



TECHNICAL WHITE PAPER

Choosing a Data Governance & Semantic Layer for the AI Era

*Strategy Mosaic vs. Snowflake Horizon Catalog
vs. Databricks Unity Catalog*

EXECUTIVE SUMMARY

Enterprises are converging on the same problem from three different directions: how do you give people, BI tools, and AI agents one governed, trustworthy view of enterprise data, no matter where that data physically lives? Three platforms answer this question in meaningfully different ways — **Strategy Mosaic**, **Snowflake Horizon Catalog**, and **Databricks Unity Catalog**.

This paper compares the three on architecture, governance depth, AI-agent readiness, ecosystem reach, and total cost of ownership. There is no universally “best” answer.

Snowflake Horizon Catalog and **Databricks Unity Catalog** are the strongest choices when an organization's data and compute are already consolidated on that respective platform.

Strategy Mosaic is the strongest choice when data and governance requirements span multiple platforms — including systems like Workday, Snowflake, Databricks, and SaaS applications — and a single semantic definition needs to travel across all of them.

The right choice depends on where your data already lives, how many platforms you run today, and how much you value platform independence versus native, zero-friction integration.

WHY THIS COMPARISON MATTERS NOW

AI agents have changed the stakes for data governance. A dashboard showing the wrong number is an inconvenience; an AI agent acting on the wrong number, or exposing data it should not have access to, is a governance and compliance incident.

All three vendors have responded in 2026 with agent-specific governance: Snowflake introduced **Horizon Context** and **Agent Identity**, Databricks introduced the **Unity AI Gateway**, and Strategy introduced **Mosaic Sentinel**. Each is racing to become the trusted control point for AI, but each starts from a different architectural position — a semantic layer built to sit above any platform, versus governance native to one cloud data platform.

PLATFORM OVERVIEWS

◆ Strategy Mosaic

Strategy Mosaic (from Strategy, formerly MicroStrategy) is a **universal semantic and governance layer** rather than a data catalog tied to one warehouse. It represents data through business-friendly objects — hierarchies, calculated metrics, and multi-form attributes — and connects to **more than 200 data sources**, including files, applications, and databases. Its Sentinel component adds real-time monitoring, PII/PHI classification, and behavioral risk detection across every connected tool, positioning Mosaic as a governance layer that sits above, not inside, any single cloud platform.

◆ Snowflake Horizon Catalog

Snowflake Horizon Catalog is the **native governance and metadata layer** embedded in the Snowflake AI Data Cloud. It provides discovery, lineage, data-quality monitoring, and policy enforcement for data that lives in or is accessed through Snowflake. The 2026 **Horizon Context** capability extends this into a

governed semantic layer for AI and BI, adding AI-generated documentation, Semantic Views, and an AI Readiness Score for each data asset. Snowflake has also opened Horizon Catalog to external Iceberg tables via Apache Polaris, extending its governance reach beyond Snowflake-managed storage.

◆ Databricks Unity Catalog

Databricks Unity Catalog is the **native governance layer** for the Databricks Data Intelligence Platform, unifying access control, lineage, and discovery for tables, files, ML models, and other assets across workspaces and clouds. Its 2026 updates extend governance across a customer's entire Databricks footprint — accounts, regions, and clouds — through a four-level namespace, and introduce the **Unity AI Gateway** to govern how agents call models, invoke tools, and connect through MCP servers. **Unity Catalog Metrics** brings first-class, reusable business metric definitions into the lakehouse.

WHERE THESE PLATFORMS SIT IN THE ENTERPRISE STACK

The diagram below illustrates the two governance patterns described above: native catalogs (Horizon Catalog, Unity Catalog) embedded inside their own platform, versus Strategy Mosaic as an overlay that spans systems of record and both data platforms at once.

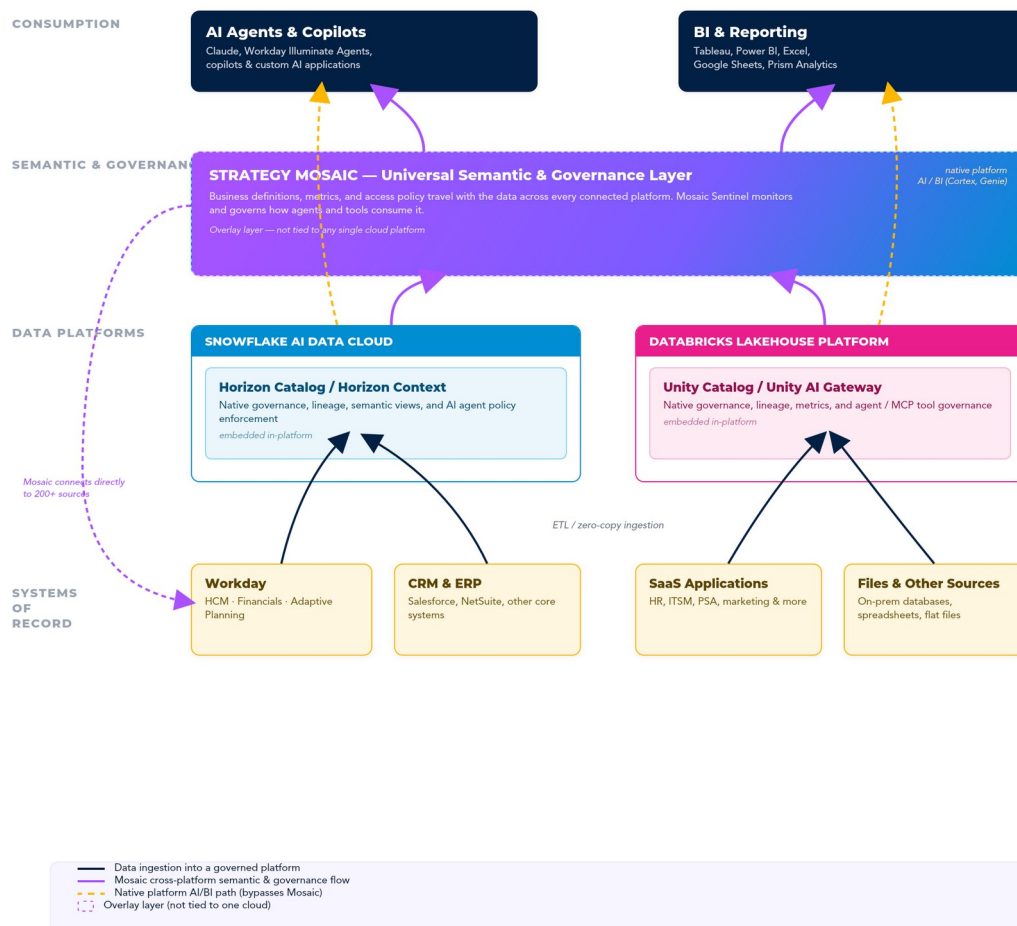


Figure 1. Systems of record feed both data platforms directly. Native catalogs govern data within their own platform and can serve platform-native AI/BI tools directly. Strategy Mosaic sits above both platforms and connects independently to source systems, giving BI tools and AI agents one governed semantic layer regardless of where the underlying data lives.

WHAT DO THESE PLATFORMS MEAN BY "CONTEXT"?

“Context” is the single most overused word in this space in 2026. Snowflake, Databricks, and Strategy all describe their governance layers as delivering “context” to AI agents — and so do the AI agents built directly into systems of record, like Workday’s Sana Self-Service Agent or Payroll Agent. These are not the same claim. At least three distinct problems get called “context,” and confusing them is one of the easiest ways to overbuy or under-govern an AI rollout.

Three Different Problems, One Word

- **Semantic / business-meaning context.** A shared, governed definition of what a business term or metric actually means — for example, ensuring “revenue” calculates the same way whether a person asks in Tableau, Power BI, or a chat agent. This is what Snowflake means by Horizon Context, what Databricks means by its new Business Glossary and Domains inside Unity Catalog, and what Strategy Mosaic means by its semantic layer of hierarchies and calculated metrics.
- **Runtime / policy context.** A real-time decision about what a specific agent, tool, or MCP connection is allowed to do right now, based on who or what is asking and what the request contains. This is what Databricks means by the Contextual Service Policies inside its new Unity AI Gateway, part of what Snowflake’s Horizon Context runtime policy enforcement does, and what Strategy’s Sentinel layer does when it monitors agent behavior.
- **Organizational / identity context.** An agent’s built-in awareness of who a person is, their role, their manager, and what they are entitled to see or do — inherited automatically because the agent is native to the transactional system of record itself, not layered on top of it. This is what Workday means when it says the Self-Service Agent “understands your organizational structure, security permissions, and unique business processes from day one,” and what ASOR means when it says agents “operate within your existing security groups and org hierarchy.”

Type of “Context”	What It Actually Means	Who Provides This	Main Limitation
Semantic / business-meaning context	A shared, governed definition of what a term or metric means (e.g., what counts as “revenue” or an “active customer”), so every tool that queries it gets the same answer.	Snowflake Horizon Context (Semantic Views); Databricks Unity Catalog Business Glossary and Domains; Strategy Mosaic’s semantic objects and hierarchies.	Governance is enforced at the meaning layer, but each platform’s version generally only covers data it can see — unless explicitly federated.
Runtime / policy context	A real-time decision about what a specific agent, tool, or MCP service is allowed to do right now, given who or what is asking and what the request contains.	Databricks Unity AI Gateway (Contextual Service Policies); Snowflake Horizon Context runtime policy enforcement; Strategy Mosaic Sentinel.	Newer capability across all three vendors — much shipped as beta or public preview in 2026; treat as emerging, not mature, until GA.

Type of "Context"	What It Actually Means	Who Provides This	Main Limitation
Organizational / identity context	An agent's built-in knowledge of who a person is, their role, their manager, and what they are entitled to see or do — inherited automatically because the agent runs inside the transactional system itself.	Workday Sana Self-Service Agent, Payroll Agent, and other Illuminate agents; ASOR's inherited permissions and org hierarchy for agent identities.	Native to one system of record and does not automatically travel to a catalog, semantic layer, or another platform unless explicitly synced or federated.

Table 1. Three distinct meanings of "context" in use across these platforms as of mid-2026.

◆ How This Compares to System-of-Record Agents

This is the distinction worth sitting with: Snowflake, Databricks, and Strategy Mosaic are primarily solving semantic and runtime context — making sure a metric means the same thing everywhere and that agent actions are policed in the moment. None of the three originates organizational or identity context; they consume it. Workday's own agents work the other way around — they are born with organizational and identity context because they run inside the system that owns the org chart, the security model, and the underlying business process, but they generally do not carry a portable, cross-platform definition of a business metric the way a semantic layer does.

In practice, these are complementary layers, not competing ones. A mature AI governance architecture typically needs all three: identity and permission context sourced accurately from the system of record (Workday, Salesforce, or similar), a shared semantic definition of what the business's numbers mean (Horizon Context, Unity Catalog's Glossary and Domains, or Mosaic), and a runtime layer that combines both to decide, in the moment, whether a specific agent action should proceed (Unity AI Gateway, Horizon's runtime enforcement, or Sentinel). Buyers should ask any vendor making a "context" claim to say specifically which of these three problems they mean — and where the other two are expected to come from.

HEAD-TO-HEAD COMPARISON

Dimension	Strategy Mosaic	Snowflake Horizon Catalog	Databricks Unity Catalog
Primary identity	Universal semantic and governance layer that sits above multiple data platforms	Native catalog and governance layer built into the Snowflake AI Data Cloud	Native catalog and governance layer built into the Databricks Data Intelligence Platform
Deployment model	Overlay / middleware; connects to 200+ sources regardless of where data lives	Embedded in Snowflake; strongest when data and compute already live in Snowflake	Embedded in Databricks; strongest when workloads run on the Databricks lakehouse
Platform independence	High — designed specifically to be platform-agnostic across warehouses, lakes, and clouds	Moderate — increasingly open via Apache Iceberg and Polaris, but governance logic is Snowflake-native	Moderate — open across clouds and Iceberg formats, but governance logic is Databricks-native
Semantic layer maturity	Deep; built on Strategy's decades-old semantic/BI engine (hierarchies, calculated metrics, multi-form attributes)	Growing quickly; new Semantic Studio, Semantic Views, and Autopilot features reduce reliance on hand-built models	Growing quickly; Unity Catalog Metrics brings first-class business metrics into the lakehouse
AI agent governance	Sentinel layer tracks and governs how agents and LLMs consume semantic definitions, independent of underlying platform	Horizon Context tracks agent access, applies runtime policy enforcement and AI Readiness Scoring within Snowflake	Unity AI Gateway governs agent and MCP tool behavior, with lineage across the full Databricks estate
Data quality / lineage	Lineage covers connected sources at the semantic layer; depth depends on source system	End-to-end lineage with root-cause analysis across ingested pipelines	Column-level lineage across a four-level namespace spanning accounts, regions, and clouds

Dimension	Strategy Mosaic	Snowflake Horizon Catalog	Databricks Unity Catalog
BI / tool connectivity	Purpose-built connectors deliver governed metrics (not just raw tables) into Tableau, Power BI, Excel, and Google Sheets via SQL/JDBC, XMLA/DAX, REST, and Python — the widest reach for consuming the semantic layer itself, independent of vendor	Excellent for raw-table access — Tableau, Power BI, and most BI tools connect natively via ODBC/JDBC. Governed Semantic Views are a different story: Power BI has no native way to query them yet, requiring a community-built connector or a bypass to raw tables	Excellent for raw-table access via native connectors and Delta Sharing. Governed Metric Views are harder to reach externally — Microsoft removed Power BI's compatibility mode for them, leaving manual SQL workarounds as the only path today
Best fit	Multi-platform estates (e.g., Workday + Snowflake + Databricks + SaaS apps) needing one governed definition of the truth	Snowflake-centric estates consolidating governance, BI, and AI workloads on a single cloud	Databricks-centric estates running data engineering, ML, and AI workloads on the lakehouse
Primary cost driver	Software licensing (user or capacity-based), typically negotiated; incremental to any warehouse spend	Included with Snowflake platform consumption; some capabilities are metered compute	Included with Databricks Premium/Enterprise tier; governed workloads still consume DBUs

Table 2. Capability comparison across the three platforms. Feature availability changes quickly; verify current GA status with each vendor before making architectural commitments.

OTHER SEMANTIC LAYER APPROACHES WORTH KNOWING

Strategy Mosaic, Horizon Catalog, and Unity Catalog are the three vendors most relevant to organizations already anchored on Snowflake, Databricks, or a multi-platform estate involving both — which is why this paper centers on them. They are not, however, the only semantic layer options on the market, and a fair evaluation should note the rest of the field. **AtScale** is Mosaic's closest direct analog: a mature, vendor-agnostic universal semantic layer with deep OLAP/MDX heritage and strong Excel and Power BI support. **Cube** takes a headless, API-first approach aimed squarely at embedded analytics and AI agents, competing directly with Mosaic on caching and MCP support rather than traditional BI governance. **dbt Semantic Layer** (built on the now open-sourced MetricFlow) appeals to teams whose metric definitions already live inside dbt's transformation layer and want metrics reviewed the same way as code. None of these displace the case for evaluating Mosaic, Horizon Catalog, and Unity Catalog against this organization's actual stack, but they are worth knowing before treating platform-native or Mosaic-style overlays as the only two categories available.

PROS AND CONS BY PLATFORM

◆ Strategy Mosaic

Strengths	Limitations
✓ Platform-agnostic by design — built specifically to unify governance across many warehouses, lakes, and SaaS systems	– Adds a distinct licensing cost and administrative layer on top of whatever platforms already exist
✓ Deep, mature semantic layer inherited from decades of BI/OLAP engineering	– Requires building or migrating semantic models rather than inheriting them for free from a single platform, though the Open Semantic Interchange (OSI) standard — which Strategy, Snowflake, and Databricks have all joined — aims to make that migration less costly over time
✓ Broadest third-party tool reach: Tableau, Power BI, Excel, Google Sheets, plus SQL, DAX, REST, and Python access	– Value depends on breadth of adoption — less compelling for organizations already fully consolidated on one cloud
✓ Governance and definitions travel with the data, reducing lock-in to any one cloud platform	– Smaller ecosystem and talent pool compared to Snowflake or Databricks native tooling
✓ Purpose-built AI governance (Sentinel) that treats semantic meaning, not just technical metadata, as the unit of control	– Real-time lineage depth for non-connected or highly custom source systems can lag native-platform catalogs
✓ In-memory caching serves repeat queries from the semantic layer itself, avoiding warehouse compute entirely rather than just discounting it	– Cost-savings claims are Strategy's own figures; results depend heavily on query repetition patterns and cache freshness configuration
✓ MCP server exposes governed metrics directly to Claude and other LLMs, so agents reason over certified definitions rather than raw SQL	– MCP support across the industry is new and evolving; enterprise-grade audit and rate-limiting maturity is still developing

◆ Snowflake Horizon Catalog

Strengths	Limitations
✓ Included with the Snowflake platform — no separate governance product to buy or integrate	– Governance logic and richest features are strongest for data that lives in or flows through Snowflake
✓ Tight integration with Snowflake compute, storage, and Cortex AI features	– Organizations with significant Databricks, on-prem, or other-cloud data still need a bridging strategy

Strengths	Limitations
✓ Rapidly maturing semantic tooling (Semantic Studio, Semantic View Autopilot) reduces manual modeling effort	– Some newer capabilities are in private or public preview, not yet generally available
✓ Strong, fast-moving investment in AI-specific governance: Agent Identity, prompt-injection detection, AI Readiness Scoring	– Deepening reliance on Horizon Catalog increases architectural dependence on the Snowflake platform over time
✓ Growing openness through Apache Iceberg and Polaris for governing data outside Snowflake-managed storage	

◆ Databricks Unity Catalog

Strengths	Limitations
✓ Included with Databricks Premium and Enterprise tiers — governance is a first-class platform capability, not an add-on	– Governance depth is strongest for Databricks-native assets; broader enterprise catalogs may still be needed alongside it for non-Databricks systems
✓ Strong fit for ML/AI development workflows already running on the Databricks lakehouse	– Feature set assumes familiarity with Databricks-specific concepts (metastores, catalogs, access modes)
✓ Extending toward cross-account, cross-region, and (soon) cross-cloud governance from a single control plane	– Governed workloads still consume Databricks Units, so cost is tied to platform usage rather than a flat governance fee
✓ Unity AI Gateway directly governs agent, model, and MCP tool behavior with column-level lineage	– Business-user-facing semantic experience (Metrics, Discover) is newer than Databricks' data-engineering capabilities
✓ Open table formats (Delta, Iceberg) support governance over data outside Databricks-managed storage	

NOT ALL 'SEMANTICS' ARE EQUAL

Every platform in this space now claims a “semantic layer.” The term has become a checkbox on a slide, but the underlying implementations differ in ways that materially change cost, governance, and how much you can trust the numbers an AI agent hands back. A semantic layer is not just a glossary of business terms. It is the combination of where the definition lives, who can change it, how fresh the underlying data is when a query hits it, and what it costs in compute or credits every time someone or something asks a question.

When evaluating a vendor's semantic layer claim, five dimensions actually separate the options:

- **Definition ownership:** Is the metric defined once and inherited everywhere, or does each consuming tool maintain its own copy that can drift out of sync
- **Data freshness pattern:** Live query against source data, scheduled batch refresh, or a cached/materialized copy that ages between updates
- **Compute and cost model:** Flat platform cost, consumption-based credits per query, or hidden compute charges that only show up at scale
- **Governance and access inheritance:** Does the semantic layer carry the source system's security model with it, or does it require rebuilding permissions in a second place
- **Definition portability:** Is the metric definition expressed in a vendor-neutral format that other tools can read, or is it locked to one platform's proprietary syntax

These distinctions matter most in exactly the scenario most enterprise clients are in today: metrics defined in Workday, Salesforce, Adaptive, and Snowflake independently, with no single system of record for what “revenue” or “headcount” actually means. A semantic layer that solves that problem looks very different from one that just adds a nicer label on top of the same fragmented data.

Definition portability is worth watching closely because it is actively changing. Snowflake launched the **Open Semantic Interchange (OSI)** in September 2025 as a vendor-neutral, YAML-based specification for semantic model constructs; Strategy joined in December 2025 and Databricks joined alongside the v1.0 release in January 2026, putting all three vendors in this comparison inside the same standard. As of mid-2026, OSI is still early: Microsoft, and by extension Power BI, has not joined, which matters given Power BI's market position. Buyers evaluating any vendor's semantic layer should ask directly whether metric definitions can be exported in an OSI-compatible format today, not on a future roadmap.

COST-BENEFIT ANALYSIS

None of the three vendors publishes simple, comparable list pricing for these capabilities, and all three are typically sold or bundled as part of a larger enterprise agreement. Rather than compare invented dollar figures, the table below compares the cost structure and where value is realized fastest for each option.

Cost Factor	Strategy Mosaic	Snowflake Horizon Catalog	Databricks Unity Catalog
License / subscription	Separate, negotiated enterprise contract; priced by users, capacity, or consumption depending on packaging	No separate license; capabilities are bundled into Snowflake accounts, with some features metered as compute	Bundled into Databricks Premium and Enterprise tiers; not purchased as a standalone product
Incremental infrastructure spend	None in the standard SaaS deployment — compute is included in the subscription. Enterprise customers can typically count cost toward existing committed cloud spend.	None beyond existing Snowflake consumption; governance features run on Snowflake compute	None beyond existing Databricks consumption; governance features run on Databricks compute
Implementation effort	Moderate to high — semantic models must be built or migrated, and connections configured to every source	Low to moderate for Snowflake-native data; higher if governing many non-Snowflake sources via Iceberg	Low to moderate for Databricks-native data; higher if governing many non-Databricks sources
Ongoing administration	Dedicated semantic-layer stewardship; pays off most when many tools and teams share one truth	Managed largely within existing Snowflake administration and RBAC practices	Managed largely within existing Databricks administration and Unity Catalog practices
Switching / lock-in cost	Lower — governance and semantics travel with the data regardless of underlying platform	Higher if governance logic is deeply embedded before a future platform migration	Higher if governance logic is deeply embedded before a future platform migration
Where value is realized fastest	Organizations already paying for multiple platforms (e.g., Workday, Snowflake, Databricks, SaaS) that need one governed number	Organizations consolidating most analytics and AI workloads inside Snowflake	Organizations consolidating most data engineering, ML, and AI workloads inside Databricks

Table 3. Cost structure comparison. Actual costs depend on data volume, number of connected sources, user counts, and negotiated enterprise terms — request a current quote from each vendor before budgeting.

◆ How to Think About Total Cost of Ownership

- **Single-platform estates.** If you are already consolidated on one platform, Snowflake Horizon Catalog or Databricks Unity Catalog will almost always show a lower incremental cost than adding Strategy Mosaic, because governance is bundled into a platform you are already paying for. The cost-benefit case for Mosaic weakens as consolidation increases.

-
- **Multi-platform estates.** If multiple platforms carry real analytics or AI workloads — for example Workday, Snowflake, and Databricks together — every additional platform without a shared semantic layer multiplies the risk of conflicting metric definitions (the classic "four definitions of revenue" problem) and the manual effort to reconcile them. Mosaic's licensing cost should be weighed against that reconciliation cost, not against a single platform's bundled catalog.
 - **Future flexibility.** A near-term platform migration or divestiture (e.g., moving from Databricks to Snowflake, or vice versa) favors a governance layer that is not embedded in the platform being replaced, since semantic definitions and access policies would otherwise need to be rebuilt.
 - **Pace of change.** AI-specific governance (agent identity, runtime policy enforcement, prompt-injection defense) is advancing quickly on all three platforms in 2026. Budget for a follow-on evaluation in twelve to eighteen months rather than treating this as a one-time decision.

DECISION FRAMEWORK

- **Choose Snowflake Horizon Catalog if:** Your analytics and AI workloads run primarily on Snowflake today, or you are actively consolidating toward it.
- **Choose Databricks Unity Catalog if:** Your data engineering, ML, and AI development already run primarily on Databricks, or you are actively consolidating toward it.
- **Choose Strategy Mosaic if:** You run multiple platforms side by side (for example an ERP/HCM system like Workday alongside Snowflake and/or Databricks), and you need one governed definition of key business metrics across all of them, independent of any single vendor's roadmap.
- **Consider combining approaches if:** It is common, and often advisable, to run a native catalog (Horizon or Unity Catalog) for platform-specific governance and a semantic layer like Mosaic above it for cross-platform consistency — the two categories are not always mutually exclusive.

CONCLUSION

All three platforms are converging on the same destination — a governed, AI-ready foundation that agents and people can trust equally — but they are approaching it from different starting points. Snowflake and Databricks are extending governance outward from their respective platforms; Strategy is extending a semantic layer downward across whatever platforms an enterprise already runs.

The right decision is not about which vendor has the most feature announcements this quarter, but about **where your data and workloads actually sit today**, and **how many platforms your governance strategy genuinely needs to span**.

This comparison reflects publicly available product information as of mid-2026. Feature availability (particularly items in preview) and pricing structures change frequently; we recommend validating current capabilities and obtaining a formal quote directly from each vendor as part of any procurement process.