



The data gateway for AI agents.

Connect AI agents to operational systems and warehouses, without giving up performance, governance, and trust.

Ai4 | 2026 · AUGUST 4-6 · LAS VEGAS

MARCOPOLO.DEV

EXPO HALL · LAS VEGAS

BOOTH

841

ABSTRACT

Enterprise AI has crossed a definitive line. Agents no longer produce insights for a person to act on. Instead they can now act autonomously, and must therefore reach live enterprise systems at runtime. These include operational systems such as the CRM, ERP, ticketing system, and application database, alongside the warehouse and its semantic layer.

Open standards like the Model Context Protocol have made the connection trivial. However, making it fast, governed, and consistent across many systems, for agents that act autonomously at scale, is not. Meeting that challenge requires a distinct piece of infrastructure we call a data gateway.

Our aim here is to define what it is, why the warehouse era's playbook of centralization no longer applies, the three responsibilities the gateway must carry, and how MarcoPolo is the first at scale approach to implement a data warehouse.

From insights for humans to actions by agents

For two decades, business intelligence turned data into knowledge so that humans could close the loop. Data became dashboards, dashboards informed decisions, and people translated those decisions into action. We built data warehouses because people needed a consistent view of the business: we centralized data, standardized metrics, and built semantic layers so that everyone looking at revenue, churn, or pipeline saw the same answer. The warehouse became the company's shared source of truth that made analytics faster and consistent.

AI agents flip that model. Instead of producing insights for someone else to execute, they can execute on the insights that they find. Every action carries just enough reasoning in the moment needed to take an action. We call it "small analytics" over a handful of records: relate these accounts, reconcile this order against the ticket blocking it, read the contract before answering the customer. Intelligence no longer lives in a separate analytical layer; it lives where the work happens. AI inherits the warehouse era's requirement for consistency, but the warehouse alone is no longer enough.

This shift in how work gets done requires data and ops teams, as well as AI developers to think about the architecture of their AI stack differently from that of the traditional BI stack. All of the work that they have already put into building a data warehouse, semantic layers and data standardization now becomes a source for the AI agents to draw upon.




Figure 1. BI centralized data into a warehouse; AI federates access through a gateway. The warehouse standardized answers; the gateway standardizes trust.

Operational systems are the center of gravity

However, an agent rarely completes real work from warehouse data alone. It also needs the hot context, the live information that lives in CRM, ticketing system, ERP apps, the application database, Slack, emails, recent events, documents, and conversations. In essence, any system that currently owns the truth for the task at hand. Operational systems hold things a warehouse intentionally abstracts away:

- **Business logic.** The workflow that determines whether an opportunity is actually qualified lives inside Salesforce automation, not in a flattened status column.
- **Permissions.** The source system knows who can see or modify a record. Downstream copies approximate those rules, but they are rarely identical.
- **Operational reality.** The latest transaction, the field someone added last week, the event that arrived thirty seconds ago often exists only in the source systems.

Warehouses were designed to transform data for analysis; agents often need the systems before that transformation. The warehouse remains essential — its semantic layer is still where standardized metrics are made consistent — but in the agent era it becomes one authoritative connection among many, not the single destination. This mirrors the shift the data mesh movement described for analytics: domain-oriented, decentralized ownership with federated governance scales better than forcing everything into one central platform (Dehghani, 2019). The gateway extends that logic from the analytical plane to AI's operational runtime.

This is already happening in the field. Anthropic introduced the Model Context Protocol (MCP) in November 2024 as an open standard for connecting AI to external systems; one year later it reported more than 10,000 active public MCP servers and over 97 million monthly SDK downloads. The workforce is already wiring AI into live operational systems at runtime — with or without an AI strategy.

Connecting is easy; performance, governance, and trust are not

Point an LLM directly at a CRM and it burns tokens exploring schemas, chooses the wrong fields, misses business rules, or retrieves records it should never see. This is not a temporary limitation. On BIRD, a benchmark built from real enterprise databases, GPT-4 reached only about 55% execution accuracy where humans reached roughly 93% (Li et al., 2023) — evidence that pointing a capable model at raw schemas is not enough to produce correct, trustworthy answers.

GPT-4 EXECUTION ACCURACY
· BIRD

~55%

HUMAN BASELINE · SAME
BENCHMARK

~93%

These are the same problems data and engineering teams already solved for the warehouse: performance, governance, and trust. The work still has to be fast and affordable; access still has to be limited to the right identity for both the agent and the human receiving the analytics; the answers still have to be correct and consistent.

Governance is now harder, not easier, because there may not be a person in the loop. LLMs process instructions and data in the same channel, so any credential or sensitive record placed in a prompt becomes an exposure — prompt injection and sensitive-information disclosure sit at the very top of the OWASP Top 10 for LLM Applications (2025). Meeting these requirements calls for least-privilege, per-session access in the spirit of zero-trust architecture (NIST SP 800-207).

This is where the old playbook must evolve. Building BI meant centralizing data into one warehouse. Building AI agents means preserving the authority of distributed systems while still meeting those three requirements across all of them at once. None of the hard problems that were solved over decades for data warehouses disappeared when AI arrived — they just moved closer to the operational systems where agents now work.

"We want to own context: a controlled, scalable way to give our agents access to real data without locking into a single LLM vendor."

ADAM ANDERSON · VP REVENUE OPERATIONS, CONSERVICE

Defining the data gateway

A data gateway has to perform those same responsibilities that we defined for a data warehouse, except it must do so across many systems instead of just one:

01 Federated access

Instead of moving every dataset into one platform, the gateway reaches into the system where the authoritative data already lives, and queries it in place. The warehouse remains an important source for standardized analytics, but it becomes one connection among many.

02 Governance without a human in the loop

Agents can act autonomously or on behalf of their users. The gateway has to enforce policy consistently across every system it fronts, apply least privilege per identity and session, and prevent credentials and sensitive information from ever entering prompts or responses.

03 Context, the hardest problem

BI standardized answers to questions people knew they would ask, by building semantic layers. AI has to answer questions that have not been modeled ahead of time. The agents discover at runtime which systems matter, which records belong together, which business rules apply, and which procedures to follow. Two people asking different questions about the same customer should still receive answers that agree, even though neither path was designed in advance. That is a harder standardization problem than BI ever faced.

The data gateway becomes the checkpoint where this happens. Every agent finds its route through enterprise systems at runtime, in response to a user's goal — and still has to be performant, governed, and trustworthy on every path it takes.

How MarcoPolo implements the data gateway

MarcoPolo is a data gateway for AI agents whereas MCP is the connection. MarcoPolo provides a workspace where AI agents can work on live data. Any AI surface, whether you use Claude, ChatGPT, Cursor, Copilot, Codex, or a custom agent, connects over MCP, and each user or agent runs in its own dedicated workspace hosted in your cloud, or ours.




Figure 2. The MarcoPolo data gateway.

MarcoPolo enables the three responsibilities as follows:

- **Federated access.** An in-pod connection CLI reaches into each source system and queries it in place; results are materialized into a persistent DuckDB workspace where they can be joined across systems. The warehouse is simply one connector among many connectors that span storage (S3, GCS), databases (Postgres), warehouses (Snowflake), and apps (Salesforce, Jira).
- **Governance.** Credentials are scoped per user and session, encrypted, and never enter the model's context window. Tenant isolation is enforced through per-session containers, with SIEM audit over every action.
- **Context.** The workspace holds schema and context, including syntax, docs, and examples, alongside query history, so an agent can discover the right systems and records at runtime. Only summaries and previews return to the model's context window, keeping reasoning focused and answers consistent across paths.
- **Performance.** Caching results in DuckDB and returning summaries rather than raw rows keeps token cost low and follow-up analysis fast.

A single request shows the pattern end to end. Asked to "show which customers churned last quarter and pull their support tickets," an agent lists available connections, describes the Salesforce schema, queries churned accounts, queries the related Jira tickets, pulls their revenue data from Snowflake data warehouse, and joins across all three through DuckDB in MarcoPolo, returning a sub-kilobyte summary to the model while the full multi-megabyte result stays cached in the workspace. Two operational systems and a warehouse, federated and governed, in a single agentic conversation.

"A Salesforce flow diagnosis that used to take six to seven hours was completed in about ten minutes in a sandbox test, a massive gain for the efficiency of the team."

ADAM ANDERSON, CONSERVICE

"We use AI to compose portable, shippable data manifests, and a cloud-agnostic token-exchange layer to keep authorization standardized across systems — it saves us design and dev hours."

MATHEUS BAFFUTO, DUPLOCLOUD

Where this is going

Enterprise AI inherits the requirements of the warehouse era — performance, governance, and trust — but it can no longer meet them by only centralizing data into a warehouse. It has to meet them at the edge, across operational systems and warehouses where the truth actually lives.

The data gateway is the runtime checkpoint that makes that possible: it is where an agent's route through enterprise systems is discovered, governed, and made consistent, every time. Integrating a data gateway into your AI strategy now will simplify how the next generation of AI agents and applications gets built.

We understand that your data warehouse took years to build, that your operational systems provide the hot context that AI agents need to act autonomously, and that a data gateway enables you to take advantage of both worlds. Working with developers, engineering leaders, and AI enablers across IT, engineering, and data teams who are connecting agentic AI to enterprise systems, we designed MarcoPolo to enable this architectural shift for AI agents.

Try MarcoPolo, or meet us at Ai4 2026 — Booth 841

August 4–6 · The Venetian, Las Vegas · contact us for a complementary consultation

MARCOPOLO.DEV

REFERENCES

1. Anthropic. "Introducing the Model Context Protocol." November 25, 2024.
2. Anthropic. "Donating the Model Context Protocol and establishing the Agentic AI Foundation." December 9, 2025.
3. Dehghani, Z. "How to Move Beyond a Monolithic Data Lake to a Distributed Data Mesh." martinowler.com, 2019.
4. Li, J., et al. "Can LLM Already Serve as a Database Interface? (BIRD)." NeurIPS 2023.
5. OWASP. "Top 10 for LLM Applications." 2025.
6. NIST. "Special Publication 800-207: Zero Trust Architecture." 2020.