cube, or as a dedicated metadata layer within your data platform. The semantic layer stores the definitions, rules, and relationships that tell every AI agent what the data actually means.



Structural: Curate the data.

Build focused datasets scoped to the AI use cases you've defined. Consolidate the relevant data into a single, navigable layer. That means clean structure, clear connections between sources, and no hidden logic. This eliminates the core issues that make enterprise data hard for AI to use reliably.



For each concept that matters to your use cases, that means documenting:

Definitions

What does “completed order” actually mean, and which system is the source of truth?

Relationships

How does this concept connect to customers, locations, SLAs or financial periods?

Business rules

What conditions change the meaning, and which versions of the data are still active?

Usage

Which teams own it, which decisions depend on it, and what breaks downstream if it changes?

To make that concrete, here's an illustrative example of what a semantic layer entry for "completed order" might look like in practice:

```
metric: completed_order  
  
description: A customer order where payment has been confirmed and  
            fulfillment has been initiated. Does not include orders in  
            "held for review" status.  
  
source_of_truth: order_management_system  
  
source_table: oms.orders  
  
filter: status_cd = 'C' AND fulfillment_initiated = true  
  
related_entities: [customer, location, financial_period]  
  
owner: Revenue Operations  
  
downstream_reports: [daily_sales_summary, regional_performance, CFO_dashboard]  
  
deprecated_codes: ['H', 'P', 'X']
```

This is what translates raw data into AI reasoning. Without it, an AI agent is technically able to query everything, but capable of understanding almost none of it. These two shifts are only effective when they happen together.

Structure without definition is faster access to the same ambiguous answers. Definition without structure is documentation sitting on top of a mess. **But together, structure and definition turn data gravity into a foundation AI can build on.**



Make Existing Platforms Work Together

Most organizations have several data gravities: a marketing cloud here, a CRM there, and a data platform underneath it all. Each one has its own concentration of data, its own workflows, and increasingly, its own embedded AI capabilities.

The first move is straightforward: **use what you've already paid for**. Most organizations dramatically underutilize the AI capabilities built into their existing platforms. Marketing clouds can handle audience segmentation and campaign planning. CRMs can summarize records, draft communications, and surface account insights. Data platforms support natural language querying and automated analysis.

These are available now, and activating them doesn't require a new vendor relationship or build project. But native AI only gets you so far. **The harder question is how agents operating in different platforms coordinate with each other.**

An agent working in your CRM and an agent working in your data platform are each close to their own data. Getting them to reason together requires a shared foundation: common definitions, consistent business rules, and clear orchestration patterns that govern how they hand off context and decisions.

This is where the semantic layer from the previous section does its work. It gives every agent in your environment a common language to work from, regardless of which platform they live in. Emerging standards like [Model Context Protocol \(MCP\)](#) and [Agent2Agent Protocol \(A2A\)](#) are making this coordination more tractable, establishing standards for how agents communicate across systems. Whether your organization has one data gravity or twenty, the principle holds: **build an environment where AI can follow your data wherever it lives.**



PATTERN 02

Start with the decision, not the technology



Specific Use Cases Beat Enterprise-Wide Programs

Most AI initiatives start with ambitious goals: “Transform customer service.” “Revolutionize marketing.” “Automate operations.”

Then they stall, because no one can evaluate whether they’re working. Vague use cases can’t be executed, measured, or optimized. **That’s why successful organizations get granular with their AI use cases especially early in their AI journey.** A well-defined use case must answer three questions:

1. • **What decision does this enable?**

2. • **What delay does it eliminate?**

3. • **How will you know if it’s working in six months?**



This means saying no to the dozens of other AI ideas floating around, and defining success in terms a CFO would recognize, not “AI maturity.”

What Good Use Case Definition Looks Like

Let's rewrite one of the examples above—"Transform customer service"—to show the difference:

What's the inspiring vision for this change?

Delight customers with exceptional service.

What granular action will we take?

Auto-classify incoming emails by urgency and topic, routing to specialist customer service representatives.

What decision will this enable?

Which service inquiries need to be addressed fastest, and the best way to address them.

What delay does this eliminate?

Reduces latency from service inquiry to resolution; decreases time required by customer service representatives to resolve issues.

How will we know if it's working in six months?

- KPIs: 40% reduction in resolution time, 15-point net promoter score increase, 25% boost in customer service representatives' satisfaction.
- Target outcomes: Customers feel heard and valued by the organization; customer service representatives handle more meaningful work.

This level of detail delivers three key benefits:

It creates provable ROI.

Each use case has quantifiable value that justifies continued investment. **Stakeholders can see exactly what they're funding and what they'll get in return.**

It enables faster learning cycles.

Success isn't determined by gut feelings.

Measurable goals eliminate guesswork and indicate when to scale success or pivot for improvement.

It guides strategic decisions.

When you understand the specifics, the line between strategic capability and commodity becomes clear. **That sharpens every downstream choice about resourcing, sequencing, and where to invest most rigorously.**

In practice:

Satair receives more than 6 million customer support emails annually, and they needed a way to address a specific, measurable problem: how do we respond to quote requests faster to meet growing demand? Merkle built Lilly, an AI assistant trained on Satair's quoting data, that now handles more than 450,000 quote emails per year.

You can read the full case study [here](#).



PATTERN 03

Redesign the work, not just the tools.



THE FIVE PATTERNS THAT MAKE ENTERPRISE AI ACTUALLY WORK

The Operating Model Matters More than the AI Model

Many companies have bold AI visions and powerful technology assets, but lack modern operating models to connect them. Without that connective tissue, even the best AI investments underdeliver.

AI applied to poor workflows doesn't fix them, it amplifies them. Yet many organizations still focus heavily on technology and platforms while underfunding work redesign efforts that determine whether any of it sticks.

AI demands new operating models because the way businesses create and deliver value is changing. **Customer experiences that were once impossible are now achievable, but only when the organization is built for humans and AI to work together, not just side by side.**

AI applied to poor workflows doesn't fix them, **it amplifies them.**

”



An effective operating model does four critical things:

It clarifies accountability.



Defining who makes decisions, who has final say, and where AI recommends versus decides.

It shows how AI complements roles.



Demonstrating how AI handles routine work while humans focus on strategy, creativity, and complex problem-solving.

It enables AI to support the way teams actually work.



Designing workflows that integrate AI into natural team dynamics rather than forcing teams to adapt to rigid systems.

It establishes governance and visibility.



Defining how humans and AI check each other's work, when to hand off, and how oversight remains clear as systems scale.

Investing in the People Side of Redesign

Work redesign requires real investment in the people living through the change. [Companies currently spend 93% of their AI investments on technology and just 7% on people](#), treating the human element as a cost to minimize rather than a capability to build. The organizations that escape pilot purgatory adjust that equation to achieve greater balance. In practice, that means:

Role redesign:

Systematically restructure job duties, responsibilities, and reporting lines so employees can clearly see how their role is evolving, not disappearing. When a marketing analyst understands they're becoming a strategic advisor who guides AI-generated insights rather than a data processor, they lean in rather than resist.

Cross-training:

Help employees grow into redesigned roles, whether that's a creative designer expanding into strategy and prompt engineering, or an analyst developing skills in AI oversight and quality evaluation.

Scaled AI training and literacy:


Go beyond "how to use the tool" to help teams understand what AI can do, where it needs oversight, and how to integrate it into daily work to actually improve outcomes.

Underpinning all of this work is codified human-AI collaboration, explicitly defining where AI acts, where humans decide, and how the two hand off. For most organizations, this starts from scratch. Operating models are rarely written down; they live in habit and institutional knowledge. AI doesn't tolerate that ambiguity, which is exactly why documenting ways of working is worth the effort.

Without that clarity documented and shared, even well-designed roles and well-trained people will default to inconsistency. With it, the operating model becomes something people can actually use.

Why **Clarity** Compounds

It's tempting to see work design and codification as bureaucratic overhead that slows teams down. With AI, the opposite is true. **Without clarity about how work gets done, AI systems produce inconsistent outputs and humans spend their time debating edge cases instead of doing meaningful work.** With clarity, AI operates reliably within defined guardrails, exceptions get resolved through documented rules, and people are freed to focus on the work that actually requires human judgment.



What Our Research Shows

The case for prioritizing work redesign is well supported. [Merkle's research](#) across 100 US C-suite leaders found that talent gaps and insufficient training are among the leading causes of underutilized AI technology, and that tech-mature organizations are 3.5x more likely to operate with a formalized change management function backed by dedicated resources.

The organizations pulling ahead share a few consistent patterns:

- **Cross-functional alignment** on a unified vision across marketing, technology, and data;
- **Formalized change management** treated as a core workstream rather than a side project;
- **Success metrics tied to business value** such as revenue impact and forecasting accuracy, rather than the vanity metrics that make AI initiatives look busy without proving they're working.

3.5x

Tech-mature organizations are 3.5x more likely to **operate with a formalized change management function backed by dedicated resources.**

”

In practice:

Merkle established unified measurement frameworks and integrated CX planning as part of Lumen's content supply chain transformation, leading to a 64% reduction in content creation time and social campaigns getting to market 3x faster.

Read the full case study [here](#).



PATTERN 04

Earn trust through measurement



THE FIVE PATTERNS THAT MAKE ENTERPRISE AI ACTUALLY WORK



Trust Must Be **Engineered**, Not Assumed

AI systems work in probabilities. The same question can produce different answers, and some will be wrong. That's not a reason to avoid AI, but a reason to build measurement into the foundation from day one. With the right evaluation framework and iterative improvements, reliability can be directly addressed.

The contrast between operating with and without measurement is stark:

Without Measurement

- Teams check everything manually, eliminating efficiency gains
- Executives can't quantify ROI or justify investment
- The organization drifts back to the old way

With Measurement

- Accuracy is baselined, failures are surfaced, improvement is continuous
- Performance is tracked against defined quality standards
- Trust compounds over time

What Measurement Actually Looks Like

Strong Evaluation Frameworks

Defined metrics tell you whether AI is working as intended, where it’s falling short, and where to focus improvement efforts.

That means tracking two things: operational performance for the teams building and managing the system, and business impact for the leaders funding it.



Operational Performance:

	What it Measures	Why it Matters
Accuracy Rate	How often the AI produces a correct or acceptable output.	The baseline signal for whether the system is working as intended.
Output Consistency	Whether the AI produces the same quality of answer across different users, times, and contexts.	Inconsistency erodes trust faster than occasional errors.
Adoption Rate	What percentage of intended users are actually using the system.	Low adoption usually signals a workflow, trust, or training problem.

Business Impact:

	What it Measures	Why it Matters
ROI	Return generated relative to the cost of building and running the system.	The primary metric for justifying continued investment.
Revenue Contribution	Incremental revenue attributable to AI-enabled decisions or actions.	Connects AI performance directly to business outcomes.

Transparency and Observability

Mechanisms make it possible to see what the AI is doing and why, so that failures can be surfaced and addressed rather than quietly eroding confidence.

Together, these capabilities create a feedback loop that turns real-world deployment into a source of continuous learning. That’s why (counterintuitively) systems built for trust don’t need to be perfect to launch. Waiting for perfect conditions means missing the opportunity to learn from actual use.



Measurement As an Always-On Requirement

Deployment is a milestone, not a finish line. Models drift, business needs evolve, and customer expectations rise. The organizations that sustain AI value treat measurement and improvement as permanent operating costs, not post-launch afterthoughts. **That means monitoring from day one, building feedback loops for users to inform system refinement, and budgeting for continuous improvement.**

The payoff is an organization that trusts its AI enough to expand it, moving from cautious pilots to scaled, embedded capabilities with the confidence that comes from knowing the system is being watched, measured, and consistently improved.

In practice:

When a major restaurant brand's AI chatbot began returning inaccurate results across its 6,000-plus locations (each with different store sizes, designs, and customer profiles), field teams lost confidence in the data. By investing in accuracy and making that accuracy visible, the brand rebuilt trust and scaled adoption to over 100,000 users who now rely on it for location health metrics.



Ebook

Read [*How Real-Time Analytics Keeps You a Step Ahead of Customer Expectations*](#) to learn how to build a real-time data foundation that makes measurement meaningful.

LEARN MORE



PATTERN 05

Embed intelligence at every altitude



THE FIVE PATTERNS THAT MAKE ENTERPRISE AI ACTUALLY WORK

Better Decisions at Every Level

The organizations creating the most value from AI aren't building one dashboard or one chatbot. They're building intelligence systems: connected layers of data, AI, and automated delivery that put the right insights in front of the people responsible for acting on them. That requires a mix of reactive and proactive intelligence.

Pull (Reactive): The user initiates

Mechanisms:

- Dashboards
- Conversational interfaces
- Ad-hoc queries

Strengths:

- Flexible
- Exploratory
- Good for novel questions and deep investigation

Limitations:

Requires someone to know what to ask and when to ask it

Best for:

- Exploratory analysis
- Root cause investigation
- Complex or ambiguous questions

Push (Proactive): The system initiates

Mechanisms:

- Scheduled briefings
- Triggered alerts
- Automated digests

Strengths:

- Consistent
- Timely
- Doesn't depend on someone remembering to check

Limitations:

Only covers what you've anticipated and configured

Best for:

- Routine monitoring
- Threshold-based alerts
- Recurring reporting

Between these two modes sits the human judgment layer: analyst review on critical decisions before action, context and interpretation AI can't provide, investigation of anomalies flagged by proactive systems, and sign-off on high-stakes recommendations. AI handles the intelligence. Your people handle the judgment.

The Right Intelligence at the Right Altitude

Different levels of the organization need different intelligence from AI. Access to everything isn't the same as access to what matters. Your C-suite doesn't need daily platform-level performance activations, and your analysts don't need weekly market positioning reports. Effective intelligence delivers what each altitude needs in the format and frequency the decision requires.

	Strategic	Operational	Execution
Who	• C-suite • VPs	• Directors • Managers	• Analysts • Specialists
Decisions	• Portfolio direction • Investment allocation • Market positioning	• Campaign optimization • Resource allocation • Vendor performance	• Daily optimizations • Tactical adjustments • Issue diagnosis
Intelligence Needed	• Synthesized trends • Cross-functional performance • Competitive signals	• Performance against targets • Anomaly detection • Forecasts	• Granular metrics • Root cause analysis • Real-time signals
Key Decision Cadence	Monthly, quarterly	Weekly, triggered	Daily, triggered

CONCLUSION

Bringing the Patterns Together to Take Action



Your Opportunity: The Five Patterns, Working Together

The five patterns aren't independent. Each one makes the others more effective to build scaled, compounding value.

When you follow your data gravity:

Your use cases have reliable, real-time data to execute against. Teams work in familiar systems, which accelerates adoption. Governed data makes performance easier to monitor, and the intelligence you deliver at every altitude is accurate enough to act on.

When you start with the decision:

Work redesign gets concrete. Employees can see exactly how their role is evolving, not just that it is. Clear success metrics give your evaluation framework something meaningful to measure. And well-defined use cases tell you precisely what intelligence each level of the organization needs, making it far easier to design systems that deliver it.

When you redesign the work:

Adoption follows. Codified human-AI collaboration creates the accountability structures that allow trust to develop. And when people understand what intelligence they need to do their jobs well, they become active participants in refining the systems that deliver it.

When you earn trust through measurement:

Your capabilities compound. Feedback loops surface data quality issues before they become adoption problems. Validated results make it easier to prioritize the next use case with confidence. And trust drives the adoption that turns AI investment into visible, lasting change and ROI.

When you build intelligence at every altitude:

Value becomes visible to everyone who needs to see it. That visibility sustains the organizational confidence to keep going.

Where Do You Stand with the 5 Patterns?

Use this questionnaire to honestly evaluate where you stand. Don't look at this as an AI maturity grade, but as a way to identify which patterns are already working in your favor and where a focused investment could move the needle most.



Data Gravity

Do we know where our data already clusters?

Can you name the 1 or 2 systems your teams open first when answering a business question? If people name different systems, you may not have data gravity yet.

Do you know the completeness, coverage, and structural integrity of your enterprise data assets?

Do you have baselines and thresholds for data assets and are you tracking them on a defined cadence to ensure gaps or errors are surfaced quickly?




Decisions and Use Cases

Can each AI initiative answer: What decision does it enable?

Can you complete: “This will help [role / function] decide or [specific thing / outcome] faster or better”? If you can only describe a capability like “a chatbot,” the use case is not ready.

Do we have a success definition?

Can you point to a specific metric such as time saved or cost reduced that you will be able to measure at 30, 60, and 90 days? “Improved efficiency” is not a definition.



Work Design

Are we redesigning the work, not just adding AI to it?

Have you redrawn the steps for at least one workflow? A sign you have not: the process looks the same with “AI assists here” inserted at one step.

Do we know which tasks AI handles vs. humans at each step?

Can you label every step as “AI,” “human,” or “human reviews AI?” If the answer is “it depends,” that ambiguity will surface as errors or adoption failure.




Trust and Governance

Do we have a way to evaluate AI output quality continuously?

If the answer to “how will you know if AI degrades?” is “someone will notice,” you do not have an evaluation framework. You need benchmarks, a schedule, and an owner.

Are governance and privacy controls built in, not bolted on?

Were access rules, data lineage, and audit trails defined before the system was built? Controls added post-launch are harder to enforce.



Intelligence

Are insights reaching the right people at the right cadence?

Does your CXO have a recurring performance dashboard and can they see real-time progress? Match format and frequency to needs.

Have we retired at least one manual reporting workflow?

Is there a recurring report someone still compiles by hand? That’s your first automation target. If you cannot name one, you’re adding AI on top of work, not replacing it.

Building an AI-Powered Enterprise

Organizations will invest trillions in AI in the coming years. The difference between transformational results and expensive failures won't be budget size or algorithm sophistication. It will be the approach.

The five patterns are clear:

1. Follow the data gravity, not the hype

Build on a foundation that's already working and expedite your path to value.

2. Start with the decision, not the technology

Create the clarity that turns AI investment from a leap of faith into a provable, repeatable business case.

3. Redesign the work, not just the tools

Build the organizational muscle to sustain AI.

4. Earn trust through measurement

Show that AI is working, with the evidence to keep improving it.

5. Embed intelligence at every altitude

Ensure that AI creates value at every level of the business.

These are proven patterns from hundreds of enterprise implementations. Organizations that embrace them now can build sustainable competitive advantage and deliver customer experiences their competitors can't match.

FOR FURTHER READING

Check Out Our Latest Reports and Guides

Infographic

The Transformation Gap: Why Enterprise AI Investment Isn't Paying Off

DOWNLOAD



Report

Five CX Imperatives for AI That Earns Its Place

DOWNLOAD



Ebook

How Real-Time Analytics Keeps You a Step Ahead of Customer Expectations

DOWNLOAD



dentsu



CARAT

DENTSU CREATIVE

dentsu **X**

iPROSPECT

↳ MERKLE

tag.

About Merkle

Merkle, a dentsu company, powers the experience economy. For more than 35 years, the company has put people at the heart of its approach to digital business transformation. As the only integrated experience consultancy in the world with a heritage in data science and business performance, Merkle delivers holistic, end-to-end experiences that drive growth, engagement, and loyalty. Merkle's expertise has earned recognition as a "Leader" by top industry analyst firms, in categories such as digital transformation and commerce, experience design, engineering and technology integration, digital marketing, data science, CRM and loyalty, and customer data management. With more than 16,000 employees, Merkle operates in 30+ countries throughout the Americas, EMEA, and APAC.

For more information, visit www.merkle.com.

All predictions and assertions made in this playbook are Merkle's opinion and should not be taken as advice or guarantees. Merkle complies with all local privacy laws in processing data, with or without AI. Examples of adaptive experiences and real-time personalization are contingent on lawful basis and contractual permissions and must comply with consent, transparency, and opt-out requirements. All AI initiatives should be handled responsibly, with appropriate testing, governance, and human involvement.