

**O'REILLY**<sup>®</sup>  
Report

# Structured for Intelligence

Why AI Needs Governed,  
Discoverable, and  
Provisioned Data

Tom Taulli, Jason Ganz  
& Tom Grabowski



---

# Structured for Intelligence

*Why AI Needs Governed, Discoverable,  
and Provisioned Data*

With Early Release ebooks, you get books in their earliest form—the author’s raw and unedited content as they write—so you can take advantage of these technologies long before the official release of these titles.

*Tom Taulli, Jason Ganz, and Tom Grabowski*

**O'REILLY®**

## Structured for Intelligence

by Tom Taulli , Jason Ganz , and Tom Grabowski

Copyright © 2026 O'Reilly Media, Inc. All rights reserved.

Published by O'Reilly Media, Inc., 141 Stony Circle, Suite 195, Santa Rosa, CA 95401.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (<https://oreilly.com>). For more information, contact our corporate/institutional sales department: 800-998-9938 or [corporate@oreilly.com](mailto:corporate@oreilly.com).

**Acquisitions Editor:** Aaron Black  
**Development Editor:** Gary O'Brien  
**Cover Designer:** Susan Thompson

**Cover Illustrator:** Susan Thompson  
**Interior Designer:** David Futato  
**Interior Illustrator:** Kate Dullea

November 2025: First Edition

### Revision History for the Early Release

2025-08-29: First Release

The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. *Structured for Intelligence*, the cover image, and related trade dress are trademarks of O'Reilly Media, Inc.

The views expressed in this work are those of the author(s) and do not represent the publisher's views. While the publisher and the authors have used good faith efforts to ensure that the information and instructions contained in this work are accurate, the publisher and the authors disclaim all responsibility for errors or omissions, including without limitation responsibility for damages resulting from the use of or reliance on this work. Use of the information and instructions contained in this work is at your own risk. If any code samples or other technology this work contains or describes is subject to open source licenses or the intellectual property rights of others, it is your responsibility to ensure that your use thereof complies with such licenses and/or rights.

This work is part of a collaboration between O'Reilly and dbt. See our [statement of editorial independence](#).

979-8-341-65299-6

[LSI]

---

# Table of Contents

<b>Brief Table of Contents (<i>Not Yet Final</i>)</b> .....	<b>v</b>
<b>1. AI Meets the Data Stack</b> .....	<b>1</b>
How AI Changes the Data Stack	3
Why Structured Context is Critical	6
When AI Gets It Wrong	8
The Industry Response	10
The Changing Role of Data Engineers	11
<b>2. The Structured Context Interface for Governance and Trust</b> .....	<b>13</b>
Defining the Interface Between AI and Structured Context	13
The Tools: Model Context Protocol as the Plumbing	14
The Governance: Making AI Safe and Compliant	17
Building Trust: The Gap Between Adoption and Confidence	18
The Payoff: Why This is Strategic, Not Just Technical	20
Building for What's Next	22

---

# Brief Table of Contents (*Not Yet Final*)

Chapter 1: AI Meets the Data Stack (available)

Chapter 2: The Structured Context Interface for Governance and Trust (available)

*Chapter 3: Discovery—Making Data Findable for AI* (unavailable)

Chapter 4: *Patterns and Priorities in the Enterprise* (unavailable)

---

# AI Meets the Data Stack

## A Note for Early Release Readers

With Early Release ebooks, you get books in their earliest form—the author’s raw and unedited content as they write—so you can take advantage of these technologies long before the official release of these titles.

This will be the 1st chapter of the final report.

If you’d like to be actively involved in reviewing and commenting on this draft, please reach out to the editor at [gobrien@oreilly.com](mailto:gobrien@oreilly.com).

For years, and still today in many organizations, pulling insights from enterprise data means wrestling with a patchwork of tools. BI dashboards, SQL editors, data warehouses, ETL pipelines, and governance systems with each piece comes with its own interface, quirks, and learning curve. Most business folks lack the technical chops to use these tools on their own, so they lean on analysts or engineers to get the job done. This usually means waiting days or even weeks for reports, only to find the insights have gone stale by the time they arrive. The system technically functions, but in practice, it’s slow, inflexible, and anything but user-friendly.

The numbers paint a stark picture of just how fragmented today’s data landscape remains. Enterprise analytics teams are now working across an average of **400 data sources**. At the upper end, nearly one in five enterprises juggle more than 1,000. And these figures come

from organizations with at least 1,000 or more employees, making these numbers even more striking—these aren't small companies struggling with limited resources, but established enterprises with significant IT investments.

A 2024 industry survey adds operational texture to the picture. It revealed that more than **70% of data teams** rely on five to seven different tools just to get through their daily workflows. About 10% are juggling more than ten. The result is mounting cognitive overload and constant integration headaches that bog down decision-making. For end users, the experience is often just as frustrating. They need to navigate a patchwork of platforms, which makes it harder to find the data or insights they need, eroding the return on the data they collect.

The productivity toll is measurable across roles. In the **2025 State of Analytics Engineering Report**, even though 70% of analytics and data professionals use AI to help write code and documentation, 57% still spend most of their time maintaining or organizing datasets, the same level as in the prior year. The promise of AI assistance augmentation hasn't yet freed these professionals from the grunt work. Data scientists face similar challenges, spending about **60% of their time** just cleaning and organizing data. Another 19% spent gathering the datasets. This leaves roughly 20% for actual analysis and insight generation, the part of the job that drives business decisions and growth.

Utilization lags accordingly. Another survey has found that between **60% and 73%** of all enterprise data never gets used for analytics. And it's not just about volume. It's about access and alignment. **83% of respondents** said their organizations suffer from data silos, and nearly all of them (97%) believe **those silos are hurting performance**. Many users don't even know what data exists within their organization, let alone how to access or apply it.

This fragmented, siloed reality is the starting point. It's where most enterprises find themselves today—wrestling with complexity instead of extracting value.

But AI is about to change everything.

# How AI Changes the Data Stack

AI is flipping the traditional model on its head. It's not just making data easier to access—it's changing how the entire data stack operates. Smart features can be woven into every layer, like natural language interfaces that let you skip the SQL, or AI assistants that handle data preparation and analysis. This makes working with data feel more like having a conversation than wrangling code. This isn't just a nice upgrade; it's a reinvention of how companies get value from their data.

## From Dashboards to Dialogue

For much of the modern data era, business intelligence (BI) has been defined by static dashboards and the technical expertise required to navigate them. Pulling meaningful insights often meant waiting in line for scarce data analysts to run SQL queries or create tailored reports. This system empowered only a fraction of the workforce—those with the tools and training to interpret raw datasets. For the rest, decision-making often lagged behind, hampered by delays and a dependency on intermediaries.

But this bottleneck is now easing. AI-powered conversational analytics lets users ask in plain English and iterate in real time while preserving definitions and controls behind the scenes, rewriting how organizations query, explore, and act on data. The global conversational AI market hit **\$13.2 billion in 2024** and is projected to soar to \$49.9 billion by 2031, growing at an annual rate of nearly 25%. Natural language processing (NLP), the backbone of these conversational interfaces, is expanding even faster, with a **projected CAGR of 38.7%**.

What's even more revealing than the market growth is how quickly enterprise behavior is shifting. Gartner **predicts** that by 2025, natural language will be the main way people interact with data systems. That change alone is expected to drive a 100x surge in data usage across organizations. As access improves, the value of data doesn't just rise; it multiplies.

## Beyond Conversational Analytics

Beyond one-off questions, agentic AI coordinates multi-step work: planning, writing code or SQL, running checks, and proposing

changes and is poised for explosive growth. Research from Capgemini found that 50% of enterprises plan to **implement AI agents in 2025**, with adoption expected to reach 82% in 2028. Nvidia CEO Jensen Huang has **said** that this technology is “a multi-trillion-dollar opportunity.”

True, such predictions should be taken with a grain of salt. This is especially the case with dynamic categories like agentic AI. But as seen with the growth in usage of tools like ChatGPT, Claude, Gemini, and Microsoft Copilot, user expectations have shifted. The natural language interface has not only become a standard for AI systems but also a key feature for many traditional applications.

Something similar may happen with agentic platforms. As major AI developers roll out new capabilities, they will be exposed to enormous user bases. This will help cement expectations for interfaces. Users will become accustomed to systems that solve problems autonomously in the background, with periodic queries for approvals. Expectations are normalizing around systems that don't just answer, they act within guardrails. It's why enterprises need to keep an eye on the emerging trends with UIs.

## The Evolution of Data Interfaces

“What were our top-selling products last quarter?”

On the surface, this is a straightforward query. But until recently, answering it was extremely challenging. It required SQL knowledge, database access, understanding of table structures, and often multiple queries to aggregate and filter the right data.

That's now changing. The rise of natural language interfaces is beginning to democratize access to enterprise data, allowing anyone—not just analysts or developers—to ask questions in plain English and get meaningful, real-time answers.

A quiet revolution is happening in how we interact with technical systems. While this trend is still in its nascent stages, there are some innovative tools shedding light on what to expect for the future.

## A Glimpse into the Future with Agentic Development

Modern AI IDEs and notebooks point to the pattern: take Cursor as an example of where interfaces are heading. Since launching in March 2023, it's grown from zero to **\$500 million ARR**, hit a \$9

billion valuation, and now processes **over one million queries per second** while producing nearly a billion lines of production code daily.

What Cursor shows us is the shift from manual prompting to goal-driven collaboration. Users state what they want. The system figures out how to get there, proposes a plan, drafts code, runs tests, and opens a PR.

Here's what this could look like for data engineering. Suppose you've been assigned to build a weekly ETL pipeline to aggregate customer activity, enforce data quality standards, calculate summary metrics, and push the final output into production. Under traditional workflows, you'd be piecing together SQL queries, Python scripts, data validation routines, and CI/CD configurations by hand. But in an agentic system, the process starts differently. You define your goal in natural language:

“Create a weekly customer activity ETL pipeline. Include data quality checks for nulls and duplicates, calculate weekly active users, and push summary tables to the analytics warehouse.”

From that point, the AI agent gets to work. It scans your project, taking into account schema definitions, naming conventions, current pipeline structures, and your warehouse configuration. This comprehensive understanding allows it to outline a detailed, coherent plan. For instance, it proposes creating a new model, *customer\_weekly\_activity.sql*, drafting Python scripts for anomaly detection, and preparing orchestration configs aligned with your tech stack. The plan comes annotated with rationale and implementation steps.

Once you approve it, the agent transitions seamlessly into automated development. It writes out the SQL aggregation logic, test scripts to check for nulls and duplicates, CI/CD configurations tailored to your stack, and optional README updates. Everything is formatted to match your project's style guidelines, versioned correctly, and staged for review.

Next comes execution and validation. The agent runs the full pipeline in a sandboxed or development environment. It executes the SQL, initiates the data quality scripts, and runs your test suite along with any CI checks. Upon approval, it applies the change and re-

runs the pipeline. This kind of real-time feedback loop ensures that problems are caught and resolved before human review, not after.

Finally, when the pipeline passes all checks, the agent handles deployment. Depending on your preferences, it might open a pull request with a summary of changes and test results, or directly push the updates to your orchestration layer. You receive a preview of diffs, validate the output, make any final adjustments, and merge. You do this without ever needing to write deployment scripts or manually trigger the pipeline.

Such a sophisticated system would lead to high levels of productivity and agility.

However, for it to operate safely and autonomously in complex environments, you need more than just instructions and advanced AI. There must be structured data and metadata (or structured context): the schemas, semantics, relationships, permissions, and lineage that describe how your data works.

## Why Structured Context is Critical

GenAI has earned its spotlight largely thanks to what it can do with unstructured data. These models are impressive at summarizing documents, generating content, translating languages, and pulling insights from dense, text-heavy sources like emails, research papers, and internal reports. The ability to sift through and make sense of unstructured information has captured considerable attention, and rightly so.

But here's where things get interesting. The emergence of agentic AI has pushed enterprise companies to rethink their entire data foundation. They're no longer just chasing better conversational interfaces. They're gearing up for intelligent enterprise automation at scale.

In this shift, structured context moves from “nice to have” to “non-negotiable.” You can think of this as the operating system for the next generation of AI. Enterprise automation depends on *structured context*; without it, agents can't function safely or effectively.

Unstructured data carries facts like call recordings or chat logs and can generate insights, sure, but structured context encodes meaning, quality, and policy (models, sources, lineage, dependencies, tags,

permissions) that drive critical business decisions. Agents need both to discover the right tables, understand relationships, plan safe actions, and explain results. Without structured context, agents remain helpful chatbots; with it, they become reliable teammates that can move beyond answers to actions.

## The Two Critical Integrations

Something that's important is having a system that integrates with the following:

### *Structured data*

This refers to the rows in your data warehouse or lakehouse. This is information from customer relationship management (CRM), enterprise resource planning (ERP), human capital management (HCM) and other enterprise systems.

### *Structured metadata*

This includes data about models, sources, lineage, dependencies, tags and governance tags. This is essentially a map of your data ecosystem.

With structured data and metadata, AI agents can discover tables, understand relationships, check quality, and plan safe actions. This allows for powerful conversational analytics, which is powered by planning, reasoning, and autonomous decision-making. But it also has a layer of governance. It is what keeps actions safe: policies and permissions embedded in metadata that determine who can access which datasets, which transformations are permitted, and how sensitive fields must be handled. In short, structured context is the difference between plausible answers and *trustworthy actions*. For example, suppose someone in marketing has a question about churn rates and prompts a standard chatbot about it. The answer will show the details about its origin, as well as the definition of the terms. This will be based on the underlying structured metadata.

Another use case involves AI-powered agentic development. A data engineer can leverage the structured metadata to carry out a complex data migration. The system will understand how to manage features like stored procedures and the differences among the databases.

## Three Key Capabilities

Structured context equips agents with three key capabilities:

1. *Memory* via metadata, so they know what assets exist and how they relate.
2. *Boundaries* via clear definitions, permissions, and rules so they don't wander outside guardrails
3. *The ability to take useful actions with validated tools to read/write safely.*

When you combine these, agents evolve from being just chatbots and start acting like reliable teammates. They can plan, reason, and execute tasks at scale. We're already seeing agents autonomously modify data pipelines, fix errors, manage migrations, and spin up new data products. They can be always-on, embedded in the infrastructure. All this is driven by structured inputs and aligned with business logic.

That's the real unlock. As companies move from pilot experiments to production-grade agentic AI, this structured foundation ensures their systems aren't just delivering impressive responses, but taking trustworthy, business-critical actions. The magic happens when you combine agentic AI with structured metadata, creating scalable, accurate, and governed systems that organizations can truly rely on. The practical "how" of exposing structured context to AI systems is covered in the next chapter.

## When AI Gets It Wrong

Weak governance and poor inputs have predictable consequences. Inside many companies, AI initiatives underperform: studies often cite that between **70% and 80% of AI projects don't succeed**. That's not just a little worse than traditional IT projects; it's nearly double the failure rate. And in most cases, it comes down to something incredibly basic: bad data. Models built on unreliable or incomplete information don't just underperform; they can quietly drain revenue. Gartner puts the average cost of poor data quality at around **\$12.8 million a year per organization**. In some cases, companies lose as much as **6% of their annual revenue from flawed AI outputs**.

## Real Consequences of AI Failures

High-profile examples show the impact. For example, a major airline was **taken to court** after its chatbot promised a bereavement fare refund to a customer who had just lost a loved one. The customer followed the advice, only to be told later that the discount didn't exist. Oddly, the company argued that the chatbot was its own legal entity and responsible for its own behavior. The tribunal didn't buy it. The airline had to pay the customer over \$600 and was held accountable for the AI's error.

Lawyers have faced sanctions for submitting briefs citing fictional cases. News sites have published AI travel content directing readers to unsafe destinations. And it's getting worse—OpenAI's newest models **hallucinate at higher rates** than predecessors, with error rates hitting 48% and 33%.

The pattern is hard to ignore. As organizations rely more heavily on AI to make sense of their data, the quality of that data and the structures in place to manage it become mission-critical. Generative models can be powerful tools, but without clear constraints and reliable inputs, they risk misleading as much as they help. When AI gets it wrong, it's not just a glitch. It can lead to legal exposure, reputational damage, and costly operational missteps. The response is not merely "better models," but *disciplined governance*, like quality checks, lineage, permissioning, and policy enforcement throughout the pipeline.

## Regulatory Considerations

Of course, when an AI implementation fails, it can trigger serious regulatory consequences. Regulators are making these guardrails explicit. Under the **EU AI Act**, especially Articles 10 and 27, these failures point to deeper compliance risks that organizations cannot afford to ignore.

**Article 10** mandates that high-risk AI systems use datasets that are complete, accurate, representative, and error-free. Organizations must document everything: data sources, annotation methods, quality checks, and bias mitigation. Without this documentation, you're non-compliant.

**Article 27** goes further, requiring Fundamental Rights Impact Assessments (FRIAs) that examine fairness, dignity, and non-

discrimination—especially critical for AI in hiring, credit scoring, healthcare, or education. Companies must map data flows, retention policies, oversight mechanisms, and risk mitigation, sometimes reporting to regulators.

While FRIAs can be aligned with DPIAs to avoid redundancy, they still require organizations to map out not just where data comes from and how it's used, but also how it might affect people, including vulnerable groups. Companies must document data flows, retention policies, oversight mechanisms, and risk mitigation steps. In some cases, they're also required to report FRIA outcomes to regulatory bodies, especially if fundamental rights could be impacted.

Regulators are also pushing for continuous oversight. This means companies must implement full data lineage tracking, robust metadata infrastructure, role-based access, and systems capable of automatically flagging issues as they arise. If a dataset changes, such as with schema shifts, updates in source systems, or inconsistent labeling, organizations are expected to respond in real time.

What the EU AI Act makes clear is that data quality and governance can no longer be treated as just an IT or performance issue; they are legal obligations that need to be baked into every layer of the AI pipeline.

Articles 10 and 27 effectively turn structured data governance into a compliance discipline. To meet these standards, companies must build systems that deliver not only clean and reliable data but also transparency, traceability, and accountability. They need to be embedded into every layer of the AI pipeline. Real compliance means more than checking boxes. It means engineering a governance framework that is automated, auditable, and built to scale.

## The Industry Response

The good news is that the technology industry is actively working to fix what's broken in this new AI-infused era. The influx of venture capital into the space is proof enough.

One major trend is the consolidation and unification of data stacks. After years of best-of-breed fragmentation, many organizations are seeking more unified platforms to reduce complexity. Vendors are converging on this context-first, governance-forward model. Quality and policy controls are moving closer to where queries run

and transformations execute, so guardrails are enforced in the flow rather than audited after the fact.

Industry observers predict that GenAI will drive a shift from highly fragmented stacks toward unified data platforms. Data quality and observability are receiving an AI boost as well. Monitoring dozens of pipelines and tables for issues can overwhelm human teams. So vendors are adding AI/ML capabilities to detect anomalies or quality issues in real-time. New features in cloud data platforms can automatically monitor freshness, null spikes, or other data health metrics and alert teams before issues propagate downstream.

Behind these improvements sit the same ingredients emphasized above: well-modeled data, accurate metadata, enforceable policies, and continuous signals about data health.

## The Changing Role of Data Engineers

No doubt, this deeper integration of intelligence into the data stack signals dramatic changes for the role of data engineers. Tasks that once defined the profession, such as building ingestion pipelines, wrangling schemas, writing ETL logic, monitoring data flows, and managing break-fix incidents, will increasingly be handled by intelligent systems. These systems will go beyond routine automation, learning to optimize themselves, detect anomalies in real time, resolve issues independently, and enforce governance frameworks.

Yet this evolution does not mean data engineers are becoming obsolete. As routine responsibilities fade into the background, engineers are shifting their focus to more strategic concerns. They're now tasked with designing resilient and adaptive data architectures, validating the integrity and semantics of AI-driven pipelines, maintaining rigorous standards for data quality, and embedding ethical and compliance principles into the core infrastructure.

In this new environment, data engineers are evolving from system operators to system stewards. The work now demands fluency in AI-native tools, semantic data modeling, governance strategy, and the supervision of autonomous agents. Collaboration is also evolving, as engineers engage more deeply with product, compliance, and analytics teams to ensure that data infrastructure aligns with broader organizational goals. The shift isn't about doing less, but about doing more of what truly matters.

But for this new era of data engineering to become reality, there must be a solid foundation where LLMs effectively interact with structured data and metadata. Otherwise, the capabilities remain shallow and disappointing, as they're not grounded in the relevant data, processes, and workflows of the enterprise.

*So what?* AI is reshaping the enterprise data stack from the ground up. What used to be a rigid setup, built for specialists and stitched together with siloed tools, is starting to feel more conversational, smarter, agentic, and quicker to respond. It's about weaving intelligence into every layer: data ingestion, governance, and how users interact with the data. But for all that to work, you need a modern metadata foundation. Success won't come just from picking the right AI tools. It'll come from rethinking the data stack itself, designing it to support trust, agility, and growth in a world where AI is quickly becoming the norm.

---

# The Structured Context Interface for Governance and Trust

## A Note for Early Release Readers

With Early Release ebooks, you get books in their earliest form—the author’s raw and unedited content as they write—so you can take advantage of these technologies long before the official release of these titles.

This will be the 2nd chapter of the final report.

If you’d like to be actively involved in reviewing and commenting on this draft, please reach out to the editor at [gobrien@oreilly.com](mailto:gobrien@oreilly.com).

## Defining the Interface Between AI and Structured Context

In just a few years, prompt engineering has turned into a cottage industry. There are many books on the topic, as well as seemingly endless blogs and LinkedIn posts. Then there are the workshops, YouTube videos and courses.

The major LLM providers, such as OpenAI, Anthropic, and Google, have published their own guides. For example, Google has written a [68-page handbook by Lee Boonstra](#). It covers topics like how to change the creativity of responses and craft prompts for multi-step reasoning.

Yet all this points to something interesting: prompt engineering is a clear sign of the limitations of chat-based LLMs. They are brittle, as they rely on probabilistic transformer models. After all, with an application like ChatGPT, Claude, or Gemini, you will likely get a different response with the same prompt. Even minor changes in a word or punctuation can have wildly different outputs.

True, this unpredictability is fine for summarizing long PDF documents or evaluating customer reviews. But it is far from ideal when it comes to working with complex enterprise workflows, which need to be reliable and consistent.

As for agentic AI, this certainly holds much promise to address the issues of chat-based interfaces. It should mean that prompt engineering will fade in importance. But successfully deploying agentic AI systems will be challenging. According to a [report from Gartner](#), 40% of these projects will be cancelled by the end of 2027. The firm lists a variety of reasons for this, including the high costs, ineffective risk controls, and unclear business objectives. Despite this, Gartner remains optimistic about agentic AI. The firm forecasts that at least 15% of daily work decisions will be made autonomously by 2028.

But success with generative and agentic AI for data engineering requires a structured context interface, a reliable way for AI systems to interact with structured data and structured metadata under policy. This will unlock major gains in productivity and agility because it brings consistency, permissions, lineage, and definitions directly into the loop where assistants and agents operate.

## The Tools: Model Context Protocol as the Plumbing

Simply put, a structured context interface is about the interactions between structured context and AI, which is usually an LLM or small language model (SLM). There are different ways to make the connections. But the one that has been gaining significant adoption is model context protocol (MCP).

### Understanding MCP

In November 2024, Anthropic introduced the open [MCP standard](#). The company defined it as a way “for connecting AI assistants to the

systems where data lives, including content repositories, business tools, and development environments.”

Before this, the development of agentic applications was cumbersome because of the need to write integrations to implement tools with LLMs. This became known as the “ $M \times N$  problem.” That is, for every  $M$  LLMs that must be supported, you need  $N$  connectors for each tool.

For example, suppose we are building an HR assistant. Let’s say you will use two models, one from OpenAI and another from Anthropic. The application will also use five tools: a payroll API, calendar, employee directory, ticketing system, and knowledge base. For this, you will have to build ten connectors, such as by using OpenAI’s and Anthropic’s SDKs.

By contrast, using MCP changes the formula to  $M+N$ . This is how it works for our HR application:

- You will deploy five MCP servers or one for each tool
- You will use two MCP clients. This includes one for OpenAI and Anthropic.

This has reduced the connectors to seven. However, there are more advantages to using MCP. Keep in mind that there are thousands of servers available as open source repositories. Moreover, OpenAI, Microsoft, and Google have adopted the standard.

Another benefit of MCP is that there is a clean separation of responsibilities between the server and client. The server focuses on the backend tool logic and data access, whereas the client makes connections with a JSON-RPC interface. This modularity provides for more maintainability and scalability, as there is much less complexity when adding models and tools.

## MCP Alternatives and Ecosystem

True, in terms of a structured context interface, MCP is not the only approach for integrations. There are alternatives. But currently, MCP is in the leadership position and is building a strong ecosystem, which includes a growing set of open-source servers and clients. As is the case with technology standards, this will help to build powerful network effects.

Something else to consider: a structured context interface is essentially an architectural pattern. This means there are various methods for the implementation. For example, dbt can operate as an MCP server, allowing for integrating AI applications with data warehouses. It's available to users via Cloud CLI, API, the Discovery API, and Semantic Layer. With it, you can access private APIs, text-to-SQL, and SQL execution. What this means is that you can connect a dbt project with structured data and metadata to any MCP client, like Claude Desktop Projects, Cursor, custom apps, or agent frameworks. This makes structured context addressable and enforceable through a consistent, auditable interface.

## How the Interface Works in Practice

With this structured context interface, an LLM will recognize when it needs to answer a question with structured data and metadata and which tool to call. These are some use cases:

### *General knowledge*

*Prompt:* What is the capital of the United States?

*Response:* The LLM will know the answer based on its training data and model knowledge.

### *External or real-time queries*

*Prompt:* What is the weather going to be in Washington, DC, tomorrow?

*Response:* The LLM will understand that this will require external information that is not likely to be in structured data. Instead, there will be a call for a web search to retrieve the weather data.

### *Enterprise metrics*

*Prompt:* How many customers do I have in Washington, DC?

*Response:* The LLM will recognize the need for structured context. For this, there will be a call for the `query_metrics` tool to access the correct information.

However, in an enterprise environment, a structured context interface needs to carry out these types of functions with strong governance and trust, whether a human or an agent initiates the change. Let's look at this in more detail.

# The Governance: Making AI Safe and Compliant

Once assistants can act (not just answer), governance becomes a prerequisite. Effective governance is at the heart of any scalable AI system. This ensures clear, transparent, and secure boundaries.

Yet despite its importance, most organizations are falling short. **Gartner's 2025 survey** paints a stark picture. Only 12% of companies have put a dedicated AI governance framework in place, and 55% still have no formal structure at all.

This gap comes with real consequences. According to Gartner, poor governance can lead to increased costs, failed AI projects, and damaged reputations. Without guardrails, AI and agent-based systems can introduce bias, violate privacy laws, trigger regulatory fines, and erode trust among both users and customers.

**Deloitte's research** adds another dimension. For those organizations that implement an effective governance system, they tend to see higher AI adoption rates and stronger revenue growth. In short, well-governed AI programs not only keep your organization out of trouble, they also build stakeholder confidence and generate tangible business value.

## Governance Best Practices

Here are a few best practices to keep top of mind with governance:

### *Policy enforcement and access control*

Role-based and attribute-based controls should govern AI agents and human users alike. This ensures only the right people and processes access sensitive data.

### *Lineage, versioning, and auditing*

Track every dataset change from ingestion to transformation to final consumption by AI agents.

### *Embedded structured metadata and quality scores*

Your structured context interface should include rich structured metadata that agents can surface alongside the data. This empowers both agents and users to assess the relevance and credibility of inputs.

### *Legal and regulatory compliance*

Codify policies to mask or anonymize personally identifiable information (PII), enforce data minimization, and honor user rights under laws like GDPR, including erasure requests.

## **Implementation in Modern Interfaces**

A modern structured context interface will enforce these approaches. For example, MCP supports OAuth 2.0 as its main authentication and authorization system. Next, a semantic layer will define an organization's metrics and dimensions, allowing for strong governance. Finally, a structured data interface will have SQL validation. This can be built into an MCP server.

## **Building Trust: The Gap Between Adoption and Confidence**

As AI systems start making more decisions on their own, trust has never been more important. Without it, even the most sophisticated tools can end up unused. A recent KPMG study, [The American Trust in AI Paradox: Adoption Outpaces Governance](#), published in 2025, makes this clear. While 70% of U.S. workers are leveraging AI's benefits and 61% say it's already improving their jobs, about 75% still feel uneasy about its potential risks. In fact, only 41% say they truly trust these systems.

This gap in trust shows that simply rolling out AI tools does not guarantee they'll be integrated effectively. Without strong governance, transparency, and clear controls, these tools can easily be pushed aside.

## **The Shadow AI Problem**

The survey also found that nearly 44% of U.S. workers are using AI outside official company channels. This often involves uploading sensitive data to public platforms. What's more concerning is that 58% rely on AI outputs without checking for accuracy or compliance, and 57% admit they've made mistakes because of AI in their daily work. Then there are 56% who say they hide their use of AI from their managers. They will pass off AI-generated work as their own.

All of this points to a deeper problem. When policies are unclear, training is lacking, and oversight is limited, AI tools turn from productivity boosters into hidden risks. Without greater transparency, solid evaluation processes, and a culture that treats AI as a governed partnership rather than a secret shortcut, organizations will continue to struggle with fragile trust and results.

## Core Pillars of Trust

Trust is about having systems and practices that give both people and AI agents confidence that data and outputs are accurate, explainable, secure, and aligned with the organization's goals and values. Here's a look:

### *Data quality and validation*

Trust starts with reliable data, which must be tested, validated, and monitored at every step. Automated checks for uniqueness, unexpected schema changes, distribution shifts, and stale or unusual data act as constant safeguards.

### *Explainability and transparency*

AI outputs only have real value if people understand how they were generated. This means using explainability tools that show the reasoning behind model predictions. They should also be backed by rich metadata and context.

### *Continuous oversight and feedback loops*

Trust must be maintained over time through constant monitoring of model performance, drift detection, and dashboards tracking indicators like test pass rates, incident resolution times, and metadata completeness.

## Building Trustworthy AI Starts with a culture of transparency

For any enterprise building AI systems people can trust, it starts with a culture of transparency and alignment. Trust needs to be woven into everyday teamwork.

More organizations are seeing the value of bringing different groups together. Data engineers, analysts, privacy and legal teams, and business leaders are working side by side to define what "trusted" really means, whether that's agreeing on clear data definitions or setting the right risk limits. When teams align on these fundamentals from

the start, they break down silos, cut down on misunderstandings, and steer clear of ethical or regulatory pitfalls.

Embedding trust early in development or “trust by design” is another critical success factor. This includes building bias checks, structured metadata tagging, and explainability reviews directly into CI/CD pipelines. These proactive measures surface potential issues sooner and foster a shared sense of accountability.

Of course, trust is only meaningful if it’s measured. It’s important to track metrics such as how often shared data assets are used, mean time to detect and fix incidents, and data retrievability rates. Direct user feedback on confidence in AI systems is becoming part of the equation, too. For example, the icons for thumbs-up/thumbs down for the responses from ChatGPT are a powerful use case of this.

Some companies are taking it further by publishing internal trust dashboards. These dashboards correlate governance activities with user confidence surveys to pinpoint where improvements are needed.

## The Payoff: Why This is Strategic, Not Just Technical

Structured context for AI insights creates a trusted, reliable data layer that organizations can utilize across teams. This common foundation provides consistent KPIs and metrics that help different departments work together more effectively while maintaining regulatory compliance.

What does this actually look like day to day? Teams across the organization—from marketing to finance to operations—can query the same governed metrics and get consistent, trustworthy answers. When everyone’s working from the same definitions and validated data, collaboration improves and compliance becomes automatic rather than an afterthought.

## The C-Suite Takes Notice

This major shift is also reaching the C-suite. Gartner [reports](#) that 70% of Chief Data & Analytics Officers (CDAOs) are now leading enterprise AI strategies. For example, the share of CDAOs reporting directly to CEOs has jumped from 21% last year to 36% this year.

These leaders understand that technology is not intrinsically valuable. It has to drive real business outcomes. Their role is to ensure that data architecture, governance, and trust are closely aligned with the company's broader goals.

## The Architecture for AI-Native Data

What does an AI-native data architecture look like in practice?

### *Unified semantic layer*

A single source of truth for metrics and dimensions that both humans and AI agents can query consistently. This eliminates the confusion that comes from multiple definitions of the same business concept.

### *Embedded governance*

Policies and controls are built into the data layer, not bolted on afterward. This means access controls, data quality checks, and compliance rules are enforced automatically as data flows through the system.

### *Real-time quality monitoring*

Automated systems that detect and often fix issues before they impact downstream users. These systems use ML to learn normal patterns and flag anomalies immediately.

### *Transparent lineage*

Complete visibility into data flow from source to consumption, enabling both debugging and compliance. Every transformation, every join, every aggregation is tracked and auditable.

### *Federated but coordinated*

Different teams maintain their domains while operating within a common framework. This balances autonomy with consistency, letting teams move fast while maintaining enterprise standards.

This isn't just technical architecture—it's organizational transformation. Companies that get this right will have AI systems that are fast, accurate, trustworthy, and compliant.

## The Business Impact

Organizations with well-implemented structured context interfaces and governance see measurable benefits:

### *Faster time to insight*

Natural language queries eliminate the analyst bottleneck

### *Higher data utilization*

When data is discoverable and understandable, it actually gets used

### *Reduced compliance risk*

Automated governance means fewer regulatory violations

### *Improved decision quality*

Trustworthy data leads to confident decisions

### *Accelerated innovation*

Teams spend less time on plumbing and more time on value creation

The companies making this investment today will be the ones whose AI systems actually work when it matters most.

## **Building for What's Next**

Prompt engineering highlighted how unreliable chat-based models can be. You're never sure what they'll come back with. Structured context interfaces shift this dynamic. They allow us to create AI workflows that are stable, scalable, and governed right from the outset.

They're strategic tools that connect business objectives with AI capabilities. As companies move from experimental pilots to systems that sit at the heart of their operations, their success will depend on how well they build and manage these interfaces. Without them, AI remains a risky black box. With them, it can become a trusted partner in making decisions that matter.

The enterprises that win in the AI era won't be those with the best models or the most data. They'll be the ones that build the strongest foundation—where structured context, clear governance, and earned trust come together to enable intelligent automation at scale.

*So what?* The choice is clear. Build the right foundation now, or spend years trying to retrofit governance and trust into systems that were never designed for them. As enterprises step into this new era, success will not come just from picking the right AI tools. It'll come

from rethinking the data stack itself, designing it to support trust, agility, and growth in a world where AI is quickly becoming the norm.