

Your Data Has a New Consumer

Why AI Agents Are Breaking Your Infrastructure



[Insert Name]

[insert title]


FIVETRAN + dbt Labs

Your Data Has a New Consumer: AI Agents

Most enterprise data infrastructure was built for human analysts working at human speed. Agents are neither. They operate continuously and act on whatever data they find, at scale, automatically, before anyone notices a mistake. The infrastructure that served your analytics teams well was simply never designed for that. In this session, we'll cover:

- The four things AI agents actually need from your data (fresh, complete, trusted, and open)
- The hidden costs of getting it wrong
- How leading organizations like NAB and LVMH are building their data foundation for AI with Fivetran + dbt Labs





40%

5%

of enterprise
applications will feature
AI agents by 2026.

Up from less than 5% last year.

**The age of AI agents is not
coming. It is already here.**




Agentic AI is already changing what's possible



Every industry, every scale, everyone's innovating



Autonomous Vehicles



AI Drug Discovery




AI Customer Service Agents



AI Drug Intelligent Supply Chains



**For years, we built
infrastructure for
one consumer**



**Today, the primary
consumer of
enterprise data are
AI agents**

**Data's role
has changed;
infrastructure
*has not.***

COST INCREASES

OPERATIONAL FAILURES

GOVERNANCE RISKS

INABILITY TO SCALE



How much
enterprise agentic
AI deployments
cost to run monthly
\$13K



1

Costs explode at scale

THE SCENARIO


You deploy an internal IT helpdesk agent for 500 employees.

WHAT THIS MEANS FOR YOUR TEAM

- ❖ Audit your data pipeline costs now — before agents multiply them
- ❖ Fragmented pipelines mean 10× the retrieval overhead
- ❖ Open, reusable data reduces redundant compute

The age of data in
a typical overnight
batch when an
agent acts on it
at noon

14hrs



2

Stale data creates serious consequences

THE SCENARIO


Your e-commerce agent is running on overnight batch data.

WHAT THIS MEANS FOR YOUR TEAM

- ❖ Batch pipelines are not fit-for-purpose for agent workloads
- ❖ Real-time CDC (change data capture) becomes a baseline requirement
- ❖ Every hour of stale data = an hour of wrong agent decisions

Data policy
violations involving
AI applications, per
month

223



3

Governance risks increase

THE SCENARIO


An AI financial reporting agent pulls from an ungoverned data source.

WHAT THIS MEANS FOR YOUR TEAM

- ❖ Data lineage and auditability are now legal requirements, not nice-to-haves
- ❖ Ungoverned pipelines create AI liability, not just data quality issues
- ❖ Governance must be built into the pipeline — not bolted on after

320x

YoY increase in AI reasoning token consumption per enterprise org



4

Low latency becomes a necessity

THE SCENARIO

Your company rolls out an internal research and analysis agent.

WHAT THIS MEANS FOR YOUR TEAM

- ❖ Provision data infrastructure for agent-scale, not human-scale
- ❖ Multi-engine, open architecture avoids single-warehouse bottlenecks
- ❖ Always-on pipelines prevent latency spikes under agent load

You don't have an AI problem.

You have a *data* problem.

The infrastructure built for dashboards and reports was never designed for what AI agents actually need — fresh, complete, governed, open data at scale.

What AI agents actually need from your data

Agents don't just analyze dashboards. They take action.

1

Fresh

Agents operate continuously and act on whatever they find. Batch pipelines mean agents make decisions on yesterday's reality and execute those decisions automatically before anyone notices.

Human: daily report → Agent: 1-minute sync

2

Complete

A human can infer missing context. An agent cannot. AI needs a unified view across every business system — CRM, ERP, support, finance — to act reliably and not create blind-spot errors.

Human: single-source view → Agent: 360° context

3

Trusted

When an agent acts on ungoverned data, consequences compound. Every agent action must be traceable, auditable, and built on a governed foundation that can defend decisions to regulators.

Human: mistake is caught → Agent: mistake is systemic failure

4


Open

AI tooling evolves fast. Proprietary formats create expensive rebuilds. Agents need data in open formats readable by any engine so the infrastructure you build today doesn't become the technical debt of tomorrow.

Human: BI tool output → Agent: open, portable data layer



Move, manage, & transform data with Fivetran + dbt Labs



TRANSFORM


- dbt Labs coordinates transformations
- Quickstart data models speed analytics
- Operational visibility into transformation runs

MANAGE

- Supports leading cloud data lakes (Amazon S3, Google Cloud Storage, Azure Data Lake Storage)
- Managed ingestion into the data lake
- Open table formats to convert data into standardized formats (Iceberg, Delta Lake)
- Unified catalog integrations

MOVE

- 700+ pre-built connectors
- High-performance, low-impact CDC
- 99.9% uptime across 5.8M+ daily syncs
- Automated schema drift and API change
- Idempotent, self-healing pipelines



AGENTS SCHEMA an open-source standard giving agents a standard place to look up context

ANY COMPUTE



PLATFORM FOUNDATION



BUILD & EXTEND
with SDKs, APIs, and 3rd party integrations



SECURITY & PRIVACY
with full compliance, encryption, and multiple deployment offerings



RELIABLE & AUTOMATED
with 99.9% uptime across 7M+ daily syncs


Companies that got the foundation right

They invested in data infrastructure first and unlocked AI at scale.



Australia's largest business bank


Built a lakehouse ingesting 200+ sources via CDC, cutting costs 50% in year one while deploying GenAI workloads that now run at enterprise scale.




LVMH

Global luxury with 70+ Maisons

Unified data across Louis Vuitton, Dior, Moët Hennessy, achieving a 360-degree customer view. Unlocked real-time SAP and marketing insights to accelerate decision-making.





KEY TAKEAWAYS



Your infrastructure was built for human analysts; it's not ready for the speed, scale, and autonomy of agents.


Success requires more than just a model. You need data that is fresh, complete, trusted, and open.

Build a new data foundation to turn AI investment into business impact.

The question isn't
whether AI agents will
run your business.

It's whether your data
will be ready when
they do.





Thank you!