Harness the power of Fabric Ontologies to boost your Agentic AI implementation

Reader Summary

The agent infrastructure is real: multi-agent orchestration, tool calling, retrieval-augmented generation, and production-grade platforms like Azure AI Foundry's Agent Service are all capable and improving rapidly. But most organizations are not seeing the results they expected from enterprise AI, and the problem is not the agents. It is the data.

Agents can query tables and call APIs,
but they do not understand your business:

- what "customer" means in your context,
- how contracts relate to service commitments, or
- what constraints should govern automated decisions.

Ontologies close that gap by making your domain knowledge machine-readable. Microsoft's introduction of Fabric IQ with ontology as its core item signals that this shift is moving from academic to operational. This post explains what ontologies are, why Microsoft is betting on them now, how they connect to the broader agent ecosystem, and why the organizations that start building semantic richness into their data today will have a compounding advantage over those that wait.

The Comprehension Gap

Something fundamental is shifting in enterprise AI, and it is not about the technology getting more powerful. It is about a mismatch between what agents can *do* and what they *actually understand*.

The agentic shift is well underway. The conversation has moved beyond chatbots and copilots to autonomous, goal-directed systems that reason across data, evaluate trade-offs, and take actions on behalf of the business. Cloud-native data platforms like Microsoft Fabric have unified storage, compute, and governance into a single estate. What used to be months of integration work across siloed systems now collapses into days, eliminating the fragmentation that made cross-domain reasoning impractical. AI orchestration has matured to where multi-agent systems can be built, deployed, and monitored in production. Azure AI Foundry's Agent Service offers a production-grade platform with native integrations across Microsoft's ecosystem. The tooling is real, and there is much more to come.

> "The agentic shift gave us the engine. **What we still lack is the map."**

And yet most organizations are not seeing the results they expected. People try an agent, ask a few questions, get answers that feel shallow or inconsistent, and walk away. There is a growing sense of disillusionment: the technology is clearly getting more powerful, but it does not feel like it is getting more useful for the work that actually matters.

The issue is not that agents cannot execute tasks. They can, and they do so increasingly well. The issue is that even the most capable agent, operating against a well-structured data estate, still lacks something fundamental: an understanding of the business it is working within. Agents can query tables, call APIs, and summarize documents. But they do not know what "customer" means in your specific context, how your contracts relate to your service commitments, or what constraints should govern their actions when something goes wrong. Without that understanding, we remain cautious about how much autonomy to grant, and rightly so. "The agentic shift gave us the engine. What we still lack is the map."

There is a pattern we see consistently in enterprise AI engagements. An organization builds an agent, connects it to their data sources, and it works well for straightforward queries. But when the questions become more nuanced, when they require reasoning across domains or evaluating business rules, the agent falters. It gives inconsistent answers. It makes recommendations that violate constraints nobody explicitly encoded. It produces results that are technically correct but operationally wrong. Users lose confidence, stop trying, and the investment stalls.

This is not a retrieval problem. It is a comprehension problem. Agents fail not for lack of data but for lack of meaning. They can access tables and columns, but they cannot reason about the relationships, rules, hierarchies, and domain semantics that define how a business actually operates. The industry has spent two years focused on the mechanics of agentic orchestration: how to call tools, how to structure retrieval, how to manage memory, how to route between agents. These are essential capabilities. But they address the how of agent execution without addressing the what. What does the agent actually understand about the domain it is operating in? Orchestration without comprehension produces agents that can move fast but cannot move wisely.

This gap between data access and domain understanding is the single most important bottleneck in enterprise AI today. And it is exactly the gap that ontologies are designed to close.

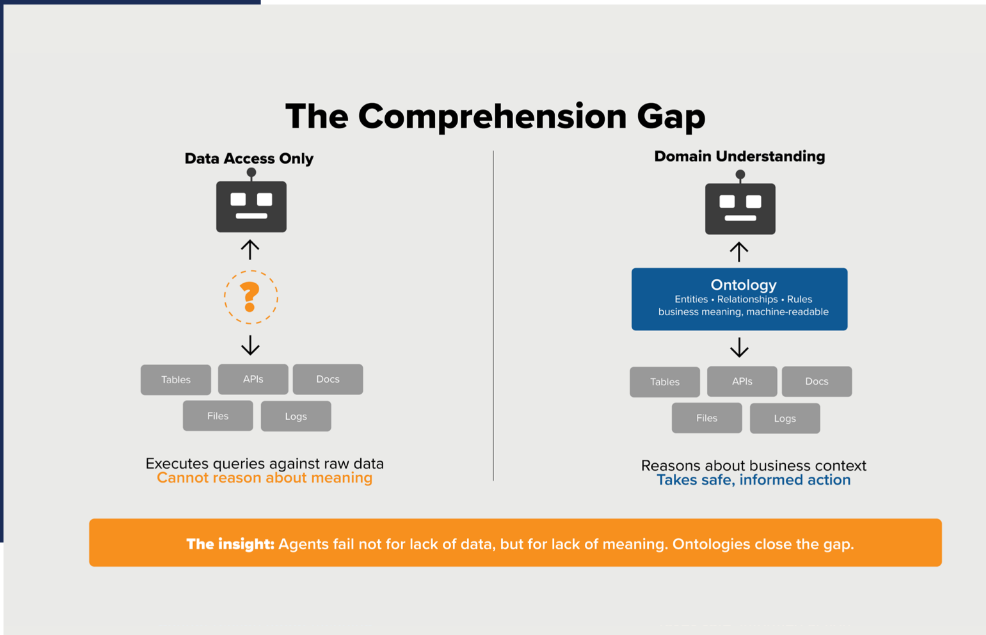

Diagram comparing agents with data access only versus agents with domain understanding through ontologies

Figure 1: *The Comprehension Gap — agents fail not for lack of data, but for lack of meaning.*

Ontologies: Making Your Business Machine-Readable

An ontology, in practical terms, is a machine-understandable vocabulary of your business. It defines the core concepts (entity types like Customer, Shipment, Product, or Sensor) along with their properties, relationships, and the rules and constraints that govern how they interact. It is not a glossary and it is not a data dictionary. It is a structured, queryable representation of how your business actually works, expressed in terms that both humans and AI agents can use for reasoning and action.

The concept of ontologies is not new. Standards like OWL (Web Ontology Language) and RDF have existed for decades in the semantic web community. Domain-specific ontology languages serve specialized industries, from healthcare (SNOMED CT, FHIR) to manufacturing. What is new is that ontologies are moving from academic and niche applications into the operational center of mainstream enterprise data platforms. Microsoft's introduction of Fabric IQ at Ignite 2025, with ontology as its core item in public preview, represents a significant signal. But the broader point is not about any single vendor's implementation. It is about a fundamental shift in what enterprises need to capture and codify in order for AI to work reliably at scale.

Here is why this matters right now: the domain knowledge that ontologies capture is novel information about your organization that AI simply does not know.

No foundation model, no matter how capable, has been trained on how your specific business defines its entities, structures its relationships, or enforces its rules. This knowledge lives in the heads of your people, scattered across documentation, embedded implicitly in your data models, and encoded in business logic that nobody has written down. Until you make it explicit and machine-readable, your agents are operating blind in the one area that matters most: understanding the business they are supposed to serve.

This is the key insight that distinguishes ontology work from other AI infrastructure investments. You are not training a model or fine-tuning a prompt. You are curating knowledge that does not exist anywhere else in a form that AI can consume. It must be thought through, debated, refined, and governed. It requires domain experts and data practitioners working together to articulate what the business actually means by its core concepts. That work is inherently human, and it is the kind of investment that compounds over time as your ontology becomes richer and your agents become more capable because of it.

Microsoft's Bet:
The Semantic Model as Ontology Seed

Microsoft's pitch to existing Fabric and Power BI customers is pragmatic: you already have a lightweight ontology in your semantic model. Power BI semantic models define entities, relationships, measures, and hierarchies that represent how your business thinks about its data. Fabric IQ's ontology can be generated directly from an existing semantic model, with entity types created from tables, properties from columns, and relationship types from the model's defined relationships. As Arun Ulag, Microsoft's President of Azure Data, put it at Ignite: "We've had a gap. These semantic models were only used for BI use cases. There's a much bigger opportunity out there, which is the opportunity to take these semantic models and upgrade them into a full ontology." That upgrade means extending beyond analytics into operational context: connecting data across the enterprise, integrating real-time streams, and defining business rules and constraints that agents can reason over.

Microsoft has been expanding the ways that semantic definitions can be expressed within the Fabric ecosystem. The introduction of dbt jobs in Fabric brings a code-driven approach to defining models, with richer metadata, documentation, and quality checks built into the development workflow. For organizations already practicing analytics engineering, this is strategically significant: every well-documented dbt model, every column-level description, every tested business rule becomes raw material for the ontology your agents will eventually need. Treating your transformation layer as a semantic investment, not just an engineering convenience, is one of the highest-leverage decisions a data leader can make right now. We will explore how to reverse-engineer existing Power BI projects into enterprise-grade dbt in a companion technical post.

The key point for leaders is this: **the domain knowledge your ontology will capture is perishable, and it is decaying right now.**

The senior analyst who knows why "active customer" excludes trial accounts in one division but includes them in another is not going to document that before she retires. The data engineer who built the original ETL and understands which upstream quirks the logic silently compensates for is interviewing at another company. Every month you defer this work, the knowledge gets harder to capture and the ontology you eventually need to build gets more expensive. Beginning the work of making your business machine-readable, starting with whatever semantic layer you already have, is more important than waiting for the perfect standard to emerge. The cost of starting imperfectly is low. The cost of starting late is compounding.

The Full Stack: Data, Context, and Action

What is emerging across the Microsoft ecosystem is a three-layer architecture for enterprise AI. Understanding it matters not because the architecture is novel, but because it reveals exactly where the gap is for most organizations, and why that gap is urgent to close.

The **data** foundation is Microsoft Fabric's OneLake: lakehouses, warehouses, eventhouses, and mirrored databases all accessible through a unified platform with consistent security and governance. This is where your data lives.

The **context** layer is the semantic intelligence that sits above the data. In Fabric, this is the ontology from Fabric IQ, which transforms raw data into business-meaningful context by defining entities, relationships, constraints, and rules that agents can reason over. Combined with Foundry IQ (which reimagines RAG as a dynamic, multi-source reasoning process) and Work IQ from Microsoft 365 Copilot, Microsoft is assembling what it calls a "shared intelligence layer" for the enterprise. The architecture ensures that agents operating anywhere in the stack reason within the same semantic guardrails.

The **action** layer is the agent runtime. Azure AI Foundry's Agent Service provides the platform for building, deploying, and governing agents with native integrations across the ecosystem. Fabric's Data Agents use the ontology to answer business-level questions across heterogeneous data sources. Operations Agents go further, continuously monitoring real-time data, reasoning over live conditions, and taking automated actions bounded by ontology-defined rules. Because the constraints and business rules live in the ontology itself, these agents can move beyond generating answers to taking safe, auditable actions.

The architectural significance is that these layers are designed to compose. An agent built in Foundry can use a Fabric Data Agent as a knowledge source, which in turn uses the ontology to resolve queries across semantic models, lakehouses, and eventhouses. The ontology's constraints travel with the data. This composability is what transforms a collection of AI capabilities into a genuinely agentic platform.

There is a strategic consideration that deserves candor. Industry analysts have noted that the deeper an enterprise builds on this semantic layer, the greater the investment in the Fabric ecosystem becomes. Ontologies create valuable organizational assets, but they also create switching costs. Leaders should pursue this path with clear-eyed awareness of both the value it unlocks and the platform commitment it implies. Building on open standards where possible and maintaining clear documentation of your ontology definitions outside of any single platform are practical mitigations. But the more pressing risk for most organizations is not over-commitment to a platform. It is under-investment in the knowledge layer while the people who hold that knowledge are still available to contribute to it.

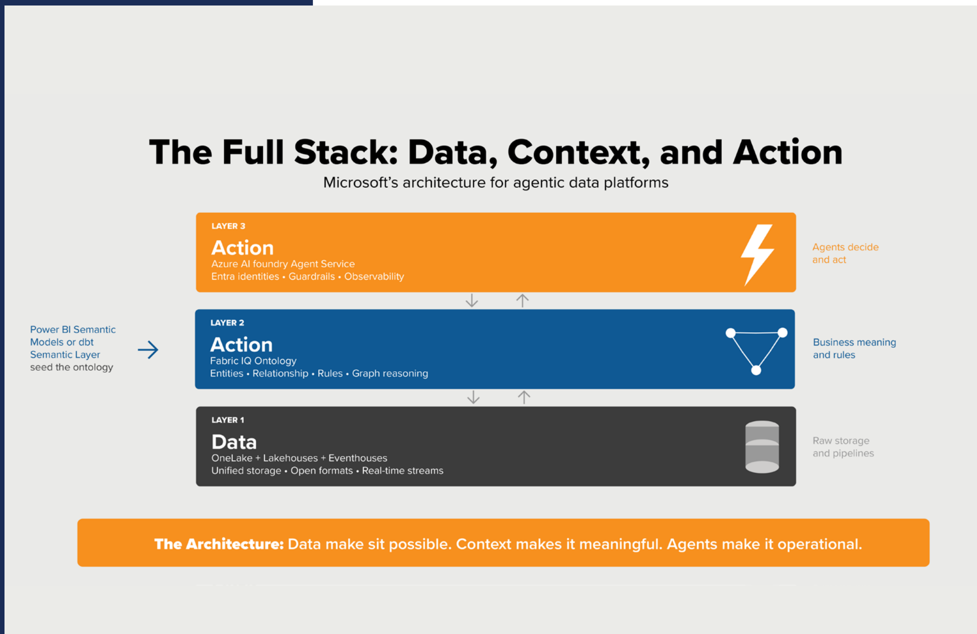

Three-layer architecture diagram showing Data Foundation, Context Layer, and Action Layer

Figure 2: *The Full Stack — data makes it possible, context makes it meaningful, agents make it operational.*

The Maturity Path: Practical Steps to Get There

This is a maturity journey, not a single deployment. Based on what we are observing across enterprise engagements, we recommend thinking about this in three stages. Each stage produces tangible artifacts that deliver value immediately while building toward the next.

1 Invest in your semantic layer with intention.

If you have Power BI semantic models, treat them as the foundation they are, but invest in making them richer, more consistent, and more deliberately designed. In practice, this means starting with an entity definition inventory: a structured audit of how core business concepts like “customer,” “revenue,” and “active” are defined across your semantic models, and where those definitions conflict. It means running a measure reconciliation process where the same metric computed differently across departments gets resolved to a single governed definition. And it means conducting documentation sprints where data engineers and business stakeholders sit together to capture the implicit rules that everyone “just knows.”

Power BI semantic models are one approach, but they have limitations, particularly around version control, testing, and cross-team governance. The introduction of dbt in Fabric provides additional opportunities to express your semantic layer through code with richer metadata and quality checks built in. The deliverable at the end of Stage 1 is a documented, reconciled semantic layer with explicit definitions for your core entities and measures, regardless of which tooling produced it.

2 Model your business domains explicitly.

Begin piloting ontology capabilities in a bounded domain where cross-domain reasoning or real-time operational context would deliver clear value. Supply chain, claims processing, and customer operations are common starting points because they involve multiple entity types with well-understood relationships and clear business rules. The concrete work here is producing a domain model that defines your entity types, their properties, their relationships, and the constraints that govern them. For a claims processing domain, that might mean 20-30 entity types (Claim, Policy, Provider, Member, Procedure, Authorization) with explicitly defined relationships and business rules like "a claim cannot be approved if the authorization has expired" or "out-of-network providers require additional review above a threshold amount." Whether you capture this in Fabric IQ's ontology modeler, in OWL, or in a structured format like JSON-LD, the act of formalizing it is what matters. Generate the initial ontology from your semantic model where you can, then extend it with constraints, streaming data integrations, and the cross-domain relationships that analytics alone never required. The deliverable is a formal ontology for one business domain, validated by the domain experts who will rely on it.

3 Pilot agent patterns against well-described data.

Start with Data Agents for conversational analytics, testing whether the ontology produces more consistent and accurate answers than direct database querying. The evaluation should be concrete: take 50 business questions that your analysts answer regularly, run them against an agent with and without the ontology, and measure accuracy, consistency, and the rate at which the agent produces answers that are technically correct but operationally wrong. Then progress to Operations Agents that can monitor conditions and recommend or take actions. This is where ontology-defined constraints become critical, because actions carry consequences and the agent needs to reason about what is permissible, not just what is possible. Microsoft now gives every agent a first-class Entra identity with full lifecycle management, access policies, and role-based controls, meaning agents are subject to the same identity governance as your human workforce. That identity layer, combined with AI Foundry's built-in guardrails, observability dashboards, and red-teaming capabilities, provides the operational governance envelope. The ontology provides the semantic governance, ensuring that the agent's reasoning is bounded by your actual business rules and constraints. At each stage, the ontology acts as both an accelerator and a guardrail, improving what the agent can do while bounding what it should do.

WHERE THIS IS GOING: THE ENTERPRISE IN AN AGENTIC WORLD

Where This Is Going: The Enterprise in an Agentic World

The organizations that invest in this foundation over the next two to three years will find themselves operating in fundamentally different ways. We see the precursors of this in current engagements. A financial planning team deploys an agent against their data warehouse and it answers surface-level questions well: total revenue by region, headcount by department. But when a VP asks **"which facilities are underperforming relative to their peer group?"**, the agent cannot answer, because "peer group" is a business concept defined by bed count, payer mix, and geography that exists only in the CFO's head and a spreadsheet on someone's desktop. The data is accessible. The meaning is not. That gap between what agents can reach and what they can understand is exactly what ontologies close, and it is the gap that determines whether your AI investment produces insight or just data retrieval.

Healthcare illustrates the pattern concretely. Today, a health system might use AI to summarize patient records or flag abnormal lab results. Useful, but limited. With an ontology that encodes care pathways, formulary rules, payer constraints, and regulatory requirements, agents can coordinate care transitions across departments, verify that discharge plans satisfy both clinical and insurance requirements, or monitor post-surgical recovery against expected trajectories and escalate proactively when patterns diverge. Healthcare already has rich domain vocabularies in SNOMED CT, FHIR, and ICD-10. The opportunity is to connect those standards to agent-consumable ontologies that encode not just what things are, but how they interact and what constraints govern them.

The pattern scales across domains. A supply chain agent grounded in an ontology that encodes supplier tiers, contractual obligations, and lead times moves from reactive alerting to proactive optimization with full traceability of its reasoning. A compliance agent with access to a regulatory ontology can reason about whether observed patterns are consistent with frameworks, not just match keywords. In every case, the differentiator is the same: the agent's value scales directly with the richness of the ontology behind it.

None of this requires waiting for some future generation of AI models. The foundation models available today are capable enough. The orchestration patterns are capable enough. What is missing, and what only the enterprise itself can provide, is the structured domain knowledge that makes these scenarios safe and reliable.

What This Means for Technology Leaders

The fundamental principle is straightforward. The better context you can provide your agents, the better they work. An agent grounded in rich domain context can do things that an agent with only data access cannot, and that gap widens as the scenarios become more consequential.

This is the same trajectory we saw with business intelligence. Semantic models started as a way to make data more accessible for reporting. Over time, they became the foundation for self-service analytics, enabling business users to ask and answer their own questions without waiting for IT. Ontologies represent the next step in that progression, but the destination is different. Where semantic models enabled self-service insight, ontologies will enable autonomous action. They are the bridge that allows us to grant agents the autonomy to reason, decide, and act on behalf of the business, because the guardrails and context they need are built into the layer itself.

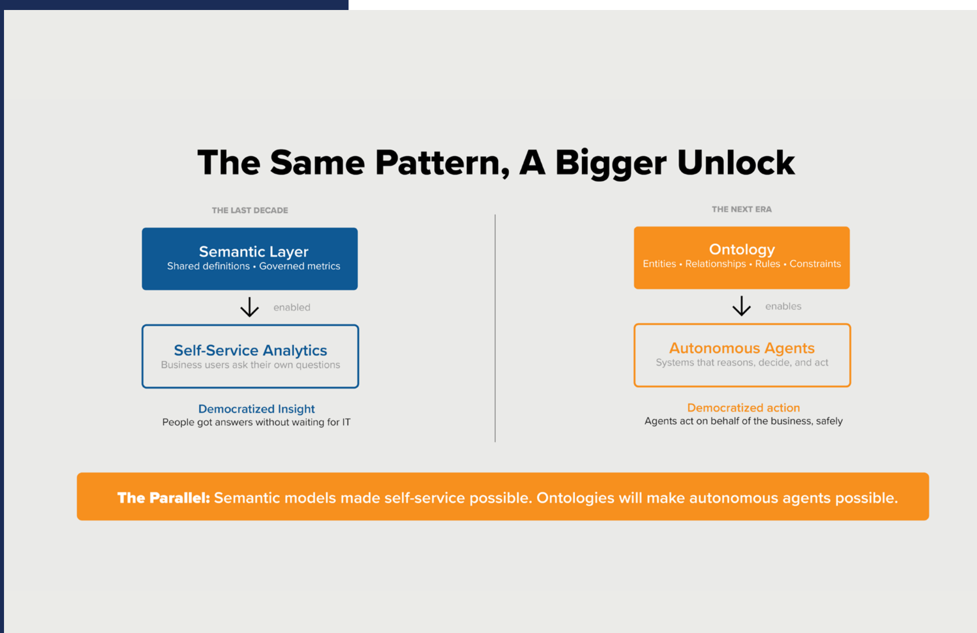

Parallel comparison showing semantic layer enabled self-service analytics while ontologies enable autonomous agents

Figure 3:
Semantic models made self-service possible. Ontologies will make autonomous agents possible.

The work of capturing that context, standardizing definitions, formalizing relationships, and documenting the rules that govern how your business operates, is foundational. It compounds over time. The opportunity is to treat your domain knowledge as the strategic asset it is, and to start giving your agents the context they need to do meaningful work.

This is also an operational question. The ontology you build is not a trade secret. It is a metamodel that enhances data discoverability across the organization, improving accuracy and accelerating integration across departments that previously operated in silos. Two organizations can deploy the same foundation models, use the same orchestration framework, and run on the same cloud infrastructure. The one with a well-governed ontology will produce agents that are faster to deploy and more accurate in their reasoning, because the agents can discover and ground themselves in business context that would otherwise require human interpretation. In many industries, ontologies are already published externally as interoperability interfaces, much like APIs for B2B systems. The advancement is not in locking knowledge away. It is in making knowledge structurally discoverable, both within the organization and across the boundaries where systems need to interoperate.

Tech Takeaways

1

Ontologies close the gap between data access and domain understanding. Agents fail not for lack of data but for lack of meaning. Ontologies make your business rules, entity relationships, and constraints machine-readable so agents can reason, not just retrieve.

2

Your semantic model is the starting point, not the finish line. Fabric IQ can generate an initial ontology from existing Power BI semantic models. Extend it with constraints, real-time data, and cross-domain relationships to move from analytics to operational intelligence.

3

Domain knowledge is perishable—start capturing it now. The people who know why "active customer" means different things in different divisions will not be around forever. Every month you defer, the ontology you eventually need gets more expensive to build.

4

The three-layer architecture—data, context, action—is the blueprint. Fabric's OneLake provides the data foundation, Fabric IQ's ontology provides the context layer, and Azure AI Foundry's Agent Service provides the action layer. Plan your investments across all three.

5

Your ontology accelerates everything else. Investing in making your business knowledge structurally discoverable means agents deploy faster, departments integrate more easily, and AI can ground itself in context that raw data alone cannot provide.

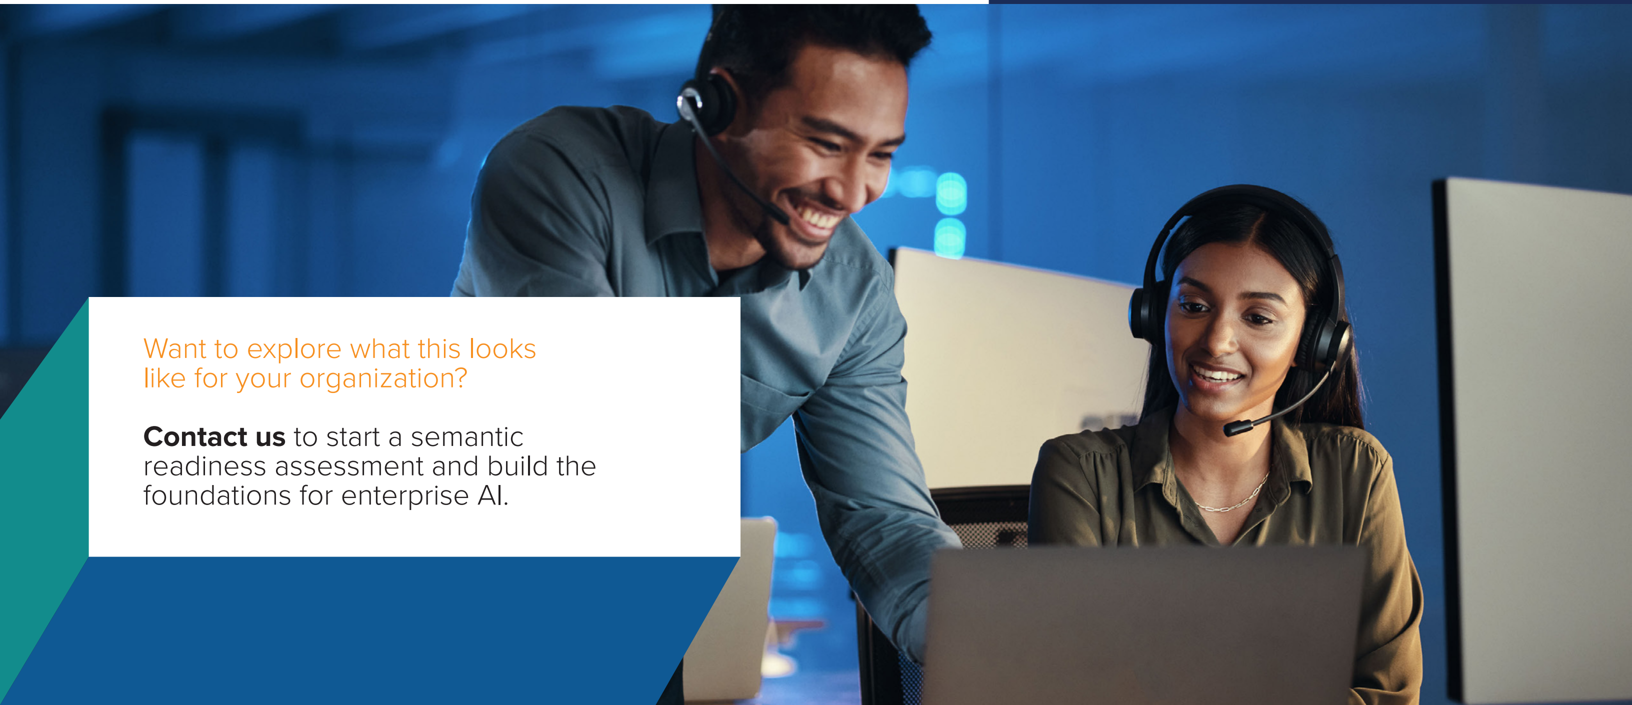

Want to explore what this looks like for your organization?

Contact us to start a semantic readiness assessment and build the foundations for enterprise AI.