

Governed Multi-Cloud Data in Excel

Exponam Analyst Intelligence vs. the Databricks Excel Add-in

A comprehensive comparison across cost, usability, flexibility and extensibility, functionality, architecture, governance, and administration.

Exponam, LLC · July 2026

Executive Summary

Earlier this year, Databricks released its own Excel Add-in into Public Preview. The Databricks launch confirms what Exponam has argued for three years: governed, business-user access to cloud data from Excel is a strategically important, permanent enterprise capability – not an edge case. Databricks has chosen to build the Databricks piece of that. They will not build the Snowflake piece, or the Fabric piece, or the Redshift piece. No cloud-platform vendor can occupy the platform-neutral position.

Exponam Analyst Intelligence is that platform-neutral layer. It connects enterprise users directly to governed data in Databricks and Snowflake today, with Microsoft Fabric in development and additional platforms to follow. It offers three governed paths to data – natural-language query, direct SQL, and a visual data builder – and it is the only product in this category with a fully local, private LLM option for regulated environments.

This paper compares the two products across the dimensions that matter to enterprise evaluation committees: cost, usability, flexibility and extensibility, functionality, architecture and performance, governance and security, and administration

TL;DR: the absence of a license fee is not the absence of a cost. The Databricks Excel Add-in is Databricks-only connector for workloads by small tech-savvy teams. Exponam Analyst Intelligence is the enterprise analyst layer – multi-platform, private-AI-powered, self-service by design, and cost-predictable at any scale.

1. The Enterprise Analyst Access Problem

Modern enterprise data lives in cloud platforms – Databricks, Snowflake, Microsoft Fabric, Redshift, BigQuery. Enterprise analysis happens in Excel. Between those two realities sits a painful, ungoverned gap that every Fortune 500 data team knows intimately: manual exports, stale CSVs, ODBC driver configurations, long IT ticket queues, and a proliferation of shadow data that no compliance officer would approve if they saw the full picture.

Business intelligence tools do not solve this problem. Tableau, Power BI, Sigma, and Qlik are visualization and dashboard platforms – powerful for dashboard creators, but demanding training, report design, and design-team support for everyone else. The analysts, finance teams, and operations leads who drive day-to-day enterprise decisions are still waiting for someone to email them a spreadsheet or update a BI report from which they can export.

In 97 of 102 recent Exponam sales conversations, prospects self-identified this exact workflow – the IT-ticket, stale-CSV, reconciliation, shadow-data cycle – as their primary pain. The pattern is platform-agnostic. Companies on Databricks described it. Companies on Snowflake described it. The tool almost always at the end of the chain: Excel.

Three concurrent market signals

- **Snowflake Ventures invested in AtScale (December 2025)** to build governed Excel and BI connectivity on Snowflake data, citing the one-billion-plus Excel user population. AtScale provides semantic-layer infrastructure; the user-facing layer above it is exactly the category Exponam occupies.
- **Databricks released its own Excel Add-in (March 2, 2026)**. A company valued at \$188 billion in its most recently announced funding round built a purpose-specific product to solve this problem for its own users. This is unambiguous market validation – and it confirms that every major cloud platform will build, buy, or partner for this capability on its own platform only.
- **Microsoft Copilot for Excel launched without live cloud-platform connectivity**. Copilot assists with formulas, summaries, and chart suggestions. It does not connect to Snowflake or Databricks data at enterprise scale, and Microsoft cannot build governed, live connectors to platforms with which it competes without creating a commercial conflict its cloud partners would not accept.

The structural implication is simple. Databricks will build an Excel connector for Databricks data. They will not build one for Snowflake. Snowflake will not cover Databricks. Microsoft Copilot will not close the data-connectivity gap. The governed, AI-powered, platform-neutral position is structurally available to one kind of company only: an independent analyst layer. Exponam is that layer.

2. What Exponam Analyst Intelligence Is

Exponam Analyst Intelligence is a Microsoft Excel add-in that connects enterprise users directly to governed data in their organization's cloud data platforms – without exports, ODBC configuration, IT intermediaries, or workflow disruption. It ships as a native VSTO Windows add-in. The product is organized around three capability layers, all accessed from a single ribbon inside Excel, all sharing the same governance model and the same data-stays-in-your-organization commitment.

CONNECT · governed multi-cloud data access

Direct, governed access to Databricks and Snowflake from a single Excel add-in, with Microsoft Fabric in development. Authentication, access control, and row-level security are enforced by the source platform's existing security model – Unity Catalog, Snowflake Horizon – and Exponam maintains no parallel permission layer. Data volumes are real: up to 10,000,000 rows per extract on an enterprise license, with no limit on the number of extracts per user. Imports land as static cell values – no live cell functions, no recalculation risk, no hidden compute consumption. Refresh runs independently of Excel recalculation: automatic at any frequency, manual, or paused, all controlled from the ribbon.

QUERY · three paths, same governance

- **Natural-language query.** Ask a question in plain English. The AI layer interprets intent, maps it to the schema, constructs a governed SQL query, executes it against the cloud platform, and returns the results to Excel together with a plain-English explanation of what is retrieved and from where. The generated SQL is open for review and lockable for sharing.
- **Direct SQL.** A full SQL scratchpad with individual and shared workspaces. Write any query the cloud platform supports – joins, window functions, subqueries, aggregations – and execute it from inside Excel against the platform's native SQL engine with full governance.
- **Visual data builder.** Point, click, filter, import. No code required. On Databricks, the visual builder uses the zero-compute Delta Sharing path – direct Parquet retrieval from cloud object storage, bypassing the SQL warehouse entirely.

INTELLIGENCE · AI inside the security perimeter

Natural-language query requires a language model to translate questions into SQL. Most products send the customer's schema – or worse, query context – to a commercial AI endpoint outside the perimeter. That is a compliance failure for financial services, healthcare, accounting, pharmaceutical, and government organizations with absolute policies against sending proprietary information to public AI APIs.

Exponam Analyst Intelligence offers a fundamentally different model. A fully local LLM runs on the analyst's machine – no data, no schema, and no tokens leave the network. SQL-specialized models from 8B to 32B parameters are supported: 8B models run efficiently on standard corporate laptops, and 14B or 32B models deliver near-zero latency on GPU-equipped machines. Organizations with existing commercial model agreements can alternatively connect their preferred provider, and Exponam passes only schema metadata – never row-level data. The Intelligence layer also executes Databricks ML Serving Endpoints directly from Excel through a guided formula builder, returning predictions inline – risk scores, forecasts, classifications – without leaving the workbook.

Zero data exposure by architecture. Exponam is not a SaaS platform. No data, no schema definition, and no query ever transits Exponam's infrastructure. Every byte moves directly between the customer's cloud platform and the analyst's machine.

3. What the Databricks Excel Add-in Is

The Databricks Excel Add-in is a Microsoft Office Web Add-in that connects Excel to Databricks data via the Databricks SQL warehouse endpoint – the same compute path that powers Databricks SQL notebooks, dashboards, and BI tool connections. Users browse Unity Catalog, apply filters, and import data through a task-pane interface; they can also write SQL directly and embed it in cells via the `=DATABRICKS.SQL()` and `=DATABRICKS.Table()` cell functions. Unity Catalog governance applies throughout. Product enables open end-user write-back capability.

As of Databricks documentation dated July 17, 2026, the add-in uses the SQL warehouse endpoint exclusively, so queries, previews, and refreshes all consume Databricks compute. AI querying, named on

the April roadmap, has not shipped in the add-in, and the AI direction of the platform suggests the eventual integration will be Genie only. There is no facility for serving external users without Databricks workspace accounts, and identity is bound to a single workspace account.

4. Cost

Exponam Analyst Intelligence pricing

Exponam Analyst Intelligence is a freemium product. The free tier is free to install and free to use – no license key, no trial period, no credit card. It supports import and query of up to one million rows with no time limit. Enterprise licensing is priced per user, per month, on a volume-tiered schedule with a \$1,000 per month minimum, and it unlocks the full ten-million-row import capacity, AI and natural-language query, ML model execution, private-LLM deployment, custom MSI installers, and dedicated support.

Users	Per user / month	Monthly total	Annual total
Freemium (to 1M rows)	Free	–	–
100 (enterprise minimum)	\$10.00	\$1,000	\$12,000
1,000	\$5.00	\$5,000	\$60,000
10,000	\$1.00	\$10,000	\$120,000
50,000	\$0.75	\$37,500	\$450,000
100,000	\$0.50	\$50,000	\$600,000

Pricing is transparent, predictable, and budgetable at any scale. A 10,000-user deployment is a fixed \$10,000 per month. There is no compute surprise.

Compute cost is scoped by path. The Delta Sharing path consumes zero DBUs at any query volume – data arrives as compressed Parquet directly from cloud object storage, and the Databricks compute plane is never invoked. The SQL-endpoint path consumes compute for query execution only. AI cost is contained to text-to-SQL translation: local models add no incremental cost, and commercial models incur translation tokens only, with no agent or MCP overhead.

The Databricks Excel Add-in cost model

The Databricks Add-in carries no per-user license fee. On first comparison, this looks attractive. The actual comparison is not license-versus-license – it is license-versus-compute consumption.

Queries, previews, and refreshes in the Databricks Add-in all run through a SQL warehouse and consume DBUs – compute accrues across the whole workflow. During trial testing in April 2026, a single analyst's casual use – four to five representative imports in one day – consumed 16 DBUs. Using a conservative 5 DBUs per analyst per day over a 20-working-day month at the \$0.70 per DBU Serverless list rate yields roughly \$70 per analyst per month; heavier use can easily be double.¹

The AI cost direction compounds this. On July 8, 2026, Databricks moved Genie to pay-as-you-go DBU pricing – 150 free DBUs per user per month, then metered per interaction. When Genie-based AI querying arrives in the add-in, it will incur metered use on top of warehouse compute. Exponam's AI cost, by contrast, is fixed at the license and contained to text-to-SQL translation.

The shared-workbook recalculation risk

Data imported via the =DATABRICKS.Table() or =DATABRICKS.SQL() cell functions lives as a formula in the worksheet. Excel formula recalculation – triggered by ordinary workbook operations, third-party add-ins, or a manual Ctrl+Alt+F9 – can silently re-execute those queries against a running SQL warehouse. In workbooks that circulate across teams, this adds unbudgeted consumption that is difficult to detect until the monthly invoice arrives. Exponam imports data as static cell values; refresh is managed from the ribbon and is unaffected by any recalculation event.

The cost gap at scale

The following table compares Exponam Analyst Intelligence enterprise pricing against observed Databricks DBU consumption at two rate assumptions. The \$25 per user figure is a floor; the \$70 per user figure is the conservative estimate based on observed trial consumption. The Exponam figures reflect the Delta Sharing path, on which query volume consumes zero DBUs regardless of scale.

Deployment size	Exponam monthly	Databricks at \$25 / user ¹	Databricks at \$70 / user ¹	Advantage
100 users	\$1,000	~\$2,500	~\$7,000	2.5x – 7x
1,000 users	\$5,000	~\$25,000	~\$70,000	5x – 14x
10,000 users	\$10,000	~\$250,000	~\$700,000	25x – 70x

At the observed trial rate, Exponam is roughly seven times less expensive at 100 users, fourteen times at 1,000 users, and seventy times at 10,000 users. Even the conservative \$25 per user floor produces multiples of 2.5x, 5x, and 25x respectively.

On the Exponam SQL-endpoint path, the differential narrows but remains substantial. Compute is consumed only when the analyst executes an import query — one warehouse call, initiated deliberately, returning data that lands in Excel as static values. There are no cell-function refreshes, no formula-recalculation events, and no background consumption of any kind; refresh is managed from the ribbon and each refresh is a single deliberate query. On the Databricks Add-in, by contrast, queries, previews, and refreshes all run through a SQL warehouse, and compute accrues across the whole workflow. An Exponam analyst on the SQL-endpoint path therefore consumes a fraction of the warehouse compute that the same analyst generates through the Databricks Add-in — and an enterprise retains the option

to move any workload to the Delta Sharing path and eliminate that consumption entirely. The absence of a license fee is not the absence of a cost.

5. Usability and Self-Service

Exponam Analyst Intelligence was designed by business users for business users. The ribbon integration, the task-pane layout, and the interaction model mirror conventions Excel users already know. Filters are applied as they are on a native Excel sheet – the same gestures, the same mental model, no new vocabulary. A finance analyst or operations manager is productive within minutes of installation, without training. Natural-language requests return a plain-English explanation of what is retrieved and from where, with the underlying SQL open for review and lockable for sharing – so a self-service user can verify the result, and a governing team can standardize it.

The Databricks Excel Add-in presents a query-builder interface in the BI-tool pattern – parameter-driven, query-centric, and structured around concepts that are second nature to a data engineer but unfamiliar to most Excel business users. It is an approach that has struggled for a decade to achieve genuine self-service adoption among business users. A Databricks-credentialed technologist will find it familiar; an analyst who does not work in Databricks will not. Every user additionally requires a provisioned Databricks workspace account before the first import.

In most enterprises, the population of Excel-centric business users – finance, operations, supply chain, compliance, planning – outnumbers the technical Databricks user base by a wide margin. Both products target the broader population in principle. In practice, Exponam is the product that population will actually use without hand-holding.

6. Flexibility and Extensibility

One add-in across platforms

Exponam Analyst Intelligence is one add-in across cloud platforms – Databricks and Snowflake today, with Microsoft Fabric in development and further platforms planned. The application layer is sufficiently abstracted that each new platform integration follows rapidly from the last. The Databricks Add-in is Databricks only, by design: a first-party connector will not extend to Snowflake, Fabric, or other competing platforms, so each additional platform in the estate requires an additional tool, with a different interface, a different account model, and a separate training burden.

Identity

An Exponam analyst can hold unlimited identity-provider identities – usable across internal and client systems alike. One install serves every environment an analyst works in: the corporate Databricks workspace, a client's Snowflake account, and a partner's shared data, each under its own governed credential. In the Databricks Add-in, identity is bound to a single workspace account.

Shared query libraries

Exponam shared query libraries are defined at any level – departments, divisions, teams, or clients – independent of platform topology. A finance team can maintain multiple governed libraries spanning its

Databricks and Snowflake sources. In the Databricks Add-in, query sharing is scoped to the Databricks workspace.

External users

Exponam serves external recipients – partners, clients, regulators – via .share files for Databricks Delta Sharing, with no workspace accounts required. The Databricks Add-in requires a Databricks workspace account for every user, which makes external distribution impractical for most organizations.

7. Functionality

AI and natural-language query

Exponam Analyst Intelligence ships natural-language query today, with any local or bring-your-own model – no row data leaves the customer network. The Databricks Add-in has no AI querying as of July 2026; the capability named on the April roadmap has not shipped, and it is expected to be Genie only, with no local or bring-your-own model option. For generation of governed, reusable sql via a text-to-sql AI generator, the Exponam model offers a zero-cost, zero token local route and minimal token, schema-only routes.

Data volume

Exponam supports extracts to 10,000,000 rows under enterprise licensing, with no limit on the number of extracts. Databricks Add-in usage exposes a one worksheet limit, approximately 1 million rows.

Metric views

Both products support Unity Catalog metric views: the Databricks Add-in imports them, including PivotTable creation added in June 2026, and Exponam accesses them natively through the SQL-endpoint path with the semantic definitions intact. The meaningful difference is scope, not mechanics. Metric views are a Databricks-specific construct: they do not govern Snowflake or Fabric data. For a multi-cloud enterprise, defining KPIs in metric views solves semantic consistency for Databricks data only – and Exponam preserves the option to extend the same analyst experience, with the same interface, across every platform where semantics live.

Write-back

The Databricks Add-in added open end-user write-back in June 2026 – a capability enterprises have typically restricted rather than distributed, because uncontrolled writes from spreadsheets into governed platforms create audit and data-quality exposure. Exponam offers write-back where required, scoped to specific use cases, so the capability arrives with the controls an enterprise deployment expects.

ML model execution

Exponam executes Databricks ML Serving Endpoints as native Excel formulas through a guided formula builder on Windows – predictions, risk scores, and classifications inline in the workbook. The Databricks Add-in has no ML execution path.

Feature comparison

Capability	Exponam Analyst Intelligence	Databricks Excel Add-in
Browse tables	Visual catalog tree across Databricks and Snowflake, with descriptions and column information shown for discovery.	Unity Catalog browser within the task pane.
Filter and select	GUI filters matching Excel conventions; drag-and-drop column reordering; saved favorites for one-click reuse.	GUI filters and column selection documented; SQL option available.
Direct SQL	Full SQL scratchpad with individual and shared workspaces; persists between sessions.	SQL query editor plus =DATABRICKS.SQL() and =DATABRICKS.Table() cell functions.
Natural-language query	Shipping – plain-English question to governed, visible, reproducible SQL; local or bring-your-own model.	Not shipped; roadmap item, expected to be Genie only.
Refresh model	Independent of Excel recalculation; auto, manual, pause, and resume from the ribbon; data lands as static values.	Manual refresh; cell functions subject to Excel recalculation, triggering warehouse queries.
PivotTable creation	Checkbox at import time.	Pivot at import; metric-view PivotTables added June 2026.
Write-back	Available where required, scoped to specific use cases.	Open end-user write-back (June 2026).
ML model execution	Databricks ML Serving Endpoints as Excel formulas via guided builder (Windows).	Not available.
External user access	Yes – .share files, no workspace accounts required.	No – every user requires a Databricks workspace account.
Data export	Unlimited CSV and .BIG export (enterprise license).	Not documented.

Databricks Add-in feature representation reflects documented capabilities as of July 17, 2026. Some entries may represent features under development rather than permanent gaps.

8. Architecture and Performance

VSTO versus Office Web Add-in

The most technically significant difference between the two products is the client framework. Exponam Analyst Intelligence uses VSTO – Visual Studio Tools for Office – on Windows: Microsoft's native, COM-based integration framework, running in-process with Excel and architected at Exponam by the co-author of the Microsoft Office COM integration layer. The Databricks Add-in uses the Office Web Add-in

framework on every platform, including Windows: JavaScript in a sandboxed browser frame, communicating with Excel through the asynchronous, rate-limited Office.js API layer.

For most light use, the Office Web Add-in framework is adequate. For high-volume data imports, the Office.js boundary introduces meaningful serialization overhead. The result is markedly faster large-dataset writes on the platform where most enterprise analysts work.

Data retrieval paths

- **Delta Sharing path (Databricks, zero compute).** Compressed Parquet delivered directly from cloud object storage via presigned URLs. The Databricks compute plane is never invoked; zero DBUs are consumed at any query volume. Requires Delta Sharing to be enabled in the customer environment – standard in most enterprise deployments, though subject to internal approval.
- **SQL-endpoint path (all supported platforms).** Queries route through the source platform's native SQL engine with the platform's full SQL syntax. Compute is consumed for query execution only, and results land as static values.
- **Natural-language path.** Questions are translated to governed SQL by a local or commercial model and executed on the same SQL-endpoint infrastructure, through the same governance. The model sees schema metadata only – never row-level data.

Published benchmarks

Exponam publishes performance benchmarks against the Open Delta COVID-19 NYT dataset, a public Delta Sharing dataset, on a standard Windows PC via the Delta Sharing path:

Record count	Time to import into Excel
100,000 rows	1 second
500,000 rows	5 seconds
1,000,000 rows	11 seconds
1,111,930 rows (full dataset)	11 seconds

These figures reflect compressed Parquet transfer, VSTO direct object-model writes, and Delta Lake partition optimization. For wider schemas of 100 or more columns, Parquet's columnar compression extends the advantage further.

9. Governance, Security, and the AI Privacy Model

Governance enforced at the source

Both products enforce governance through the source platform's existing security model. For Databricks workloads, both use Unity Catalog for access control, row-level security, and audit – neither maintains a parallel permission layer, and the products are at parity on this dimension for Databricks

data. Exponam extends the same principle to Snowflake through Snowflake Horizon, which the Databricks Add-in does not support.

No data transits Exponam's infrastructure

Exponam Analyst Intelligence is installed software, not a SaaS platform. No data, no schema definition, and no query ever leaves the organization's security perimeter to transit Exponam-controlled infrastructure. For infosec review, the attack surface is the customer's cloud platform plus the analyst's endpoint – no third-party cloud in between. For vendor management, Exponam is not a processor of customer data, so the risk-assessment scope is correspondingly narrower. Exponam's Security Center documents the architecture, data flow, authentication model, and NIST SP 800-218 (SSDF) attestation, and is available at support.exponam.com.

The AI privacy model

The default Exponam deployment runs a fully local LLM on the analyst's machine – no schema metadata, no query context, no tokens, and no data leave the network. Organizations with commercial AI agreements can connect their preferred provider instead, in which case only schema metadata is transmitted – never row-level data. Organizations that can use neither on policy grounds can disable natural-language query entirely and use the direct-SQL or visual-builder paths, which require no model at all.

The Databricks Add-in has no AI querying today, so its privacy behavior cannot yet be evaluated. Its expected Genie-based direction runs in the Databricks cloud, with per-interaction metering. For regulated industries, the distinction is not between AI available and AI unavailable – it is between AI that deploys inside the security perimeter today and AI that either has not shipped or will run outside it.

10. Installation and Administration

- **Exponam Analyst Intelligence** installs from a single code-signed installer at exponam.com – under five minutes to first governed data, with no tenant-level approval required for individual install. For IT-managed rollouts, custom MSI installers support silent deployment via SCCM, Intune, or any standard Windows tooling, independent of tenant Marketplace policy. Data recipients require no workspace accounts and no tenant-wide consent.
- **The Databricks Excel Add-in** installs from the Microsoft Office Marketplace where tenant policy permits, via admin center central deployment where it does not, or through a legacy manifest sideload. It additionally requires a Databricks workspace account for every user and, on current Excel desktop releases, tenant-wide Microsoft Entra administrator consent per Databricks documentation – an approval step that in regulated tenants routes through security review.

For self-service individual adoption, Exponam remains the faster path. For managed rollouts, the comparison comes down to deployment tooling fit – MSI pipelines versus Marketplace vetting, manifest manipulation, and Entra consent – and to the per-user workspace account requirement that applies to the Databricks Add-in only.

11. Where the Databricks Add-in Has Genuine Advantages

- **Small teams already inside the Databricks workspace.** For teams below 100 users who operate daily in Databricks, run SQL-scale queries rather than million-row imports, and hold Marketplace install permissions, the add-in is a reasonable and pragmatic choice. It avoids the \$1,000 per month Exponam enterprise minimum and uses governance the team already knows.
- **Open write-back.** Organizations that specifically want broadly distributed end-user write-back – rather than scoped, controlled write-back – will find the Databricks capability more permissive out of the box.
- **Mac-primary deployments.** On macOS both products use the Office Web Add-in framework, and the performance comparison is more balanced. Exponam's cost and multi-platform advantages still apply, but VSTO-specific and AI-related capabilities are Windows-first, and Mac-heavy deployments should test both products against representative workloads.
- **Vendor scale.** Databricks is a company valued at \$188 billion in its most recently announced funding round, and evaluation committees legitimately weigh vendor stability. Exponam counters with more than three years of Fortune 500 production deployment, an architecture proven to 100,000-plus seats, and a zero-data-exposure model – but an evaluator comparing vendor scale alone will see the difference.

12. Decision Framework

The products compete directly for Databricks-to-Excel workloads. For every other platform, only Exponam Analyst Intelligence is in the market. The table maps realistic enterprise scenarios to the better fit.

Scenario	Better fit	Primary reason
Multi-cloud enterprise (any combination of Databricks, Snowflake, and Fabric)	Exponam	The Databricks Add-in is Databricks only; no cloud vendor builds the neutral layer.
Cost-sensitive; minimizing Databricks compute spend	Exponam	Zero DBUs on the Delta Sharing path; fixed per-user cost; no recalculation consumption risk.
Large deployment of non-technical business users	Exponam	Excel-native self-service; no query builder to learn; no per-user workspace accounts.
Regulated industry requiring private AI	Exponam	Local-LLM option shipping today; the Databricks Add-in has no AI querying.
Large-volume imports (1M+ rows regularly)	Exponam	10M-row capacity; compressed Parquet path; VSTO write performance on Windows.
External data sharing (partners, clients, regulators)	Exponam	.share files work without workspace accounts.

Scenario	Better fit	Primary reason
Running Databricks ML models from Excel	Exponam	ML Serving Endpoints as native formulas; not available in the Databricks Add-in.
Analysts holding identities across internal and client systems	Exponam	Unlimited IdP identities per analyst; Databricks identity is bound to one workspace account.
Small Databricks-native team; permissive IT	Databricks	The add-in avoids the Exponam enterprise minimum; DBU consumption should be monitored.
Metric-views access on Databricks	Either – parity	Both products access Unity Catalog metric views.
Broadly distributed end-user write-back	Evaluate	Databricks ships open write-back; Exponam scopes write-back to specific use cases by design.
Mac-heavy workforce	Databricks	Both use the Office Web Add-in framework on Mac; test representative workloads.

13. Summary Specifications

Specification	Exponam Analyst Intelligence	Databricks Excel Add-in
Release status	Generally available; 3+ years in enterprise production.	Public Preview (per documentation dated July 17, 2026).
Supported platforms	Databricks and Snowflake; Microsoft Fabric in development; additional platforms planned.	Databricks only.
Data paths	Delta Sharing (zero DBU) + SQL endpoint + visual builder.	SQL warehouse endpoint only.
Compute consumption	Zero on Delta Sharing; query execution only on SQL endpoint.	Queries, previews, and refreshes all consume DBUs.
License cost	Free to 1M rows; enterprise from \$10 per user per month at 100 users to \$0.50 at 100,000; \$1,000 monthly minimum.	No license fee; consumption billed via DBUs; observed ~\$70 per user per month at list rates. ¹
AI querying	Shipping – local (8B–32B) or bring-your-own commercial model; schema metadata only; no row data leaves the network.	Not shipped; roadmap, expected Genie only; Genie metered pay-as-you-go as of July 8, 2026.
Rows per extract	10,000,000 (enterprise license); unlimited extracts.	1,000,000 (one worksheet)

Specification	Exponam Analyst Intelligence	Databricks Excel Add-in
Client framework	Native VSTO on Windows	Office Web Add-in on all platforms, including Windows.
Identity	Unlimited IdP identities per analyst (Entra ID, Okta, OAuth 2.0, OIDC federation) across internal and client systems.	Bound to a single Databricks workspace account.
Query sharing	Shared libraries at any grain – departments, divisions, teams, or clients.	Scoped to the Databricks workspace.
External user access	Yes – .share files without workspace accounts.	No – workspace account required per user.
Governance	Enforced at the source platform – Unity Catalog, Snowflake Horizon; no parallel permission layer.	Unity Catalog.
Deployment	Single installer or custom MSI (SCCM, Intune); independent of Marketplace tenant policy	Marketplace, admin center central deployment, or manifest sideload; manifest file modification
Security posture	Installed software; no data transits Exponam infrastructure; NIST SP 800-218 attestation; Unity Catalog governance.	Office Web Add-in framework; Unity Catalog governance.

¹ Based on observed April 2026 trial consumption of 5 DBUs per analyst per day at the \$0.70 per DBU Serverless list rate over 20 working days. Actual consumption varies with query frequency, warehouse type, dataset size, and contracted rates; formula recalculation in shared workbooks adds further unplanned consumption. Cost-at-scale figures reflect the Exponam Delta Sharing path; on the SQL-endpoint path Exponam adds compute for query execution only and the differential narrows but remains substantial. Trial observations predate subsequent preview releases; organizations should validate against their own workloads.

About Exponam

Exponam, LLC is the foremost firm at the intersection of modern cloud data platforms and Microsoft Office. Our engineers include co-authors of core technology upon which Microsoft Office still relies, with decades of experience spanning relational databases, data lakes, encryption, and financial-services data delivery. Exponam Analyst Intelligence is deployed in Fortune 500 environments including Nasdaq, Eversource Energy, and Mizuho on an architecture that supports deployments of 100,000-plus seats with no re-architecture required.