

ARCHITECTURE BRIEF

The Four Levels of Agentic AI Architecture

Why the industry's current path cannot reach the destination, and what the agentic enterprise actually requires



— ARCHITECTURE BRIEF

A decade-old lesson worth remembering

In 2012, the consensus playbook for moving to the cloud was "lift and shift." Take your virtual machines, move them to AWS, declare victory. A decade later, the companies that won the cloud were not the ones who got there first. They were the ones who eventually realized that lift-and-shift was the starting line, not the destination. The real winners re-architected for what cloud actually offered.

The agentic AI era is now in its own lift-and-shift moment. The majority of enterprise AI initiatives in 2026 are smart chat interfaces bolted onto unchanged data architectures; agents as queries, not actors. They demo well. They do not scale to the agentic enterprise.

THE ARGUMENT OF THIS PAPER

Four distinct architectural levels define how vendors organize data, compute, and control for agentic workloads. The boundaries between them are structural, not gradual. Vendors cannot incrementally evolve between levels without abandoning the commercial premise of their current one. Where your AI strategy sits on this taxonomy today will largely determine whether it can serve the agentic enterprise tomorrow.

◆ WHAT YOU WILL FIND IN THE NEXT FIVE PAGES

- A four-level taxonomy classifying the dominant approaches to agentic AI architecture in 2026.
- Each level's architectural premise, technical strengths, and the workloads it cannot serve.
- A structural argument for why Level 1 and Level 2 vendors cannot reach Level 4 without rewriting their commercial models.
- The announcement of Acceldata's Autonomous Data & AI Platform — the first Level 4 architecture in the market, built on a Level 3 foundation laid since 2021.


WHO THIS PAPER IS FOR

CIOs, CDOs, chief architects, and heads of data and AI evaluating where to commit architectural investment for the next three to five years. If you are currently evaluating Databricks, Snowflake, Google's Agentic Data Cloud, or building agentic capabilities on top of your existing data estate, this paper is designed to help you locate that decision on the architectural map.

THE FOUR LEVELS AT A GLANCE

Data · Compute · Control

Every agentic AI architecture organizes three core elements: data (where information lives), compute (where workloads execute), and control (how governance, observability, and orchestration are applied). The four levels are distinguished by where each of these elements sits relative to the vendor boundary, and by how much intelligence the control plane itself carries.



Read the progression left to right

What moves outside the vendor box as you progress is data first (Level 2), then compute (Level 3), then governance itself (Level 4). What moves inside the control plane at the same time is intelligence: from passive configuration in Levels 1–3 to active, autonomous operation in Level 4.

THE ARCHITECTURAL PREMISE OF EACH LEVEL

LEVEL 1 Single Plane

The vendor's premise is that consolidation wins. Bring all workloads to one platform, governed in one place, optimized in one runtime. Easy to operate, hard to escape.

LEVEL 2 Split Plane

The vendor's premise is that data won't consolidate but agents will. Let data live where it sits; run the agents in our cloud. Half federated, half consolidated.

LEVEL 3 Federated Plane

The architecture's premise is that neither data nor compute will consolidate, and that the customer must control the orchestration layer that spans them. Federation is the architecture, not a feature.

LEVEL 4 Autonomous Plane

The control plane itself becomes agentic. Observability that diagnoses, governance that adapts, optimization that re-replaces workloads — all running across the federated estate without human direction.

THE CONSOLIDATION ERA: WHERE THE MARKET IS TODAY

Today's incumbents


Levels 1 and 2 represent today's incumbents. Both architectures are commercially successful, technically mature, and increasingly inadequate for cross-cloud, hybrid, and sovereignty-sensitive agentic workloads.

LEVEL 1

Single Plane

Databricks · Snowflake

ONE VENDOR · ONE CSP BOUNDARY



Data, compute, and control consolidated in one vendor's perimeter, bounded by one CSP. Agent-native primitives built into the lakehouse or warehouse.

TECHNICAL PROS

- Real transactional capability for agents (Postgres-on-lakehouse).
- Branching and scale-to-zero make ephemeral agent environments cheap.
- Single governance model in-platform; strong economics for in-platform workloads.
- Lowest-latency path for native data.

TECHNICAL CONS


- Vendor-managed control plane; customer does not own the orchestration layer (sovereignty concern).
- Single-CSP scope per workload, cannot span clouds; no on-prem substrate.
- Lock-in is multiplicative across data, compute, governance, and ML.
- Single platform = single blast radius.
- Performance ceiling on extreme-TPS workloads (Postgres semantics).

LEVEL 2

Split Plane

Google Agentic Data Cloud

DATA OUTSIDE, CONTROL INSIDE



Data plane federated across clouds via Iceberg; agent runtime and control plane consolidated in one cloud. Internalizes the data-gravity lesson, not the compute-gravity lesson.

TECHNICAL PROS

- Best-in-class models and silicon (TPU, Gemini 3.x).
- No forced data movement; Iceberg cross-cloud federation is real.
- Strong agent governance primitives (Identity, Registry, Gateway, Observability).
- A2A protocol for cross-platform agent interoperability.

TECHNICAL CONS

- Agent runtime consolidated to GCP. Serious agent inference economics tied to TPU.
- Cross-cloud query latency is real (egress + RTT) for hot paths.
- Knowledge Catalog assumes Google as the semantic-layer owner.
- Workspace pull is a one-way ratchet for productivity workloads.

WHY LEVELS 1 AND 2 ARE UNSTABLE

Level 1 cannot serve a multi-cloud, sovereignty-bound enterprise. Level 2 admits this for data, then contradicts itself by consolidating the agents.

THE FEDERATION ERA: WHERE THE ARCHITECTURE IS GOING


Foundation and destination

Levels 3 and 4 start from a different premise: enterprise data and compute will not consolidate, and the control plane is the only layer that can and must span them. Level 3 is the foundation Acceldata has been building since 2021. Level 4 is the destination we are announcing.

LEVEL 3

Federated Plane

Acceldata Foundation · 2021–2026



Federated compute runs where data lives across clouds and on-premises. Customer-controlled unified control plane spans the heterogeneous estate, including on-premise.

TECHNICAL PROS

- Matches the empirical reality that enterprise data and compute do not consolidate.
- Resilient to single-vendor risk and pricing concentration.
- Aligned to data gravity, residency, and sovereignty requirements.
- Agentic workloads do not trigger platform migration projects.


TECHNICAL CONS

- Workload-level optimization stays in the underlying platforms; enterprises retain platform tuning responsibility.
- Requires auditability and control, often through manual processes.
- Runtime governance relies on human operation.

LEVEL 4 · ANNOUNCING

Autonomous Plane

Acceldata Autonomous Data & AI Platform



Federated compute and data. The control plane itself becomes agentic; observability, governance, and optimization run autonomously across the estate. Any workload, trusted data, anywhere.

TECHNICAL PROS

- Compounding intelligence; control plane learns across the install base.
- Removes human bottleneck in operating distributed estates at scale.
- Self-healing data quality and adaptive policy.
- Workload placement, cost optimization, and incident response at machine speed.
- The only architecture that can credibly scale governance and observability to thousands of agents across hundreds of data sources.

TECHNICAL CONS

- Early-stage deployments operate in supervised mode until the control plane has learned the environment.
- New operating category; internal stakeholders (risk, compliance, audit) need orientation on what runtime governance changes about their oversight model.

— WHERE EACH LEVEL BREAKS

What it cannot serve at all

The most useful test of any architecture is not what it handles well, it is what it cannot serve at all. The use cases below are structurally unsupported at each level. They are not roadmap items.

1

Level 1 · Single Plane
WHAT IT CANNOT SERVE

- ✗ Agents spanning AWS + Azure + on-prem with low-latency requirements.
- ✗ High-throughput transactional AI; fraud scoring at 50K TPS, real-time bidding, IoT triggers.
- ✗ Regulated workloads requiring physical separation (DORA, GDPR strict interpretation, healthcare PHI isolation).
- ✗ Edge or air-gapped on-prem agents.

2

Level 2 · Split Plane
WHAT IT CANNOT SERVE

- ✗ On-prem agentic workloads (no Google footprint outside GDC appliance).
- ✗ Workloads where data and compute must co-locate outside GCP for compliance, latency, or sovereignty.
- ✗ Compliance regimes that bar GCP entirely (parts of EU public sector).
- ✗ Accounts where Google is competitively or politically blocked.

3

Level 3 · Federated Plane
WHERE IT IS NOT THE RIGHT CHOICE

- ✗ Greenfield single-vendor deployments where consolidation is genuinely the right answer (small org, one cloud, one data platform).
- ✗ Use cases where deep platform-specific optimization beats portability.
- ✗ Customers who want a single throat to choke and accept the lock-in trade.

4

Level 4 · Autonomous Plane
WHERE IT IS NOT THE RIGHT CHOICE

- ✗ Static, single-platform, low-velocity environments where automation overhead exceeds value.
- ✗ Highly bespoke regulated workflows where every decision must have prospective human accountability documented in advance.
- ✗ Customers who treat the control plane as a feature catalog rather than as an operating layer.

◆ THE PHILOSOPHICAL THREAD: RUNTIME GOVERNANCE

What ties Levels 3 and 4 together, and what neither Level 1 nor Level 2 can fully deliver is a shift from design-time governance to runtime governance.

Design-time governance

Design-time governance assumes humans write the rules, define the catalogs, approve the access requests, and the system enforces what was decided in advance. This was sufficient when data pipelines were authored once and ran on a schedule.

Runtime governance

Runtime governance assumes the system observes behavior, adapts policy, and intervenes as workloads execute. It is the only governance model that can keep pace with autonomous agents acting on the estate at machine speed.

Runtime governance is not a Collibra or Alation feature. It is not a Unity Catalog upgrade. It is an architectural property of a control plane that spans the heterogeneous estate at runtime, observes behavioral telemetry continuously, and acts on what it sees. That property is structurally unreachable from Level 1, partially reachable from Level 2, foundational to Level 3, and fully realized at Level 4.

RUNTIME GOVERNANCE REACHABILITY

Level 1
Structurally unreachable

Level 2
Partially reachable

Level 3
Foundational

Level 4
Fully realized

— CONCLUSION · THE ARCHITECTURAL ENDPOINT

Why the path matters more than the pace

Every CIO and CDO in 2026 is being asked the same question: What is your agentic AI strategy? The temptation is to answer in features which models, which agents, which copilots. The architectural question is harder and more consequential: which level of architecture are you betting on, and can it reach the agentic enterprise you will need in three years?

◆ THE STRUCTURAL ARGUMENT, RECAPPED

Level 1 cannot become Level 3. Becoming federated would require abandoning the consolidation thesis that is Databricks' and Snowflake's entire commercial premise.

Level 2 is unstable. Google has conceded the data-gravity argument but not the compute-gravity argument. The position will resolve toward one side over the next two to three years.

Level 4 is unreachable from Level 1 or Level 2. An autonomous control plane needs visibility across the full heterogeneous estate to learn meaningfully. Vendors whose control plane stops at their walls cannot get there without rewriting their commercial models.

Level 4 requires Level 3 as a foundation. Federation is the substrate on which autonomy is built which is why Acceldata's five years of Level 3 deployment is the basis for what we are announcing now.

◆ THE ACCELDATA COMMITMENT

Since 2021, Acceldata has built and deployed a Level 3 architecture for some of the most demanding data estates in financial services, healthcare, retail, and the public sector. Federated compute running where data lives and a unified control plane spanning AWS, Azure, GCP, and on-premises. That foundation is the basis for what comes next.

ANNOUNCING THE AUTONOMOUS DATA & AI PLATFORM

Built on five years of Level 3 deployment. The first Level 4 architecture in the market. Compute planes run any workload. Trusted data, anywhere. Runtime governance, observability, and optimization operating autonomously across the heterogeneous estate.

Speak with us about how a Level 4 architecture changes what your enterprise can do with agentic AI.

◆ A CLOSING THOUGHT

“

The companies that won the cloud were the ones who recognized lift-and-shift for what it was: a starting line. The same recognition is now required of every enterprise building on agentic AI. **Where you start matters. But where you can architecturally arrive matters more.**