



The CDO's Guide to Context Engineering



Table Of Contents

1.	The New Mandate	3
2.	Why Your AI Initiatives Are Stalling	4
3.	What Context Engineering Actually Means for the Enterprise	6
4.	The Organizational Reality	9
5.	Making the Case to the Board	12
6.	Where Your Existing Investments Get You — And Where They Fall Short	16
7.	A Strategic Framework for Context Engineering	20
8.	Building the Right Team and Operating Model	25
9.	Evaluating Your Architecture Approach	29
10.	Context as Competitive Advantage	33

1. The New Mandate

The CDO role has changed – and the timeline has compressed.

A decade ago, the mandate was to build data infrastructure. Stand up a cloud warehouse. Implement a catalog. Establish governance. That work mattered, and in most enterprises it's still ongoing. But the board isn't asking about infrastructure anymore. They're asking about AI.

The numbers tell the story. Gartner's 2025 CDAO Survey found that 70% of CDAOs now lead their organization's AI strategy. 58% have AI explicitly within their scope of responsibilities, up from 34% just two years earlier. The 2026 AI & Data Leadership Executive Benchmark Survey – covering 110 Fortune 1000 and global leadership brands – is even more definitive: 99% of firms say data and AI investment is a top organizational priority, with 91% increasing their spending. 90% have appointed a CDO, and for the first time, 70% describe the role as “successful and established” – double the figure from 2023.

The CDO has moved from the back office to the center of the enterprise's most consequential bet. And the focus of that role has shifted decisively: 86% of CDOs now report that their primary function is innovation and business growth, up from 62% just three years ago. The defensive era – regulatory compliance, efficiency, risk management – isn't gone, but it's no longer the headline.

The expectations match the elevation. Deloitte's 2025 Tech Value Survey shows 74% of organizations invested in AI and GenAI this year – nearly 20 points higher than any other technology category. The average allocation is 36% of digital initiative budgets. These aren't pilot budgets. This is the board betting on AI

as the path to competitive advantage – and looking to the CDO to deliver.

The challenge is that delivering AI at enterprise scale requires something most organizations haven't built yet: context infrastructure. The data is there. The models are there. What's missing is the layer that makes AI accurate inside your specific business – the business rules, the relationships, the definitions, the institutional knowledge that turn a generic model into one that understands your organization.

That layer is what context engineering builds. And it lands squarely on the CDO's desk – because no other role owns the combination of data, governance, business knowledge, and cross-functional authority needed to make it happen.

This guide is a strategic framework for leading that effort. Not a technical primer, but a playbook for the CDO who needs to make the case, build the team, choose the approach, and deliver results on a timeline the board will accept.

2. Why Your AI Initiatives Are Stalling

The progress is real – and it’s not enough.

The 2026 Benchmark Survey shows AI adoption accelerating at a pace few predicted. Firms with AI in production at scale jumped from 5% to 39% in just two years. 94% now have AI capabilities in some form of production. Only 6% remain in the early stages of experimentation.

And yet. Only 54% of organizations report a high or significant degree of business value from these investments – up from 48% last year, but still barely half. The rest are in the gap: AI is deployed, money is being spent, but the returns aren’t materializing at the level the board expected.

MIT’s GenAI Divide report puts a finer point on it: 95% of enterprise AI pilots deliver zero measurable ROI. IBM’s CEO Study confirmed that only 16% of AI initiatives achieve scale beyond the pilot stage. The pattern is recognizable to every CDO: a successful pilot, an enthusiastic sponsor, a stalled rollout.

The architecture wasn’t built for this

The most telling finding from the Benchmark Survey isn’t about AI at all. It’s about people: 93% of respondents say cultural challenges – not technology – are the greatest impediment to AI adoption. Only 7% blame the technology.

This isn’t just a change management story. It’s an architecture story in disguise. When business processes can’t absorb AI outputs, when

definitions conflict across departments, when the context needed to interpret a question correctly lives in someone’s head rather than in a system – those aren’t “culture” problems. They’re symptoms of an infrastructure that was built for dashboards and reports, not for agents that need to reason across systems in real time.

Your data stack was designed for a world where analysts mediated every question. Data flowed from source systems through ETL pipelines into a warehouse, where humans built queries and BI tools rendered charts. AI agents don’t work that way. An agent receiving a natural language question needs to figure out which systems to query, how to join them, what definitions to apply, which business rules are relevant, and what the person asking actually means – all autonomously, all in real time.

Most enterprise data stacks don’t have a layer that delivers that context. The data is accessible. The knowledge around it is not.

The context gap

Industry survey data paints a consistent picture: while 84% of organizations are working with AI, only 17% are seeing value. Cisco’s 2025 AI Readiness Index – covering 8,000 leaders – found that only 13% qualify as “Pacesetters” with disciplined AI readiness. Those Pacesetters are 4x more likely to move pilots into production and 50% more likely to see measurable value.

What separates the 13% from the rest isn’t model sophistication or data volume. It’s

whether the organization has made its institutional knowledge – the definitions, the rules, the relationships, the decision patterns – accessible to AI.

The Benchmark Survey reinforces this from another angle: 93% of firms say that AI interest has led to a greater focus on data, with 39% saying data has become a top priority. Companies are recognizing that the AI challenge is fundamentally a data and context challenge. But recognizing it and solving it are different things.

The pilot-to-production gap

The failure modes are specific and predictable:

Architectural shortcuts that don't survive production. Pilots bypass existing patterns to move fast – custom connections, hardcoded logic, manual data preparation. When the pilot moves to production, these shortcuts turn into tightly coupled dependencies that break downstream systems.

No ownership of AI output quality. Many teams never define accountability for AI behavior once the system is live. When no one owns output quality – and the output is generated dynamically rather than from a pre-built dashboard – degradation goes unnoticed until it causes damage.

Governance that was skipped during the pilot. A prototype in a sandbox often violates GDPR, SOC 2, or internal data policies the moment it touches production data. The Benchmark Survey shows that while 89% of firms say AI guardrails are now in place, the gap between having a policy and enforcing it consistently across dynamically generated AI outputs remains wide.

Cross-functional alignment that never happened. The pilot had a technical champion

who got it working. Production requires business units to agree on definitions, governance to approve access patterns, and IT to support the infrastructure. Without that alignment, the initiative stays a pilot – technically successful, organizationally stuck.

Why this is a CDO problem

There's a reason AI delivery has become a CDO mandate rather than staying with the CTO or CIO. The models are available. The infrastructure exists. What's missing is the bridge between raw data and business meaning – and that bridge requires someone who understands both sides.

The CDO is the only executive who can connect business terminology with data architecture, enforce governance across organizational boundaries, and navigate the political dynamics of unifying competing definitions. As Gartner states: "No other role than the CDAO has the responsibility of many of the key enablers of AI, which include data governance, D&A ethics, and data and AI literacy."

The 2026 Benchmark Survey offers some encouragement here: 81% of respondents now view the CDO as a permanent C-suite role, up from 71% last year. CDO tenure is lengthening – 26% now have five or more years in the role, up from 17%. The organizational credibility is there.

But so is the accountability. Gartner predicts that by 2027, 75% of CDAOs who fail to demonstrate AI's positive impact will be reassigned or removed from the C-suite. CDO turnover already exceeds 50% within the first two years in many organizations.

The mandate is real. The timeline is short. The question isn't whether to act – it's how to act with a strategy that delivers results using the investments and authority you already have.

3. What Context Engineering

Actually Means for the Enterprise

There's no shortage of buzzwords in enterprise AI. Context engineering is not one of them – or at least, it shouldn't be. It describes something specific, and that specificity is what makes it useful.

In June 2025, Andrej Karpathy – former AI director at both OpenAI and Tesla – wrote a post that reframed how the industry thinks about building AI applications: “In every industrial-strength LLM app, context engineering is the delicate art and science of filling the context window with just the right information for the next step.” Shopify CEO Tobi Lutke championed the concept around the same time. Forrester titled a November 2025 report “The Year Context Became King.”

For a CDO, the definition needs to be more operational: **context engineering is the practice of making your organization's institutional knowledge – business rules, relationships, definitions, decision patterns – usable by AI at scale.**

This is the knowledge that currently lives in people's heads, in Slack threads, in spreadsheets, in BI tools that don't talk to each other. Context engineering is the discipline of making it machine-readable, AI-accessible, and governable.

This is not prompt engineering

The distinction matters. Prompt engineering is tactical – optimizing individual interactions between a user and a model. Context

engineering is architectural – building the systems and infrastructure that deliver the right information to the model at the right time, across every interaction, for every user and agent.

As CIO.com describes it: “Early generative AI was stateless, handling isolated interactions where prompt engineering was sufficient. However, autonomous agents are fundamentally different.” Autonomous agents don't get hand-crafted prompts for each question. They need a system that dynamically assembles the right context based on the question, the data sources involved, and the person or process asking.

Ardoq frames the stakes clearly: “If organizational context is clear, coherent, and well-architected, AI amplifies clarity, accelerates execution, and compounds returns. If context is fragmented, chaotic, and siloed, AI amplifies dysfunction at a terrifying scale and speed.”

That framing should resonate with any CDO who's watched an AI agent confidently return the wrong revenue number to a board member.

The five layers of context

Not all context is equal. Our five-level context architecture provides a useful framework, identifying the layers that help AI answer enterprise questions accurately -- and what each layer contributes:

- **Level 1: Technical metadata** -- schemas, data types, table structures. Most enterprises have this well-covered. On its

own, nowhere near enough for accurate business answers -- AI can find the data but has no idea what it means.

- **Level 2: Relationships** -- join paths, foreign keys, cardinality rules. How tables and systems connect. Often undocumented or locked in ERD diagrams that aren't machine-readable. A meaningful improvement, but still produces results that require heavy manual review.
- **Level 3: Business definitions** -- glossaries, certified terms, ownership. What "revenue" or "customer" actually means in your organization. Partially captured in data catalogs, but coverage is typically incomplete. This is where AI starts producing answers that feel directionally right -- useful, but not yet trustworthy enough for decision-making.
- **Level 4: Semantic layer** -- metric calculations, fiscal calendars, dimensional

hierarchies, business rules. Locked inside individual BI tools, often inconsistent across them. At this level, AI accuracy becomes genuinely useful -- answers are reliable enough that business users start trusting them for real work.

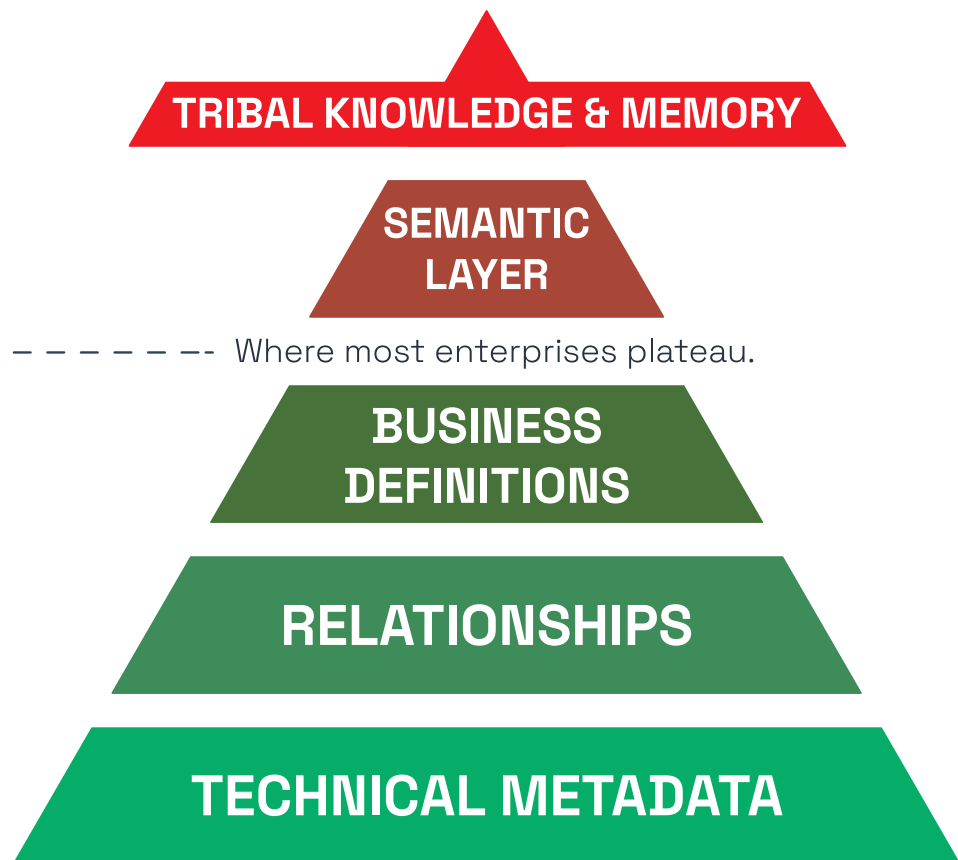
- **Level 5: Tribal knowledge and memory** -- user intent, query patterns, persona preferences, decision precedents. Lives almost entirely in people's heads. The hardest layer to capture, and the one that closes the remaining gap -- producing answers that match what your best analyst would give.

The pattern is consistent: the layers closest to raw data are well-covered and accessible. The layers closest to business meaning are sparse, fragmented, and largely invisible to AI. That's why most organizations plateau in the middle of the stack -- they've solved for the bottom but haven't addressed the top.

Well-covered,
machine-readable



Sparse, fragmented,
invisible to AI



Why this is a CDO problem, not a data engineering problem

Context engineering requires three things that no other role in the enterprise fully owns:

Business knowledge. The semantic layer, business definitions, and tribal knowledge that make AI accurate come from the business, not from IT. Someone needs to bridge those worlds – to translate what “revenue” means to the CFO into something an AI agent can apply consistently. That’s not a data engineer’s job.

Governance authority. Unifying context means resolving conflicts – different definitions of the same term across business units, competing metric calculations, inconsistent access policies. Resolving these conflicts requires organizational authority that goes beyond any single team. The CDO is the executive with the mandate to set and enforce data standards across the enterprise.

Cross-functional coordination. Context doesn’t live in one department. Finance owns some of it. Sales owns some of it. Operations owns some of it. The analyst who’s been at the company for ten years owns a lot of it. Making this knowledge explicit and accessible requires coordination across every function – and a CDO who can navigate the political dynamics of asking business units to codify and share knowledge they’ve historically kept to themselves.

As one enterprise AI advisor put it: “Who owns context engineering in your organization? If the answer is unclear, you have an organizational challenge before you have a technical one.”

The CDO who recognizes this – and acts on it – positions themselves at the center of the organization’s most strategic initiative. The one who waits for someone else to own it will watch AI stall on their watch.

4. The Organizational Reality

Context engineering sounds like a technology initiative. It isn't. It's an organizational initiative that happens to require technology – and that distinction determines whether it succeeds or fails.

The 2026 Benchmark Survey makes this stark: 93% of data and AI leaders say cultural challenges, not technology, are the greatest impediment to AI adoption. That number has been consistently above 90% for most of the survey's fifteen-year history. The technology has transformed; the organizational challenge hasn't moved.

For a CDO leading context engineering, the organizational landscape is the terrain you need to navigate. Here's what it actually looks like.

The political landscape

Data governance initiatives are among the most politically charged projects in the modern enterprise. Unlike a technology implementation that primarily affects IT, context engineering touches every business process, challenges established workflows, and redistributes decision-making authority.

The dynamics are predictable:

Territorial protection. Business units have built their own reporting, their own definitions, their own “single source of truth.” Asking them to unify those definitions with other departments isn't a technical request – it's a power question. Whose definition wins? Whose dashboard becomes the standard? Whose team's work gets superseded?

Resource competition. Every business unit wants more from the data team. Context

engineering requires data team capacity, domain expert time, and governance bandwidth – all of which are already oversubscribed. Prioritizing one domain for context engineering means deprioritizing another, and every stakeholder believes their domain should go first.

Authority challenges. The CDO's mandate is broad but often ambiguous. Unlike the CFO, who has clear authority over financial reporting, or the CIO, who owns the technology stack, the CDO's authority to enforce definitions and standards across business units varies widely. In many organizations, it's influence-based rather than structural – which means every decision requires negotiation.

Change resistance. Senior domain experts have spent years building their understanding of the data. Formalizing that knowledge into a system can feel threatening – as if their expertise is being commoditized or their judgment is being questioned. Getting these people from resistant to invested is one of the most important things a CDO can do, and one of the hardest.

None of this is unique to context engineering. Every CDO who's led a governance initiative, a catalog implementation, or a data quality program has navigated the same terrain. The difference is that context engineering touches more surface area – it requires not just metadata but business rules, not just definitions but decision logic, not just documentation but tribal knowledge. The political footprint is larger.

The talent gap

Context engineering requires a skillset that is rare: people who understand both the business

and the data deeply enough to translate between them.

The 2026 Benchmark Survey shows 93% of firms report that AI interest has led to a greater focus on data – but having the organizational will doesn't solve the talent constraint. Only 29% of organizations have in-house GenAI expertise. 51% of CEOs are hiring for roles that didn't exist last year.

The specific gap is in what you might call "context translators" – people who can sit with a finance team and understand what "recognized revenue" means operationally, then turn that into a machine-readable rule that an AI agent can apply. Data engineers don't typically have the business depth. Business analysts don't typically have the technical vocabulary. The people who bridge both worlds are the most valuable and most overextended members of your organization.

This is also a retention risk. When a senior analyst who knows the institutional history – which tables to trust, which joins are valid, what the CFO actually means when they say "revenue" – leaves the organization, that context leaves with them. Every departure is a context loss event. Context engineering is, in part, a strategy for capturing that knowledge before it walks out the door.

The tool sprawl problem

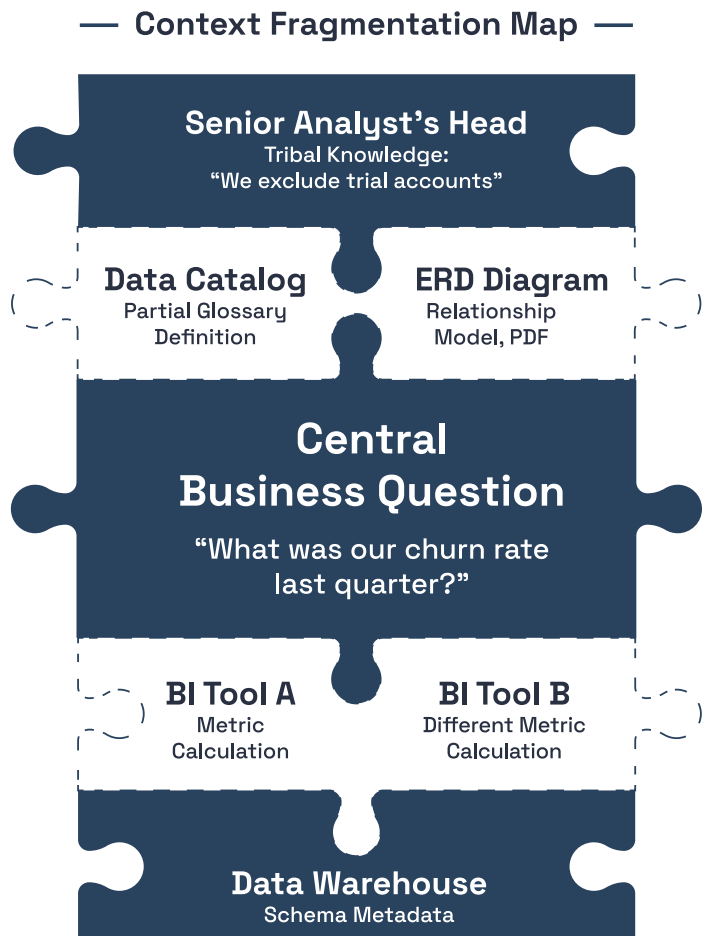
Most enterprises have invested heavily in tools that capture some form of context. The problem is that each tool captures a different slice, in a different format, with different assumptions – and none of them talk to each other.

A typical environment includes 3 to 4 BI tools, each with its own metric definitions and calculation logic. One or two data catalogs, with partial coverage of business definitions. A transformation layer like dbt or a platform-native semantic model. Multiple data

warehouses and SaaS applications, each with their own schemas.

The result is what practitioners call semantic sprawl: the same metric defined differently in different tools, serving different teams, with no mechanism for detecting or resolving the inconsistency. Revenue in Tableau doesn't match revenue in Power BI, which doesn't match what the notebook-based data science team calculates, which doesn't match what an AI agent generates when querying the raw schema.

Each of these tools is doing its job correctly within its own scope. The problem is that no one owns the cross-tool consistency – and AI agents don't stay within the scope of a single tool. They query across systems, combine data from multiple sources, and need a unified set of definitions to produce accurate answers.



The tool sprawl that was manageable in a BI-driven world becomes a critical vulnerability in an agent-driven world.

The tribal knowledge risk

The most valuable context in your organization is the most fragile. It lives in the heads of people who've been around long enough to know the exceptions, the workarounds, the "we always do it this way" logic that no documentation captures.

The impact of codifying this knowledge is well-documented. At one large international professional services firm with over 1,000 clients, only 5% of staff – senior consultants – could effectively analyze client data. After formalizing the tribal knowledge into accessible systems, junior staff could perform 85% of analyses independently (up from 10%), time to complete deliverables dropped by 40%, and onboarding time was cut from six months to six weeks.

That's the upside of capturing tribal knowledge. The downside of not capturing it is that every senior departure, every reorganization, every team change creates a permanent gap in your organization's ability to answer questions accurately. And the pace of change in most enterprises means those gaps accumulate faster than any manual documentation effort can fill them.

Context engineering treats tribal knowledge as a first-class component of the architecture rather than an afterthought – capturing it through query patterns, user feedback, and expert validation rather than relying solely on formal documentation projects that rarely keep pace.

The governance paradox

Here is the tension every CDO faces: AI adoption requires broader access to data. Broader access requires stronger governance. And stronger governance, if implemented clumsily, kills the speed and self-service that make AI valuable in the first place.

The 2026 Benchmark Survey shows 79% of firms say responsible AI investment is a top corporate priority, and 89% report having guardrails in place. But traditional governance was built for a world of structured dashboards and pre-defined reports – where a human analyst vetted every output before it reached a decision-maker. AI-generated insights are different. They're dynamic, created in response to novel questions, and often never reviewed by a human before being acted upon.

This means governance can't be an after-the-fact review process. It has to be embedded in the architecture – enforced at query time, applied consistently across every agent and user, and transparent enough that anyone can trace how an answer was generated.

Context engineering is how you resolve the paradox. A well-built context layer applies the right definitions, the right access policies, and the right business rules automatically – without requiring a human gatekeeper for every query. It enables broad access and strong governance simultaneously, because the governance is encoded in the context itself.

The alternative – manual review of every AI-generated insight – doesn't scale. And the other alternative – broad access with no governance – doesn't survive the first compliance incident.

5. Making the Case to the Board

Every CDO knows the feeling: you understand exactly what the organization needs to build, but the board wants to see a business case — not an architecture diagram.

Context engineering is an infrastructure investment. It doesn't produce a product. It doesn't close a deal. It makes every AI initiative in the portfolio more likely to succeed — which means the business case has to be framed differently than a typical project pitch. The CDO who walks into the board meeting talking about “semantic layers” and “metadata aggregation” will lose the room. The one who walks in talking about the cost of wrong answers, the speed of AI deployment, and the compounding advantage of institutional knowledge will get the budget.

Here's how to frame it.

Infrastructure, not a project

The most important framing decision is positioning context engineering as infrastructure rather than a project. Projects have defined scope, a deliverable, and an end date. Infrastructure is a capability that compounds over time and enables everything built on top of it.

The analogy that resonates at the board level: the shift from on-premises servers to cloud. Nobody built a business case for cloud migration based on one application. They built it on the premise that cloud infrastructure would make every future application faster to deploy, cheaper to operate, and easier to scale. Context engineering is the same — it's the layer that makes every AI initiative in the organization more likely to reach production and more likely to be accurate when it gets there.

This reframing also changes the ROI conversation. Projects get evaluated on direct returns: “we'll spend X and get Y.” Infrastructure gets evaluated on what it enables: “without this layer, every AI initiative continues to stall at 50-65% accuracy and never reaches production. With it, we reduce time-to-production from months to weeks and move accuracy above the threshold where business users actually trust the output.”

The cost of wrong answers

The board understands risk. The most effective business case for context engineering starts not with what it delivers, but with what the organization is losing without it.

Direct cost of AI rework. When an AI system returns inaccurate answers, the failure cascades. Analysts spend hours verifying outputs that should have been trustworthy. Data teams get pulled into debugging queries that broke because the AI joined tables incorrectly or applied the wrong definition. Business users lose confidence and revert to manual processes — which means the AI investment is producing negative productivity, not the gains that were projected.

The Benchmark Survey quantifies the scale of the opportunity: firms that have achieved data-driven status report measurably better financial performance. Organizations in the “AI Achievers” category — those seeing significant business value — are growing the gap with their peers, not closing it. The cost isn't just rework. It's the widening competitive distance between organizations that have made AI work and those still debugging pilots.

Decision latency. When business leaders can't trust AI-generated insights, they fall back on the old model: request a report from the data team, wait for it to be built and validated, review it, ask follow-up questions, wait again. The 2026 Benchmark Survey found that 86% of CDOs now describe their primary mandate as innovation and business growth — but if the organization can't get trustworthy answers in real time, the CDO is presiding over a bottleneck instead of an accelerant.

Compliance and governance exposure. AI systems that generate answers without proper context — applying wrong definitions, ignoring access policies, combining data that shouldn't be combined — create compliance risk that scales with adoption. One incorrect AI-generated report in a regulated industry doesn't just produce a wrong number. It creates an audit trail of an automated system making decisions without proper controls. With 79% of firms saying responsible AI is a top corporate priority and 89% reporting guardrails in place, the gap between stated governance posture and actual enforcement becomes a liability.

Erosion of AI credibility. Perhaps the most expensive consequence is the one that's hardest to quantify. When AI initiatives repeatedly fail to deliver accurate results, the organization develops "AI fatigue" — skepticism toward any AI-related initiative, regardless of its merit. The Benchmark Survey shows that while 94% of firms have AI in some form of production, only 54% report high or significant business value. That 40-point gap represents organizations where AI is deployed but not trusted — and where every new AI proposal faces an increasingly skeptical audience.

Connecting to board-level metrics

The business case should map directly to metrics the board already tracks. Context engineering doesn't create a new category of value — it unlocks value that's currently blocked.

Time to insight. The interval between a business question being asked and a trustworthy answer being delivered. Without context infrastructure, this is measured in days or weeks (analyst-mediated). With it, it's measured in seconds (agent-mediated, with the context to be accurate).

AI deployment velocity. How quickly new AI use cases move from concept to production. Without context engineering, each new use case requires its own context assembly — custom prompts, manual definitions, domain-specific logic built from scratch. With a shared context layer, each subsequent use case inherits the context from previous ones and deploys faster. The second domain is 2x faster than the first. The fifth is 5x faster.

Analyst productivity. Data teams today spend a disproportionate amount of time answering ad hoc questions, validating AI outputs, and debugging incorrect results. Context engineering shifts these teams from answering questions to curating the knowledge that lets AI answer questions — a fundamentally higher-leverage activity.

AI adoption rate. The percentage of the organization actively using AI tools for real decisions. This metric stalls when outputs

can't be trusted. It accelerates when users consistently get accurate, explainable answers. The Pacesetters in Cisco's AI Readiness Index — the 13% with disciplined AI readiness — are 4x more likely to move pilots into production because they've solved the trust problem that keeps everyone else stuck.

Governance posture. The board cares about whether AI systems are operating within policy. Context engineering enables governance that's embedded rather than bolted on — definitions, access controls, and lineage applied automatically at query time rather than reviewed manually after the fact. This shifts the governance conversation from “are we compliant?” to “compliance is built into the architecture.”

Positioning against the alternative

The board will ask: what happens if we don't do this? The answer needs to be specific.

Continued pilot failures. Without context infrastructure, every AI initiative is a standalone project that has to solve the context problem from scratch. The 95% pilot failure rate isn't random — it's structural. It will persist until the underlying architecture changes.

Growing AI skepticism. Each failed initiative makes the next one harder to fund. Organizations that don't invest in context infrastructure will find themselves in a negative cycle: failed projects erode credibility, eroded credibility reduces investment, reduced investment guarantees more failures.

Competitive disadvantage. The Benchmark Survey shows that AI Achievers — the organizations seeing real business value — are pulling ahead. The gap is compounding. Organizations that delay context engineering don't just stay in place; they fall further behind


as their competitors' context layers get richer with every interaction.

Talent attrition. Top data and AI talent wants to work on production systems that deliver impact, not on pilots that get shelved. Organizations stuck in the pilot loop lose their best people to competitors who've figured out how to get AI into production.

Peer benchmarks and market signals

Boards are influenced by what their peers are doing. The market signal is clear:

- **99%** of Fortune 1000 firms say data and AI investment is a top organizational priority (2026 Benchmark Survey)
- **91%** are increasing their AI and data spending year over year
- The jump from **5% to 39%** of firms with AI in production at scale — in just two years — shows the leaders are moving fast
- Gartner's warning that **40% of agentic AI projects will be canceled by 2027** due to cost, value, and risk problems tells the board exactly what happens without the right infrastructure
- Forrester projects global tech spend growing **7.8% in 2026 to \$5.6 trillion**, driven primarily by AI — the investment wave is here, and the question is whether you're building on solid ground or sand
- BARC research shows that while **50% of enterprises are already using AI agents, only 27% have deployed them for BI and analytics** — a significant gap that points directly to the context problem. Agents are succeeding in domains with simpler



context requirements; analytics, where accuracy depends on business definitions, relationships, and institutional knowledge, is where they stall. Closing that gap is what context engineering is for.

The framing for the board: context engineering isn't a new bet. It's the missing layer that makes the AI bets you've already placed actually pay off. You've invested in the models. You've invested in the data. You haven't yet invested in the layer between them — and that's why the returns aren't materializing.

6. Where Your Existing

Investments Get You

— And Where They Fall Short

Most enterprises don't start from zero. Years of investment have gone into building out metadata repositories, governance programs, semantic layers, and business glossaries across the organization. That work is real, and it represents a significant head start.

The core issue is that — much like your enterprise data itself — context currently sits fragmented across the enterprise. Business logic lives in BI tools. Definitions live in catalogs. Relationships live in ERD diagrams or in people's heads. Each system captures a different slice of context, in a different format, with different coverage. None of them were designed to work together, and none of them were built for AI consumption.

Let's take a look at where context currently resides and what each of these systems gives you — and where they fall short.

Data catalogs

What they give you. Data catalogs were the first serious enterprise attempt to make metadata useful at scale. A well-implemented catalog provides asset discovery, business glossary terms, ownership information, lineage at the table or column level, and — in the best cases — usage metrics that indicate which assets are trusted and actively used.

Where they fall short. Catalog coverage is almost always partial. Industry benchmarks

consistently show that business-level descriptions cover 30-50% of enterprise data assets at best. The other half — often the most complex, cross-functional, or recently created assets — remains undocumented. More importantly, catalogs were designed for human consumption: an analyst browsing for the right table to use in a query. They weren't built for AI consumption — an agent that needs to dynamically assemble the right context for an arbitrary question in real time.

The definitions in a catalog are also typically static. They describe what a term means in the abstract (“Revenue: the total income generated from sales”) but don't encode how to calculate it, which filters to apply, which tables to join, or how the definition varies by business unit. An AI agent needs operational context, not dictionary entries.

BI tools and semantic layers

What they give you. This is where some of the most valuable business logic in the enterprise lives. BI platforms like Tableau, Power BI, and Looker — along with semantic modeling tools like dbt, AtScale, and Cube — encode metric calculations, dimensional hierarchies, fiscal calendars, and business rules that analysts have refined over years. This is institutional knowledge in executable form.

Where they fall short. The problem is fragmentation. The average enterprise runs 3 to

4 BI tools, and each one has its own version of “revenue,” its own fiscal calendar logic, its own customer segmentation. These definitions were built to serve specific teams and use cases — not to be consistent with each other.

When an AI agent needs to answer a question, which BI tool’s definition should it use? The answer depends on who’s asking, what they mean, and which data sources are involved — contextual decisions that no single BI tool is equipped to make because none of them has visibility into the others.

There’s also a lock-in dimension. Business logic embedded in a BI tool is typically accessible only through that tool’s interface or proprietary API. It can’t be easily extracted, federated, or reused by an AI agent that needs to reason across multiple systems. The knowledge is there, but it’s trapped.

Data quality programs

What they give you. Clean data is a prerequisite for accurate AI. Data quality programs — profiling, validation, monitoring, remediation — ensure that the values in your systems are correct, complete, and consistent. This work is foundational, and organizations that have invested in it are meaningfully ahead.

Where they fall short. Data quality solves for whether the data is correct. It doesn’t solve for whether the data is correctly interpreted. A table can pass every quality check — no nulls, valid ranges, referential integrity intact — and still produce a wrong answer if the AI applies the wrong definition, joins it to the wrong table, or uses it in a context where a different metric should apply.

Put simply: clean data with wrong context still produces wrong answers. Data quality and context engineering are complementary, not substitutes. You need both.

Data governance frameworks

What they give you. Governance programs establish policies for data access, classification, retention, and usage. They define roles, responsibilities, and approval workflows. In mature organizations, they include data stewardship programs that assign accountability for specific domains. The 2026 Benchmark Survey shows 89% of firms report having AI guardrails in place — a significant achievement.

Where they fall short. Most governance frameworks were designed for a world where humans mediated every data interaction. An analyst submits a request, a steward reviews it, access is granted, the report is built and reviewed before distribution. Governance was a checkpoint in a human workflow.

AI agents don’t go through checkpoints. They receive a question, assemble data from multiple sources, apply logic, and return an answer — all in seconds, with no human in the loop. Governance for AI needs to be embedded in the architecture, applied automatically at query time, not enforced through manual review processes that can’t keep pace with machine-speed decision-making.

The other gap is that governance policies typically address access (who can see what data) but not interpretation (how data should be used, combined, and contextualized). An AI agent may have access to all the right tables and still produce a wrong answer because it lacks the contextual rules that govern how those tables should be combined for a specific question. Access governance and context governance are different problems — and most organizations have only solved the first.

Semantic models and ontologies

What they give you. Some organizations have invested in dedicated semantic models or ontologies — formal, structured representations of business concepts and their relationships. These might be enterprise-wide efforts (a corporate ontology) or domain-specific models (a finance semantic layer in dbt, a product taxonomy, a customer data model). Where they exist, they represent some of the highest-quality context in the organization: explicitly defined, logically structured, and often maintained by dedicated teams.

Where they fall short. Coverage is almost always narrow. A semantic model typically covers one domain or one use case well, but doesn't extend to cross-domain questions.

The finance semantic model doesn't know how to connect to the supply chain model. The product taxonomy doesn't account for how Sales categorizes customers differently than Marketing does.

There's also a maintenance problem. Ontologies are expensive to build and expensive to maintain. When the business changes — a new product line, a reorganization, an acquisition — the ontology needs to be updated, and that update requires specialized skills that most organizations don't have in abundance. The result is models that are accurate for the domain they were designed for but increasingly stale or incomplete at the boundaries.

The honest picture

Here's the assessment in a summary view:

Investment	Context it captures	Typical coverage	AI-ready?
Data catalogs	Asset descriptions, glossary terms, lineage	30-50% of assets	Partial — static, human-oriented
BI tools / semantic layers	Metric definitions, calculations, business rules	Per-tool, inconsistent across tools	No — siloed, proprietary
Data quality programs	Value correctness, completeness, consistency	Varies by domain	No — solves data, not context
Governance frameworks	Access policies, classification, stewardship	Broad but enforcement-oriented	No — designed for human workflows
Semantic models / ontologies	Formal concept relationships, domain logic	Narrow, domain-specific	Partial — high quality but limited scope

The pattern is clear: each investment captures a different slice of context, in a different format, with different coverage, and none of it is optimized for AI consumption. The individual pieces are valuable. What's missing is the ability to unify them — to aggregate context from all of these sources into a single, dynamic layer that an AI agent can use at query time.

That's the gap context engineering closes. Not by replacing your existing investments, but by connecting them — extracting the context they've captured, resolving the inconsistencies between them, and making the unified result available to every agent and user in real time.

What this means for the CDO

The strategic implication is important: you don't need to start from zero, but you can't get there with what you have today. The path forward is aggregation, not reconstruction. The CDO's job is to recognize what already exists, identify the gaps, and build (or buy) the layer that connects them.

The next chapter lays out a phased framework for doing exactly that — starting with one domain, proving the model, and scaling from there.

7. A Strategic Framework for Context Engineering

The chapters above laid out the problem, the organizational reality, and the gaps in your current stack. This chapter is the playbook — a phased framework for building context engineering as an enterprise capability.

The core principle is aggregation, not reconstruction. You don't need to catalog every table, define every term, and map every relationship before you start. You need to aggregate the context that already exists across your tools and systems, fill the critical gaps, and let the system get smarter through usage. Perfection on day one is neither possible nor necessary. What matters is getting to production-grade accuracy fast enough to prove value and build momentum.

Phase 1: Anchor

Objective: Pick one high-value domain, aggregate existing context, and demonstrate measurable accuracy improvement.

This is the proof-of-concept phase — but unlike an AI pilot that works on pre-tested data and falls apart in production, the goal here is to prove that the *context infrastructure* works. The AI model isn't what's being tested. The question is whether assembling context from your existing systems produces materially better answers.

Selecting the right domain. The first domain should meet three criteria:

1. **High business visibility.** Pick something

the board cares about — finance, sales pipeline, customer metrics. If the first win happens in a domain nobody's watching, it won't generate the momentum you need.

2. **Sufficient existing context.** Choose a domain where some business definitions already exist in a catalog or BI tool, where the data model is reasonably well understood, and where you have domain experts who can validate results. You're aggregating existing context, not building it from scratch.
3. **Known pain points.** The best first domain is one where people are already frustrated — where analysts spend hours reconciling numbers, where different teams get different answers to the same question, where executives have lost trust in the data. Solving a problem people already feel is the fastest way to build credibility.

What Phase 1 delivers:

- Aggregation of technical metadata, relationships, and available business definitions from existing systems across the selected domain
- Identification and filling of critical context gaps — the definitions, join paths, and business rules that don't exist in any current tool
- A baseline accuracy measurement and a post-implementation measurement that demonstrates improvement

- Using Prometheus’s five-level context architecture as a benchmark: the goal is to move from Level 1-2 accuracy (technical metadata and basic relationships) to include business definitions and semantic layer within the first domain quickly

The target: demonstrate to stakeholders that context engineering produces measurably more accurate AI answers in a real domain, with real data, on real business questions — fast enough to maintain executive confidence and build momentum for the next phase.

Phase 2: Operationalize

Objective: Turn the anchor domain from a proof of concept into an operational capability with feedback loops, governance, and conflict resolution.

Phase 1 proved the model works. Phase 2 makes it sustainable. This is where context engineering shifts from a project to a capability — and where most organizations either build lasting infrastructure or let the initial win fade.

Feedback loops. Production usage generates signal. When users ask questions, their queries reveal intent — which tables they expect to be used, which definitions they assume, what “revenue” means to them specifically. When they flag an answer as incorrect, that’s explicit context: something in the system’s understanding is wrong and needs to be updated. Both signals — implicit (usage patterns) and explicit (corrections) — need to be captured and fed back into the context layer.

Conflict resolution. As context coverage expands, conflicts will emerge. Finance defines “customer” one way; Sales defines it differently. The BI tool calculates “churn” using one formula; the data science team uses another. These conflicts aren’t bugs — they’re the reality of any organization with multiple functions. The

system needs a formal mechanism for resolving them: domain-level authority (Finance owns the financial definition), use-case-level routing (the Sales definition applies when a Sales user asks), and an escalation path for genuine ambiguities that require human judgment.

Governance integration. Context governance is different from data governance, but it needs to connect to the same framework. Who can modify a business definition? Who approves a new join path? How are changes audited and versioned? These questions need clear answers before the system scales beyond a single domain. The governance model from Chapter 4 — distributed contribution, centralized standards — applies here: domain experts contribute and maintain context, a central team sets the rules for how conflicts are resolved and changes are approved.

Measuring operational health. The key metrics shift from “did accuracy improve?” (Phase 1) to “is accuracy being maintained and are the feedback loops working?” Track: context coverage by layer, time to resolve flagged inaccuracies, number of active contributors, and — most importantly — user trust scores. If people are using the AI-powered answers for real decisions, the system is working. If they’re still exporting to Excel and running their own calculations, it isn’t.

Phase 3: Scale Across Domains

Objective: Expand to adjacent domains, leverage cross-domain context, and accelerate time-to-accuracy for each new domain.

This is where the compounding advantage begins — and where the CDO’s strategic role becomes most critical.

Why the second domain is faster than the first. The first domain required building the infrastructure: establishing connections to source systems, setting up the context aggregation pipeline, implementing feedback loops, defining governance processes. The second domain inherits all of that. It also inherits shared context — fiscal calendars, customer definitions, regional hierarchies, access policies — that doesn't need to be rebuilt.

The acceleration ratio is what matters: each subsequent domain should take meaningfully less time than the one before it, because shared context — fiscal calendars, customer definitions, regional hierarchies — carries over rather than being rebuilt. By the fifth domain, onboarding should be a fraction of the effort required for the first.

Cross-domain context emerges through usage.

When an AI agent answers a question that spans Finance and Sales — “What was the pipeline-to-revenue conversion rate by region last quarter?” — it needs context from both domains simultaneously. These cross-domain relationships aren't typically documented anywhere. They emerge through usage: the system learns that Finance's revenue data joins to Sales' pipeline data through a shared account hierarchy, and that “region” means geographic territory in Sales but legal entity in Finance.

This is one of the most valuable capabilities that context engineering builds over time — and one that's nearly impossible to replicate through manual documentation. No one is going to sit down and pre-map every possible cross-domain relationship. But a system that learns from actual usage patterns builds these connections organically.

The CDO's role in scaling. Scaling is as much a political exercise as a technical one. The CDO needs to:

- **Sequence domains by business impact.** Not every domain is equally valuable. Prioritize based on where AI accuracy has the highest business consequence — typically starting with Finance, Sales, and Customer domains, then expanding to Operations, HR, and Supply Chain.
- **Manage stakeholder expectations.** Each domain's leaders will want to know why they are (or aren't) next. Transparent prioritization criteria — business value, data readiness, stakeholder willingness — help depoliticize the sequencing.
- **Celebrate and publicize wins.** Every domain that reaches production-grade accuracy is evidence that the model works. Share results internally. Let the success in Finance build demand from Sales. Organizational pull is more powerful than top-down push.

Phase 4: Compound (Ongoing)

Objective: Let the flywheel work. Shift from building context to benefiting from context that builds itself.

By Phase 4, the context layer is mature enough that it improves through normal usage rather than requiring dedicated building effort. Every query teaches the system something about intent. Every correction refines a definition. Every new user adds a perspective. Every cross-domain question maps a relationship that didn't exist before.

The CDO's role shifts from driving adoption to governing a system that's largely self-sustaining — ensuring quality, resolving edge cases, and directing the context infrastructure toward new strategic priorities (new business lines, acquisitions, regulatory changes).

Key decision points across all phases:

Phase	Key question	Measure of success
Anchor	Does aggregating context measurably improve accuracy?	Baseline vs. post-implementation accuracy in one domain
Operationalize	Are feedback loops and governance working?	Sustained accuracy, active contributors, user trust
Scale	Is each domain faster than the last?	Time-to-accuracy per domain, cross-domain query success
Compound	Is the system learning from usage?	Context growth rate, reduction in manual interventions



Key activities:

- Pick one high-value domain (finance, sales, customer metrics)
- Aggregate existing metadata and business definitions from current systems
- Measure baseline accuracy vs. post-implementation accuracy

Primary Metric

Accuracy improvement in one domain

Key activities:

- Implement feedback loops (usage patterns + user corrections)
- Establish conflict resolution (e.g., Finance vs. Sales definitions)
- Integrate context governance (who modifies definitions, how changes are audited)

Primary Metric

User trust scores

Key activities:

- Expand to adjacent domains, sequenced by business impact
- Leverage shared context that carries over between domains
- Enable cross-domain queries

Primary Metric

Time-to-accuracy per new domain

Key activities:

- System improves through normal usage without dedicated building effort
- CDO shifts from driving adoption to governing a self-sustaining system
- Direct context toward new strategic priorities (acquisitions, regulatory changes)

Primary Metric

Reduction in manual interventions

Manual approach 6-12 months per domain.

Anchor

Operationalize

Scale

Compound

Platform-accelerated approach weeks per domain.

Anchor

Operationalize

Scale

Compound

A note on timeline

The phases above are sequential, but the time they take varies dramatically depending on approach. Building context manually — cataloging tables, documenting definitions, mapping relationships by hand — takes 6-12 months per domain in most enterprise environments. For an organization with 15 to 20 business domains, that's a multi-year program before you reach meaningful coverage.

A platform-accelerated approach that aggregates existing metadata, automates relationship discovery, and learns from usage can compress Phase 1 from months to weeks — and because shared context carries over between domains, each subsequent phase accelerates further. The difference between a 12-month timeline and a 12-week timeline for the first three phases isn't incremental. It's the difference between a board that sees results this year and one that's still waiting for a roadmap to deliver.

We'll deep dive on this later in the guide. For now, the key point is that the framework is the same regardless of approach — the question is how fast you move through it.

You can't anticipate every question

One of the most important design principles for context engineering — and one that's easy to overlook — is that you cannot anticipate every question your organization will ask.

Traditional approaches to making data accessible try to pre-define every possible query path: build the dashboards, define the metrics, create the reports. Context engineering takes a fundamentally different approach. Instead of trying to predict what people will ask, you build a system that learns from what they actually ask.

This means the architecture has to be designed for dynamic learning from the start. When a user asks a question the system hasn't seen before — a new combination of metrics, a cross-domain join that wasn't pre-mapped, a business term used in an unfamiliar context — the system needs to handle it gracefully and learn from it. Did the user accept the answer? Did they correct it? Did they rephrase and try again? Each of these signals is context that makes the next similar question more accurate.

This is a philosophical shift for many data organizations. The instinct is to document everything upfront, to have complete coverage before going live, to build a perfect ontology before letting AI use it. That instinct is understandable — but it's also the reason manual approaches take years and still leave gaps. The world moves faster than any documentation effort can keep pace with.

The better approach: start with what you have, go to production quickly, and let the system learn from real usage. An 80% answer that improves every day is more valuable than a 95% answer that took two years to build and started decaying the moment it was finished.

8. Building the Right Team and Operating Model

Context engineering doesn't belong to a single team — and the CDO who tries to staff it like one will hit a ceiling fast.

The work cuts across data engineering, analytics, governance, and the business itself. It requires technical depth and business knowledge in the same conversation. It needs contributions from dozens of domain experts who have other full-time jobs. And it has to operate continuously, not as a one-time documentation sprint.

The operating model matters as much as the technology. Here's how to structure it.

The model: distributed contribution, centralized governance

The operating model that works for context engineering mirrors what successful CDOs have already built for data governance — but with a wider scope.

Centralized governance means a small core team sets the standards: how definitions are structured, how conflicts are resolved, how changes are approved, what quality thresholds context must meet before it's published to the AI layer. This team doesn't create all the context. It creates the framework that everyone else contributes to.

Distributed contribution means the people who actually know the business — finance

analysts, sales operations leads, supply chain planners — contribute and validate context within their domains. They're the ones who know that "revenue" in their world means recognized revenue net of returns, that the fiscal quarter starts in February, that the West region was restructured last year and the historical data needs to be handled differently.

This distribution isn't optional. A central team of five or ten people cannot possibly document the institutional knowledge across 15-20 business domains, hundreds of data sources, and thousands of tables. The math doesn't work. The only scalable model is one where domain experts contribute context as part of their existing workflows, guided by centralized standards.

Roles and responsibilities

You don't need to hire a new team. You need to assign clear responsibilities across roles that already exist — and formalize accountability where it's currently implicit.

The CDO (or CDAO). Owns the strategy, sets priorities, secures funding, reports progress to the board. The CDO's job isn't to manage the day-to-day context engineering work — it's to ensure organizational alignment, resolve cross-functional disputes, and keep the initiative connected to business outcomes. This is an executive leadership role, not a program management role.

Context engineering lead. In the early phases, this doesn't need to be a new role — assign

the responsibility to a senior data or analytics leader who already has cross-domain visibility and credibility with both the business and the technical teams. What matters is that someone is explicitly accountable for context quality: the completeness of definitions, the accuracy of relationships, the health of feedback loops, the speed of conflict resolution. As the initiative scales across domains and the workload grows, this may evolve into a dedicated role — but forcing a formal hire before you've proven the model adds cost and organizational friction you don't need yet. Crawl, walk, run.

Domain context owners. One per business domain — a senior analyst or data lead who understands that domain deeply enough to validate definitions, approve changes, and serve as the authoritative voice when conflicts arise. These aren't new hires. They're your existing domain experts with an explicit mandate to contribute to and maintain context within their area. The key is making this responsibility formal and recognized, not an informal ask on top of their existing workload.

Data engineering. Responsible for the technical infrastructure: connecting source systems, building context aggregation pipelines, ensuring metadata extraction works reliably, maintaining performance. Data engineers don't define business context — they build and maintain the plumbing that makes it accessible.

Analytics and BI teams. These teams are the primary consumers of context-powered AI — and also the primary source of feedback. When an AI answer is wrong, the analytics team is usually the first to know. Their role is to validate outputs, flag inaccuracies, and contribute the metric definitions and calculation logic they've built in their BI tools to the shared context layer.

Business stakeholders. Department heads and functional leaders who approve definitions for their domain, resolve conflicts when multiple definitions exist, and — critically — champion

adoption within their teams. Without business stakeholder buy-in, context engineering remains a data team initiative that the business ignores.

Staffing without hiring

The most common objection: “We don't have the headcount for this.”

The answer is that context engineering, done right, *reduces* the total workload on data teams — it doesn't add to it. Today, your analysts spend a significant portion of their time answering ad hoc questions, reconciling conflicting numbers, debugging AI outputs, and explaining why two reports show different figures. That's the tax you pay for not having context infrastructure.

Context engineering redirects that effort. Instead of answering the same question about revenue definitions for the tenth time, an analyst codifies the definition once and the system applies it to every future question. Instead of debugging why the AI joined two tables incorrectly, a data engineer maps the correct join path once and it's enforced consistently.

The staffing model isn't “add new people to do context engineering.” It's “shift existing people from reactive work (answering, debugging, reconciling) to proactive work (curating, validating, governing).” The former scales linearly — every new user generates more ad hoc requests. The latter scales logarithmically — every definition captured reduces future requests.

For most organizations, the practical staffing looks like:

- **1 context engineering lead** (assigned to an existing senior data/analytics leader; formalize as dedicated role only as scale demands)

- **1 domain context owner per priority domain** (existing senior analysts, 10-20% of their time)
- **Existing data engineering capacity** for infrastructure and pipeline work
- **Existing governance team** for standards and conflict resolution
- **No new hires required** for the first three phases — the work is redistributed, not added

Embedding context in existing workflows

The biggest risk to the operating model isn't understaffing — it's asking people to adopt an entirely new workflow on top of what they're already doing. Context contributions need to be embedded in the work people already do, not bolted on as a separate process.

For analysts: when they build a new metric or update an existing one in their BI tool, the definition gets captured in the context layer. Not as a separate documentation task — as a natural byproduct of the work they were going to do anyway.

For data engineers: when they onboard a new data source or modify an ETL pipeline, relationship mapping and schema-level context capture happens as part of the standard pipeline process, not as a follow-up ticket that sits in the backlog.

For domain experts: when they review and approve a business definition, it happens in a lightweight interface connected to their existing tools — not in a separate system they have to remember to log into.

For end users: when they ask a question and the answer is wrong, correcting it should

take one click — a thumbs down, a brief note, a suggested fix. That feedback is context. Making it frictionless is what turns every user into a contributor.

The operating model succeeds when contributing context feels like a natural part of the job rather than an additional burden. If people have to go out of their way to participate, they won't — and the system will starve for the human input it needs to improve.


Managing change

Context engineering touches how people work with data, which means it's a change management initiative whether you plan for it or not.

The most effective change management for context engineering isn't a communications campaign or a training program — it's results. When a sales VP asks a question and gets an accurate, explainable answer in seconds instead of waiting three days for an analyst to build a report, they become an advocate. When a finance director sees that the AI applies the same revenue definition their team uses — correctly, every time — they stop being skeptical and start asking what other questions it can answer.

The CDO's change management playbook:

1. **Start with pain.** Pick the first domain based on where people are already frustrated. Solving a problem they feel gets you advocates without a roadshow.
2. **Show, don't tell.** Don't present slides about what context engineering will do. Demo it. Let people ask their own questions and see the answers. Live demos with real data are more persuasive than any deck.

- 
- 3. Empower the champions.** Every domain has one or two people who are excited about AI. Give them early access, make them the domain context owners, and let them pull their colleagues in. Peer influence beats top-down mandates.
 - 4. Make the feedback loop visible.** When someone flags a wrong answer and the system improves, acknowledge it. People contribute when they see their input makes a difference.
 - 5. Don't force adoption.** Offer context-powered AI as an option alongside existing tools. When it's consistently better, adoption follows naturally. Mandating adoption before the system is ready backfires every time.

9. Evaluating Your Architecture Approach

At some point, strategy meets infrastructure. The framework from Chapter 7 describes *what* needs to happen. This chapter is about *how* — specifically, the architectural decisions that determine whether context engineering becomes a sustainable capability or an expensive experiment.

The most consequential decision isn't build versus buy. It's the architectural model: do you consolidate into a single vendor stack, or do you keep your existing tools and add a unification layer on top?

The consolidation temptation

The major platform vendors have a compelling pitch: bring your data into our ecosystem and we'll handle everything — storage, compute, transformation, semantic layer, AI agents, governance. Snowflake Cortex, Databricks Genie, Microsoft Fabric, and Google's BigQuery stack all offer some version of this vertically integrated story.

The appeal is real. One vendor. One contract. One support relationship. Fewer integration points. And the platforms are genuinely capable — they're investing billions in AI capabilities and shipping features at a pace that's hard to match.

The tradeoff is equally real:

Data centralization. Vertical integration requires consolidating your data — or at least your AI workloads — into one platform. For

organizations with data distributed across multiple warehouses, SaaS applications, on-premise systems, and cloud platforms, this is a multi-year migration project with significant cost and risk. And even after consolidation, the SaaS applications and operational systems that generate the data still live outside the platform.

Context lock-in. When your business definitions, metric calculations, and semantic models are encoded in a single vendor's proprietary format, they become an asset you can't move. Switching costs escalate with every definition you add. The context you've built — arguably your most valuable AI asset — becomes inseparable from the platform you built it on.

Scope limitation. No single platform sees the full picture. A Snowflake-native semantic layer doesn't know about the business rules encoded in your Tableau dashboards. A Databricks-native agent doesn't have access to the glossary definitions in your data catalog. Vertical integration solves the problem within one ecosystem but leaves gaps at the boundaries — and enterprise reality is mostly boundaries.

Timeline mismatch. CDOs facing 12-18 month timelines from the board can't afford a multi-year consolidation project as a prerequisite for AI accuracy. The board wants results now, not after a data migration program that was planned for 2028.

For organizations that are already deep in a single platform ecosystem and plan to stay there, the vertical approach can work. But for the majority of enterprises — those running

3-4 BI tools, 2+ warehouses, dozens of SaaS applications, and some on-premise systems that aren't going anywhere — consolidation is aspirational, not actionable.

The best-of-breed reality

Most enterprises live in a best-of-breed world. They chose Snowflake for the data warehouse, Tableau and Power BI for visualization, dbt for transformation, a catalog for discovery, Salesforce for CRM, NetSuite for ERP — each best in its category, none designed to work seamlessly together.

This is the environment that context engineering needs to operate in. Not a clean single-vendor stack, but a heterogeneous

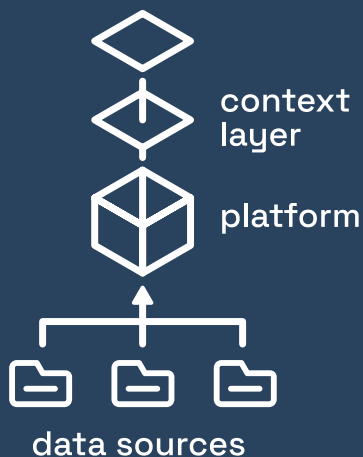
landscape of tools and systems that were selected independently and evolved over time.

The advantage of this reality is that the context already exists — it's just distributed. Business logic lives in BI tools. Definitions live in catalogs. Relationships live in data models. Tribal knowledge lives in query patterns and analyst notebooks. The question isn't whether you have enough context. It's whether you can unify it.

The unification layer approach


The alternative to consolidation is adding a layer that sits across your existing tools and aggregates context from all of them — without requiring you to move data, migrate definitions, or abandon investments that are working.

CONSOLIDATION APPROACH



Requires data centralization, single-vendor lock-in, multi-year timeline.

UNIFICATION LAYER APPROACH



Data stays in place, aggregates existing context, production-grade in weeks.

This approach has specific architectural requirements:

Dynamic context assembly. Rather than requiring all context to be pre-defined and manually maintained, the system needs to assemble the right context at query time — pulling from technical metadata, business definitions, semantic models, and learned patterns based on who's asking and what they need. This is the difference between a static ontology that someone has to keep current and a dynamic system that adapts to the question.

Existing metadata leverage. The system should aggregate context from the tools you already have — extracting business logic from BI tools, definitions from catalogs, relationships from data models, schemas from databases. If a platform requires you to rebuild all of this from scratch, you've lost the primary advantage of the aggregation approach and you're back to a 6-12 month timeline per domain.

Feedback capture and learning. Every user interaction is a context signal. The system needs to capture corrections, learn from query patterns, and improve over time without requiring manual intervention. This is what turns the four-phase framework from Chapter 7 into a flywheel rather than a perpetual maintenance project.

Conflict resolution. When definitions conflict across source systems — and they will — the platform needs a mechanism for resolving them: domain-level authority, user-level routing, and transparent audit trails that show which definition was applied and why.

Federated data access. If the system requires you to centralize your data before it can add context, it's not a unification layer — it's another consolidation play. True unification means querying data where it lives, applying context at query time, and respecting the governance and access policies of each source system.

Full lineage and explainability. Every AI-generated answer needs a traceable path: which data sources were used, which definitions were applied, which joins were executed, which business rules governed the calculation. This isn't just a governance requirement — it's a trust requirement. Business users won't rely on answers they can't verify, and auditors won't accept outputs they can't trace.

The manual path

Some organizations will consider building the context layer themselves. It's worth understanding what that entails.

Manual context engineering means assigning teams to document definitions, map relationships, build semantic models, and maintain all of it as the business evolves. The typical enterprise reality:

- **6-12 months per domain** to reach production-grade context coverage
- **Dedicated resources** — 3-5 full-time equivalents per domain during the build phase, plus ongoing maintenance
- **Cross-domain gaps** — each domain is built independently, so cross-domain relationships require a separate integration effort
- **Maintenance burden** — schemas change, business rules evolve, teams reorganize. Without automated mechanisms to detect and adapt to change, context decays and requires continuous manual updating
- **No learning loop** — manual documentation captures a point-in-time snapshot. It doesn't improve through usage, capture query patterns, or learn from corrections

For an organization with 15-20 business domains, the math is daunting: 6-12 months per domain, sequentially, with compounding maintenance as earlier domains age while later domains are built. This is the multi-year program that Chapter 7's timeline section warned about — and the reason most manual approaches stall after the first two or three domains.

The manual path isn't wrong for everything. Some foundational work — defining core business terms, establishing governance standards, mapping the most critical relationships — needs human judgment and can't be fully automated. But relying on manual effort as the *primary* approach to context engineering at enterprise scale is a timeline and staffing bet that most CDOs can't afford to make.

Questions to ask when evaluating approaches

Whether you're evaluating a platform, a build-it-yourself approach, or some combination, these questions cut to what matters:

On time to value: - How quickly can we demonstrate accuracy improvement in a single domain? - Does the approach require data migration or consolidation as a prerequisite? - Can we start with existing metadata, or do we need to build context from scratch?

On scalability: - Does context from one domain carry over to the next, or is each domain a

standalone effort? - How does the system handle cross-domain queries that span multiple business functions? - What's the realistic staffing requirement to expand beyond the first two or three domains?

On sustainability: - Does the system learn from usage, or does it require continuous manual maintenance? - How are conflicts between competing definitions detected and resolved? - What happens when schemas change, teams reorganize, or business rules evolve?

On governance and trust: - Can we trace every AI-generated answer back to its source data, definitions, and business rules? - Are access policies enforced at query time, or do they require manual configuration per use case? - Does the system integrate with our existing governance framework, or does it create a parallel one?

On preservation of existing investments: - Does the approach leverage the context we've already built in BI tools, catalogs, and semantic layers? - Are we required to abandon or replace any current tools to adopt this approach? - Does it work with data where it lives, or does it require centralization?

The CDO who walks into an architecture evaluation with these questions — and insists on specific, demonstrable answers rather than roadmap promises — will make a better decision than one who evaluates based on feature checklists and vendor demos.

10. Context as

Competitive Advantage

AI models are commoditizing. The performance gap between frontier models narrows with every release. The model that was state-of-the-art six months ago is available as an API for a fraction of the cost today. Within two years, the difference between the best available model and the fifth-best will be negligible for most enterprise use cases.

This is good news — it means the AI capability layer is becoming a utility. But it also means the model is no longer a differentiator. Every competitor has access to the same models, the same architectures, the same cloud platforms.

What they don't have access to is your context.

The real moat

Your organization's institutional knowledge — how you define revenue, how you segment customers, how your fiscal calendar works, which join paths produce valid results, which business rules apply in which situations, how your best analysts interpret ambiguous questions — is unique. It was built over years of decisions, corrections, exceptions, and accumulated experience. It can't be downloaded. It can't be licensed. It can't be replicated by a competitor, no matter how much they spend on models or infrastructure.

Context is the layer that turns a generic AI into one that understands your business. And the organizations that build that layer first don't just get a head start — they get a compounding advantage that widens over time.

The context flywheel

This is the dynamic that separates organizations that get incremental value from AI and those that get exponential value:

Accuracy drives adoption. When AI consistently returns accurate, trustworthy answers, people use it. Not because they're told to — because it's faster and more reliable than the alternative. Adoption isn't a change management problem when the product works.

Adoption drives signal. Every query is a signal. Every correction is a signal. Every time a user accepts an answer, rephrases a question, or flags a result, the system learns something about how the organization thinks about its data. More users means more signal.

Signal drives richer context. Those usage patterns get captured — query intent, definition preferences, cross-domain relationships, persona-specific interpretations. The context layer doesn't just maintain itself; it grows. Relationships that no one would think to document manually emerge through natural usage.

Richer context drives better accuracy. And the cycle repeats. Each pass through the flywheel produces more accurate answers, which drive more adoption, which generate more signal, which build richer context.

The flywheel has a critical property: it's self-reinforcing. Once it reaches a certain velocity, it compounds without proportional additional investment. The organization that starts

the flywheel spinning six months before its competitor doesn't just have a six-month lead — it has six months of accumulated context, usage patterns, and learned relationships that the competitor has to build from zero.

The cost of waiting

Every quarter of delay has a compounding cost that goes beyond the obvious metrics of stalled AI projects and unrealized ROI.

Decision traces are being lost. Every day, your organization makes decisions that reflect institutional knowledge — which data to trust, how to interpret ambiguous results, which exceptions to apply. If those decisions aren't being captured in a context layer, they exist only in the moment and in the memories of the people who made them. When those people move on, the knowledge goes with them.

Query patterns are going uncaptured. Your analysts are asking questions right now that reveal exactly how the organization thinks about its data. Which metrics get combined. Which definitions get applied in which situations. Which cross-domain relationships matter for real business decisions. Without a system that captures these patterns, every query is a learning opportunity wasted.

The competitive gap is widening. The 2026 Benchmark Survey shows that AI Achievers — the organizations reporting significant business value from AI — are pulling away from their peers. The gap between the 39% with AI in production at scale and the rest isn't just a deployment gap. It's a context gap. The leaders have built the infrastructure that makes AI work in their specific business. They're generating signal, capturing feedback, and compounding their advantage with every interaction. The organizations still in the pilot loop aren't standing still — they're falling behind.

The CDO's defining moment

The CDO role has been evolving for over a decade — from data governance administrator, to analytics enabler, to AI strategy owner. The 2026 Benchmark Survey captures the current state: 86% of CDOs now describe their primary function as innovation and business growth. 70% say the role is “successful and established.” The organizational credibility is there.

Context engineering is the initiative that converts that credibility into impact. It's the bridge between the data infrastructure you've built and the AI outcomes the board is demanding. It's the capability that makes every other AI investment in the portfolio more likely to succeed.

And it's the initiative that, done well, makes the CDO indispensable — not as the person who manages data, but as the person who made the organization's institutional knowledge a strategic asset.

The technology is ready. The organizational need is acute. The framework is clear. The question is whether you start building the context flywheel now — or wait, and watch the gap widen.

About Promethium

Promethium enables enterprises to get accurate insights by “talking” to all their business data. Its AI Insights Fabric gives AI models, agents, and business tools seamless access to both distributed data and the context surrounding it — dramatically reducing the cost and effort of context engineering and accelerating the path to production-level AI outcomes.

Rather than requiring data centralization or months of manual preparation, Promethium lets business and data teams build data products on the fly and enable self-service with AI-ready data across all tools, models, and data platforms. The result is a fundamentally more productive data organization — one where institutional knowledge compounds with every interaction rather than decaying between documentation cycles.

A Gartner Cool Vendor backed by Insight Partners, .406 Ventures, and Zetta Venture Partners, Promethium is trusted by leading enterprises like National Grid to deliver faster, more accurate insights at scale. Learn more at promethium.ai.

