RJL SOLUTIONS

TECHNICAL GUIDE / SECURITY AND RETRIEVAL

Patterns and Product Fit

Classification, enforcement, retrieval, SharePoint, gateways, vector stores and semantic layers, with two lean stacks for smaller companies.

CYERA · SNOWFLAKE · DATABRICKS · PURVIEW · DATA.WORLD · ATSCALE

Richard White
RJL Solutions · rjlsol.com

OCTOBER 2026

01 / HOW TO USE THIS PAPER

Patterns first, products second.

Many AI security failures sit in the gaps between products. A label never reaches the warehouse. A search index ignores row access. Someone asks a gateway to do a retrieval system's job. This paper sets out the patterns that close those gaps, then shows how current products fit them. I wrote it for architects, security engineers and data platform leads.

Three rules hold throughout

- **One classification vocabulary.** Sensitivity is defined once and carried as tags and labels to every place data is stored, indexed or answered.
- **Enforce where data is read.** Masking, row access and document permissions apply at the platform and at retrieval, before a model sees anything. A catalog, a semantic layer or a gateway is not the enforcement point.
- **Meaning is context, not authority.** A semantic layer or ontology helps the system choose the right answer. It never decides who may see the data.

Data products come in three delivery types, as Data Products in Practice defines them: operational (transactional included), analytical and knowledge. Gold is the publishing convention for analytical products. Operational products may be fed from gold or from the system of record, and knowledge products from governed documents, each under its own contract.

How to read the product statements

Each product claim was checked against the vendor's documentation or announcement in October 2026 and carries a numbered source. Status marks: GA generally available, PREVIEW public or private preview, BETA beta, ANNOUNCED announced but not yet shipped. Products change quickly: confirm current behavior before you design around it.

DISCLOSURE The author has worked closely with Cyera, ServiceNow, Snowflake, Databricks and Microsoft as a customer and practitioner. No vendor reviewed or paid for this paper. Products are named as examples of patterns, not endorsements.

---

02 / CLASSIFICATION

One vocabulary, many enforcement points.

You discover sensitivity in one place and enforce it in many. Decide which system is authoritative for each step, and which way tags and labels flow. Skip that decision, and a warehouse tag, a document label and a catalog field drift apart. Then nobody can say which one a control actually reads.

ROLE	EXAMPLE PRODUCTS	WHAT IT IS AUTHORITATIVE FOR
Discover and classify across the estate	Cyera (data security posture management)	Where sensitive data lives across clouds, SaaS, Microsoft 365, Snowflake and Databricks, and who can reach it [1][2][3]
Label documents and email	Microsoft Purview sensitivity labels and DLP	The label on a Microsoft 365 document, and the data loss prevention policy that reads it
Tag structured data	Snowflake object tags; Databricks Unity Catalog governed tags	The tag a masking or row access policy reads at query time [6][7]
Publish and explain	data.world (ServiceNow) [11]	What each tag means, who owns it, and where it applies. It shows enforcement; it does not perform it [9][10]

Recommended flow

01 **DEFINE** Security and data governance own one sensitivity scheme, mapped once to Purview label names and to platform tag values.

02 **DISCOVER** The posture tool finds and classifies sensitive data across platforms, and can find data the platforms' own classifiers miss.

03 **APPLY** Labels land on documents and tags land on columns, through the integrations in section 03 or a governed pipeline.

04 **ENFORCE** Platform policies read tags; retrieval filters read document permissions and labels.

05 **PUBLISH** The catalog harvests tags and policies so people can see what is protected and why.

06 **RECONCILE** A scheduled check compares discovered sensitivity with applied tags and labels, and opens a ticket for every gap.

Cyera also feeds ServiceNow. It can raise ITSM incidents for findings, send classification to the CMDB and take asset ownership back [4]. That's the natural route for the reconcile step's tickets.

03 / ENFORCEMENT AT THE DATA

Make the tag do the work.

A tag protects data only where a correctly bound policy enforces it. On every approved serving path, verify tag propagation, policy arguments and effective enforcement.

CAPABILITY	SNOWFLAKE	DATABRICKS
Tag-driven column masking	Tag-based masking policies GA. Enterprise Edition or higher [12]	ABAC column mask policies matched by governed tags GA since May 2026 [7]
Tag-driven row access	Row access policies GA (Enterprise Edition) [13]; attached through tags PREVIEW since July 2026 [6]	ABAC row filter policies GA [7]
Built-in classification	Platform classification of sensitive data	Agentic data classification applies governed tags continuously GA; custom classifiers BETA [7]
Cyera integration	Applies a Snowflake tag to flagged columns, which drives native tag-based masking GA (June 2026) [1] [8]	Scans Unity Catalog to discover and classify [3]. Writing tags back is not documented
Catalog integration	data.world harvests object tags, tag-based masking and row access policies [9]; a premium tag sync automation that writes tags back to Snowflake was released in beta in 2024 [10]	A premium data.world automation publishes descriptions and comments to Databricks tables and columns [47]

Design notes

- On Databricks, let Unity Catalog classification apply the governed tags. Use Cyera's findings to check coverage and raise gaps, since nobody has documented a Cyera write-back to Unity Catalog.
- Databricks row filters and masks need Databricks Runtime 12.2 or later. Dedicated access mode needs 15.4 or later with serverless compute [14]. Older runtimes fail safe and return no data.
- Test the negative. A role without rights should see masked values and no restricted rows, through SQL, BI, notebooks and every AI path.

04 / RETRIEVAL FOR AI

Filter before the model sees anything.

Retrieval-augmented generation fails safely only when the search step returns nothing the user couldn't open directly. The most common gap is a search index that runs with its builder's privileges and ignores the policies on its source.

STORE	DOCUMENTED BEHAVIOR	WHAT YOU MUST ADD
Snowflake Cortex Search	Runs searches with the owner's rights. The querying role needs only USAGE on the service, its database and schema, and can query any indexed data whatever its own privileges on the source [15][41]. Cyera showed masked values returned unmasked this way in 2025 [17]	A server-side filter on an access attribute, derived from the verified caller, never supplied by the client [16]
Databricks AI Search (formerly Vector Search)	Indexes are governed by Unity Catalog. Row and column level permissions are not supported [18]	Application-level access filters through the filter API, on columns carried from the source
Azure AI Search	Security filters on identity strings GA; Entra identity enforcement at query time PREVIEW; SharePoint permission ingestion PREVIEW; Purview label enforcement PREVIEW [19]	Choose the generally available filter pattern for production unless you accept preview terms

The pattern

01 INDEX Carry each chunk's source, version, effective date, sensitivity label and access list into the index.

02 IDENTIFY Resolve the caller from the verified session, including group membership, in the service layer.

03 FILTER Build the access filter server side, and apply it in the search call. Never pass it through the client or the model.

04 SYNC Propagate deletions and permission changes to the index, and measure the lag.

05 PROVE Run a reconciliation test: for sampled users, every returned chunk must be openable by that user at the source.

NOT THE GATEWAY'S JOB An AI gateway sees requests. It doesn't see the index's document permissions. It can pass the user's identity through. Permission trimming belongs in the retrieval call.

05 / SHAREPOINT AS A SOURCE

Fix oversharing before you index it.

SharePoint is usually the biggest and messiest knowledge source. Copying it into an AI index keeps its oversharing, and often loses its permissions. Choose the route by how much of the permission model you need to keep.

ROUTE	PERMISSIONS	WATCH FOR
Microsoft 365 Copilot Retrieval API	Security-trims results for the calling user; delegated permissions only; returns sensitivity label metadata and can filter by label [20]	Requires the user's delegated token. An invalid KQL filter runs unscoped, so validate it [48]. Needs a Microsoft 365 Copilot add-on license, or pay-as-you-go PREVIEW for SharePoint and connectors [48]
Azure AI Search SharePoint indexer	Ingests user, Entra group and SharePoint group permissions PREVIEW [21]	Parent-scope permission changes need a resync; 'Anyone' and organization links, guests and Entra groups nested in SharePoint groups are not supported [21]
Snowflake Openflow Connector for SharePoint	A variant indexes files and their permissions into Cortex Search GA [22]	Folder moves and renames are not captured incrementally; confirm how permission changes sync [22]
Databricks Lakeflow Connect SharePoint	Does not ingest file-level permissions [23]	Use only for libraries everyone in scope may read, or add your own access list

Before indexing

- Find overshared sensitive content and missing or wrong labels. Cyera identifies mislabeled content, and who can reach it, Copilot users included [2].
- Restricted Content Discovery hides a site's content from organization-wide search and Copilot while it's under review. It doesn't change permissions. People with access still have it [24].
- Apply sensitivity labels before content is indexed, so the label can travel with each chunk.

RULE If a route cannot carry the permissions, restrict it to content that everyone in scope may already read.

Inspect traffic. Do not trim permissions there.

An AI gateway is the single controlled route to models and tools. It enforces which models and tools may be used, inspects prompts and responses, logs every call and caps cost. It complements retrieval filtering. It doesn't replace it.

PRODUCT	DOCUMENTED CAPABILITIES
Databricks Unity Gateway	Guardrails for PII exposure, prompt injection, jailbreaks and unsafe content; model access governed through Unity Catalog; MCP tools enabled or disabled individually; unified tracing of model and tool activity; hard spend caps. Cyera is listed as an upcoming security integration [5]
Azure API Management AI gateway	Token quotas and per-consumer token metrics; load balancing with circuit breakers; a content safety policy; REST APIs exposed as MCP servers; prompt and completion logging to Azure Monitor [25]
Snowflake Cortex Guard, AI Guardrails and AI Gateway	Cortex Guard filters harmful responses GA (2024) [26]; Cortex AI Guardrails protect Snowflake CoWork and Cortex Agents against prompt injection and jailbreaks GA (May 2026) [43]; Cortex AI Gateway, for agent access to models, tools and MCP servers with spend limits, ANNOUNCED August 2026, public preview pending [44]
Cyera Agent Guardian	Discovers agents and MCP servers, governs permitted actions and blocks risky tool calls. Launched August 2026, covering platforms including Snowflake Cortex AI and Microsoft 365 Copilot [45]; general availability is not stated. Cyera's integrations page lists connectors for Claude Enterprise, Gemini, ChatGPT Enterprise and Microsoft Agent 365 [3]

What belongs where

CONTROL	WHERE IT BELONGS
Which documents and rows a user may see	Platform policies and the retrieval filter
Which models and tools an application may call	Gateway
PII, injection and unsafe-content inspection	Gateway, plus platform guardrails
Cost limits and the record of each call	Gateway, linked to the request evidence
Which agents exist and what data they can reach	Posture and inventory (for example Cyera)
What a coding agent in the IDE is told and may run	IDE settings: version-controlled rules guide the model; enforced MCP allowlists, run modes and ignore files limit tools and reads; none replaces authorization at the source [53, 54]

07 / CHOOSING A VECTOR STORE

Pick by permission model, not by benchmark.

For enterprise retrieval, how a store handles permissions, deletions and freshness matters more than raw search speed. Prefer the store that sits next to the governed data, unless its permission model can't meet the requirement.

REQUIREMENT	SNOWFLAKE CORTEX SEARCH	DATABRICKS AI SEARCH	AZURE AI SEARCH
Governed by the platform catalog	Snowflake object and privileges [15]	Unity Catalog object [18]	Azure resource and roles
Per-user document filtering	Server-side filter on an attribute [16]	Application filter through the filter API [18]	Security filters GA; Entra enforcement PREVIEW [19]
Source row and column policies apply	No: evaluated as the owner, not the caller [15][41]	No: not supported [18]	Not applicable: index fields only
SharePoint permissions carried	Openflow connector with ACLs GA [22]	Not by the SharePoint connector [23]	SharePoint indexer PREVIEW [21]
Freshness	Refreshes from the source query	Delta Sync from a source table with change data feed [18]	Indexer schedule or push
Best fit	Snowflake-centered estates	Databricks-centered estates	Microsoft 365 and Azure-centered estates

Search engines such as OpenSearch and Elasticsearch offer document-level security, with licensing that varies [49][50]. Most purpose-built vector databases offer metadata filtering, not an enterprise permission model. With those, permission trimming, deletion propagation and audit become your application's job, so test them as such.

TEST For each store, delete a source document, revoke a user and change a label. Measure how long each change takes to reach answers. Fail the release if it exceeds the agreed limit.

08 / SEMANTIC LAYERS AT SCALE

One definition, every tool, every platform.

A semantic layer makes business definitions explicit, and reusable by BI tools and AI agents. On one platform, the native layer is usually enough. Across several warehouses and BI tools, a universal layer saves you defining revenue three times.

OPTION	WHAT IT OFFERS	FIT
Snowflake semantic views	Semantic objects in Snowflake read by Cortex Agents, Cortex Analyst and SQL [27]	Snowflake-centered estates
Databricks metric views	Unity Catalog metric views queried from SQL, dashboards, Genie Agents and alerts, and from BI tools including Power BI, Tableau and Sigma [28]	Databricks-centered estates
AtScale	Universal semantic layer with consistent metrics across Snowflake, Databricks and BigQuery, served to Power BI, Tableau, Excel and AI agents [29]; SML modeling language open-sourced in 2024 [30]; MCP server in the Databricks MCP Marketplace since December 2025 [31]	Several platforms or BI tools, or migration from legacy OLAP cubes
dbt Semantic Layer	Metrics defined in dbt with MetricFlow, open-sourced under Apache 2.0 in October 2025 [32]	Teams that already model in dbt

Interoperability

The Open Semantic Interchange, started by Snowflake and now joined by more than 50 organizations, was accepted into the Apache Incubator as Apache Ossie in July 2026. It's an open specification for semantic layers and ontologies [33]. AtScale joined in January 2026 [29]. Converters exist, for dbt for example. Treat native support in products as something to verify. Don't assume it.

KEEP THE BOUNDARY

A semantic layer must never become a second access-control system. Confirm the user's identity reaches the platform, so masking and row access still apply. AtScale's MCP announcement relies on Unity Catalog for access control [31]. Verify identity passthrough for your own deployment.

09 / A LEAN STACK ON SNOWFLAKE

For PE and VC-backed and smaller companies.

A portfolio company doesn't need the full reference implementation. One platform, its native features and a small team can take a request from source data to a governed AI answer. Add a separate product only where the platform has a real gap.

STEP	NATIVE OPTION	NOTE
Ingest	Openflow connectors, including SharePoint with permissions GA [22]	Start with the two or three sources the first use case needs
Transform to gold	dbt Projects on Snowflake GA (November 2025) [34]	Gold publishes analytical products; bronze and silver stay internal
Meaning	Semantic views [27]	Define the handful of metrics the first use case depends on
Protect	Object tags with masking and row access policies [12][13]; row access attached through tags PREVIEW [6]	Requires Enterprise Edition or higher
Retrieve	Cortex Search with a server-side access filter [15][16]	Never rely on the service's owner's rights alone
Answer	Cortex Agents and Snowflake CoWork (formerly Snowflake Intelligence) GA (November 2025) [35][42]; Cortex AI Guardrails GA [43]	Keep the first release read-only
Serve applications	Hybrid tables for operational products that need fast reads and writes GA, AWS and Azure commercial regions [40]	Only when an application needs them
Evidence	Access History [36]	Requires Enterprise Edition or higher

EDITION MATTERS Masking policies, row access policies and Access History all need Snowflake Enterprise Edition or higher [12][13][36]. A Standard Edition account can't use native masking, row access policies or Access History.

The team

Three part-time roles cover it. A data engineer owns pipelines and gold products. An analytics or AI lead owns definitions and the assistant. A security owner owns tags, policies and access reviews. Each needs a named business counterpart who approves meaning.

The same journey, native to the lakehouse.

The Databricks version follows the same steps. Unity Catalog carries governance from tables to models, agents and the gateway, which keeps a small team's control surface in one place.

STEP	NATIVE OPTION	NOTE
Ingest	Lakeflow Connect managed connectors	The SharePoint connector does not carry file permissions [23]
Transform to gold	Declarative pipelines into Unity Catalog tables	Gold publishes analytical products; bronze and silver stay internal
Meaning	Unity Catalog metric views [28]	Define the handful of metrics the first use case depends on
Protect	Governed tags, ABAC row filter and column mask policies, and agentic classification GA [7]	Check runtime and access-mode requirements [14]
Retrieve	AI Search (formerly Vector Search) with application-level filters [18]	No native row or column permissions in the index
Answer	Genie Agents (formerly Genie spaces) GA [37][46] and Agent Bricks Knowledge Assistant GA [38], through Unity Gateway guardrails [5]	Keep the first release read-only
Serve applications	Lakebase synced tables: a read-only, low-latency copy of gold for operational products [39]	Only when an application needs fast lookups

SMALL-TEAM ADVANTAGE — Tags, policies, model access, tool permissions and traces all sit under Unity Catalog. One security owner can review them in one place. When the team is three people, that matters more than any single feature.

The team

The same three part-time roles apply: data engineering, analytics or AI, and security, each with a business counterpart who approves meaning.

Add components when a trigger appears, not before.

The lean stacks are a starting point. They aren't a ceiling. Each addition below has a clear trigger. Without the trigger, the component adds cost and one more place for controls to drift.

TRIGGER	ADD	WHY
A second warehouse or several BI tools	A universal semantic layer such as AtScale	Keeps one definition of each metric across platforms and tools [29]
Sensitive data across many clouds and SaaS tools	Data security posture management such as Cyera	Adds discovery across platforms and maps who can reach sensitive data [1][3]
Large document estates with complex permissions	A retrieval route that carries permissions, such as the Copilot Retrieval API or a permission-aware index	Native connectors may drop permissions [20][23]
Many teams publishing data products	An enterprise catalog such as data.world	Shows owners, meaning, tags and policies in one place [9]
Agents that write, not only read	A separate command path, approvals and duplicate protection	A read path must never become an unreviewed write path
Regulators, auditors or investors asking for evidence	Request evidence and release gates	Reconstruct any answer: caller, sources, versions, policy decisions

The full design is in The Reference Implementation. For regulated firms, AI Governance Evidence maps the same evidence to supervisory expectations.

Prove the negative before the positive.

Every pattern in this paper has a test that should fail safely. Run them before the first release and after every material change. A result that only shows the assistant answering well proves nothing about what it refuses.

TEST	EXPECTED RESULT
A user without rights asks for a masked field	Masked value through SQL, BI, semantic layer and every AI path
A user asks about a document they cannot open	Not retrieved, not summarized, not cited
A document is deleted at the source	Gone from answers within the agreed window
A user's access is revoked	Fresh authorization denies the next request; otherwise measure and approve the revocation lag
A label changes from internal to confidential	New handling applies within the agreed window
A tagged table lacks a matching policy argument or signature alias	Binding gap blocks release; test denial with the intended caller [52]
A document contains instructions to the assistant	Treated as content; no change to tools or permissions
A gateway guardrail or policy service is unavailable	The request is restricted or refused, never passed unchecked
A reviewer reconstructs one real request	Caller, sources, versions, policy decisions and result are all on record
The response to an approved write is lost	Reconciled by operation ID; the action is neither repeated nor reported as failed

Any write also needs an action contract. That means approval bound to the exact proposal, with an expiry; one operation ID per intended write; reconciliation of an unknown outcome before any retry; and a decision, made in advance, between rollback and a compensating action [51].

For Snowflake tag-based row access in open preview, no matching argument or signature alias means nothing enforces the policy. Check column types and binding coverage before testing outcomes [52].

START HERE Pick one use case. Run these ten tests on it, and fix what fails before adding users. The tests take days. Recovering from a disclosure takes far longer.

ARCHITECTURE UPDATE / DATA SECURITY TO RUNTIME

Use discovery findings to improve permitted use.

DSPM discovers sensitive data and exposure conditions. Runtime services enforce the permissions that make AI use defensible.

RESPONSIBILITY	CONNECTION TO THE ARCHITECTURE
Discover and assess	A DSPM solution helps identify sensitive data, excessive exposure and changing security posture across supported stores. Findings need validation and accountable remediation.
Connect the catalogs	Map findings to stable product and source identifiers. Publish reviewed classification and lineage references to enterprise and platform catalogs through supported integrations. Define conflict resolution and synchronization ownership.
Implement policy as code	Owners approve purpose and access rules. Engineers version policy bindings and tests, review them in the IDE/CI workflow, deploy to supported enforcement points and check drift. A label alone is not an access decision.
Protect retrieval	Preserve source permissions and relevant classifications through chunking and indexing. Apply tenant and user filters before retrieved text reaches the model; address revoked rights, derived stores and caches.
Protect action and output	Use scoped tool identities, command authorization, approved egress and protected logs. Test prompt injection, prohibited extraction, cross-tenant access and uncertain command results.

Cyera is an example of the DSPM category, not a required component or a claim that it automatically enforces every RAG path. Evaluate the current connectors, licensing and integration behavior for the chosen stack. Enterprise catalogs, Microsoft Purview and platform catalogs can hold complementary metadata; they do not automatically share one policy engine.

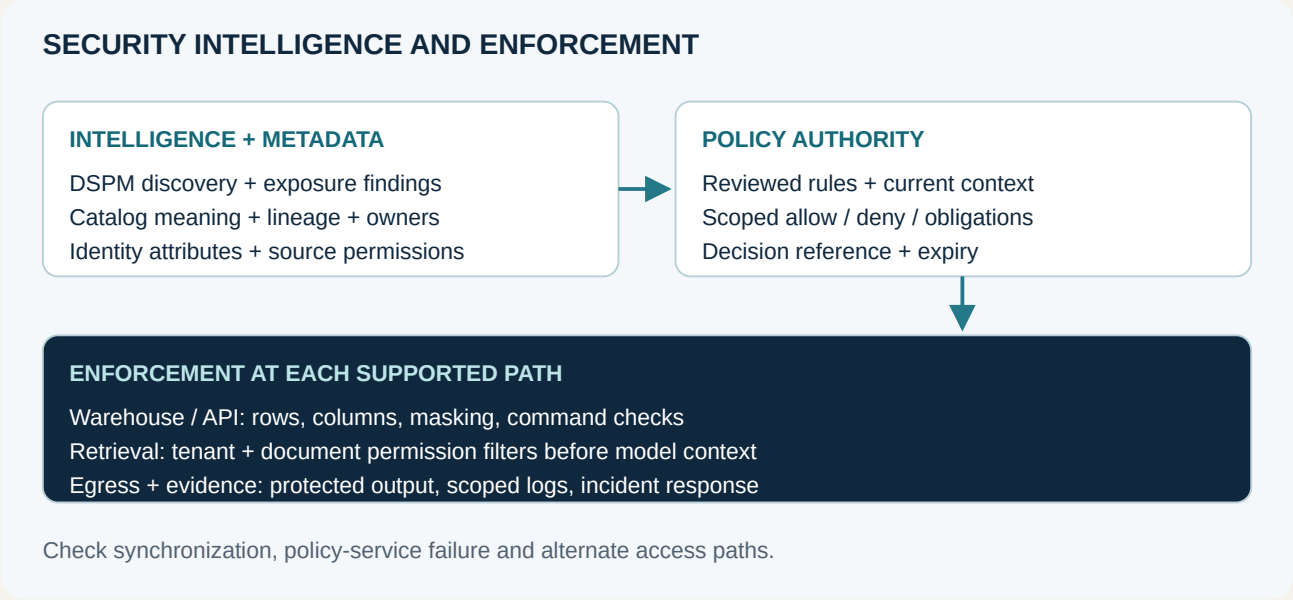
The result should be more usable, appropriately protected data: fewer unknown copies, clearer ownership and tested routes into analytics and AI. A smaller firm can start with existing discovery and native access services when those meet its scope.

Design references checked 29 October 2026: Data mesh principles · Document-level access control · RAG security. These additions are RJL implementation recommendations, not a vendor certification or a compliance determination.

ARCHITECTURE UPDATE / ENFORCEMENT BOUNDARIES

Make protection independent of the prompt.

Security intelligence improves the decision. The serving path must enforce the resulting rule.



RJL capability view. Products can combine responsibilities; verify coverage and integration for each actual path.

For an illustrative quarterly-spend question, expose an approved aggregate view with defined currency, period and entity scope. Remove unnecessary account identifiers at the serving layer. Apply field protection there when detail is permitted. Do not rely on an MCP response telling the model to remember masking in generated SQL.

An unsalted SHA-256 hash of a predictable identifier is not adequate proof of anonymity. Choose suppression, tokenization or a reviewed keyed pseudonymization design for the purpose, with protected secrets or mappings, linkage analysis and access controls. Pseudonymized data can remain identifiable; assess it accordingly.

A gateway can enforce provider, credential, budget and egress controls within its coverage. It cannot retroactively make unauthorized retrieval safe. Ensure document filters precede model context, and separately test tool authorization, cache reuse, deletion and permission revocation.

Technical reference: ICO guidance on pseudonymisation, reviewed 29 October 2026.

Design references checked 29 October 2026: Data mesh principles · Document-level access control · RAG security. These additions are RJL implementation recommendations, not a vendor certification or a compliance determination.

Checked in October 2026.

Numbers in the text refer to this list. Confirm current vendor documentation before relying on a release status or behavior.

1 Cyera + Snowflake: govern every agent, secure every dataset (June 2026). https://www.cyera.com/blog/cyera-snowflake-move-at-the-speed-of-ai

2 Cyera + Microsoft partnership. https://www.cyera.com/partnership/microsoft

3 Cyera integrations. https://www.cyera.com/integrations

4 Cyera and ServiceNow (ITSM and CMDB) via Cyera integrations. https://www.cyera.com/integrations

5 Databricks: AI governance at Data + AI Summit 2026, Unity Gateway (16 June 2026). https://www.databricks.com/blog/ai-governance-data-ai-summit-2026-whats-new-unity-ai-gateway

6 Snowflake release note: tag-based data protection policies, public preview (21 July 2026). https://docs.snowflake.com/en/release-notes/2026/other/2026-07-21-tag-based-policies-preview

7 Databricks: ABAC policies, governed tags and data classification generally available (13 May 2026). https://www.databricks.com/blog/abac-row-filtering-and-column-masking-policies-governed-tags-and-data-classification-are-now

8 Cyera and Snowflake press release (2 June 2026). https://www.cyera.com/press-releases/cyera-and-snowflake-advance-secure-ai-agent-adoption-for-the-enterprise

9 data.world: Snowflake tag and policy harvesting. https://whatsnew.data.world/enhance-your-data-governance-with-snowflake-tag-and-policy-harvesting-F7UIw

10 data.world March 2024 product launch: Snowflake tag sync automation (beta, premium). https://whatsnew.data.world/data.world-march-product-launch-3FYEgM

11 ServiceNow acquires data.world (2025). https://www.constellationr.com/insights/news/servicenow-launches-workflow-data-network-acquires-dataworld

12 Snowflake: dynamic data masking (Enterprise Edition). https://docs.snowflake.com/en/user-guide/security-column-ddm-intro

13 Snowflake: row access policies (Enterprise Edition). https://docs.snowflake.com/en/user-guide/security-row-intro

14 Databricks: row filters and column masks requirements. https://docs.databricks.com/aws/en/data-governance/unity-catalog/filters-and-masks/

15 Snowflake: Cortex Search overview (owner's rights, privileges). https://docs.snowflake.com/en/user-guide/snowflake-cortex/cortex-search/cortex-search-overview

16 Snowflake: getting started with access controls for Cortex Search. https://www.snowflake.com/en/developers/guides/getting-started-with-access-controls-for-cortex-search/

17 Cyera: unexpected behavior in Snowflake Cortex AI (April 2025). https://www.cyera.com/blog/unexpected-behavior-in-snowflakes-cortex-ai

18 Databricks AI Search (formerly Vector Search). https://docs.databricks.com/aws/en/ai-search/ai-search

Checked in October 2026.

Numbers in the text refer to this list. Confirm current vendor documentation before relying on a release status or behavior.

19 Azure AI Search: document-level access overview. https://learn.microsoft.com/en-us/azure/search/search-document-level-access-overview

20 Microsoft 365 Copilot Retrieval API. https://learn.microsoft.com/en-us/microsoft-365-copilot/extensibility/api-reference/copilotroot-retrieval

21 Azure AI Search: SharePoint indexer permission metadata (preview). https://learn.microsoft.com/en-us/azure/search/search-indexer-sharepoint-access-control-lists

22 Snowflake: Openflow Connector for SharePoint. https://docs.snowflake.com/en/user-guide/data-integration/openflow/connectors/sharepoint/about

23 Databricks: SharePoint connector limitations. https://docs.databricks.com/aws/en/ingestion/lakeflow-connect/sharepoint-limits

24 Microsoft: Restricted Content Discovery. https://learn.microsoft.com/en-us/sharepoint/restricted-content-discovery

25 Azure API Management: AI gateway capabilities. https://learn.microsoft.com/en-us/azure/api-management/genai-gateway-capabilities

26 Snowflake release note: Cortex Guard generally available (24 July 2024). https://docs.snowflake.com/en/release-notes/2024/other/2024-07-24-cortex-llm-updates

27 Snowflake: overview of semantic views. https://docs.snowflake.com/en/user-guide/views-semantic/overview

28 Databricks: Unity Catalog metric views. https://docs.databricks.com/aws/en/business-semantics/metric-views

29 AtScale joins the Open Semantic Interchange (27 January 2026). https://www.businesswire.com/news/home/20260127372541/en/

30 AtScale open-sources SML (September 2024). https://www.businesswire.com/news/home/20240910415282/en

31 AtScale MCP Server in the Databricks MCP Marketplace (1 December 2025). https://www.atscale.com/?p=28219

32 dbt Labs open-sources MetricFlow (14 October 2025). https://www.getdbt.com/blog/dbt-labs-affirms-commitment-to-open-semantic-interchange-by-open-sourcing-metricflow

33 Apache Ossie (incubating): the new name for Open Semantic Interchange (July 2026). https://ossie.apache.org/updates/ossie-enters-apache-incubator/

34 Snowflake release note: dbt Projects on Snowflake generally available (6 November 2025). https://docs.snowflake.com/en/release-notes/2025/other/2025-11-06-dbt-projects-on-snowflake-ga

35 Snowflake release note: Snowflake Intelligence, now Snowflake CoWork, generally available (4 November 2025). https://docs.snowflake.com/release-notes/2025/other/2025-11-04-snowflake-intelligence

36 Snowflake: Access History (Enterprise Edition). https://docs.snowflake.com/user-guide/access-history

Checked in October 2026.

Numbers in the text refer to this list. Confirm current vendor documentation before relying on a release status or behavior.

37 Databricks: AI/BI Genie (Genie spaces) generally available (June 2025). https://www.databricks.com/en/blog/aibi-genie-now-generally-available

38 Databricks: Agent Bricks Knowledge Assistant generally available. https://community.databricks.com/t5/announcements/databricks-agent-bricks-knowledge-assistant-is-now-generally/td-p/148506

39 Databricks: serve lakehouse data with Lakebase synced tables. https://docs.databricks.com/aws/en/oltp/projects/sync-tables

40 Snowflake: Unistore hybrid tables generally available. https://www.snowflake.com/en/news/press-releases/snowflakes-unistore-unifies-transactional-and-analytical-data-with-the-general-availability-of-hybrid-tables/

41 Snowflake: query a Cortex Search service (privileges). https://docs.snowflake.com/en/user-guide/snowflake-cortex/cortex-search/query-cortex-search-service

42 Snowflake renames Snowflake Intelligence to Snowflake CoWork (3 June 2026). https://itbrief.com.au/story/snowflake-renames-ai-tools-amid-broader-platform-push

43 Snowflake release note: Cortex AI Guardrails for Snowflake CoWork and Cortex Agents generally available (14 May 2026). https://docs.snowflake.com/release-notes/2026/other/2026-05-14-cortex-ai-guardrails-si-cortex-agents

44 Snowflake launches Cortex AI Gateway (3 August 2026). https://cybermagazine.com/news/snowflake-launches-cortex-ai-gateway-to-govern-agentic-ai

45 Cyera launches Agent Guardian (3 August 2026). https://www.cyera.com/press-releases/cyera-launches-agent-guardian-to-secure-the-autonomous-workforce

46 Databricks: AI/BI and Genie One release notes 2026 (Genie spaces renamed Genie Agents, July 2026). https://learn.microsoft.com/en-us/azure/databricks/ai-bi/release-notes/2026

47 data.world: Databricks Publisher automation, updating descriptions and comments. https://docs.data.world/en/251703-updating-descriptions-comments-for-databricks-tables-and-columns.html

48 Microsoft: Copilot Retrieval API overview (licensing, filter behavior). https://learn.microsoft.com/en-us/microsoft-365-copilot/extensibility/api-reference/retrieval-api-overview

49 OpenSearch: document-level security. https://docs.opensearch.org/latest/security/access-control/document-level-security/

50 Elasticsearch: setting up field and document level security. https://www.elastic.co/docs/deploy-manage/users-roles/cluster-or-deployment-auth/controlling-access-at-document-field-level

51 Amazon Builders' Library: making retries safe with idempotent APIs. https://aws.amazon.com/builders-library/making-retries-safe-with-idempotent-APIs/

52 Snowflake documentation: tag-based row access, open preview; policy argument and signature-alias matching. https://docs.snowflake.com/en/user-guide/tag-based-row-access-policies

53 Cursor documentation: project rules in .cursor/rules; guidance for the model, not a security control. https://cursor.com/docs/rules

54 Cursor documentation: security and privacy hardening (MCP allowlists, run modes, hooks, .cursorignore limits). https://cursor.com/docs/enterprise/security-hardening

ABOUT

Richard White

I was global firm chief data officer at McKinsey & Company, owning data, AI and platforms end to end, and before that chief data officer at New York Life, where I built the chief data office and reported data risk to risk committees at multiple levels. Earlier I led data at the Federal Reserve Bank of New York and enterprise data governance at Citigroup. I founded RJL Solutions in 2024. My book, *The Day After the Demo*, is free at rjlsol.com/book.

RJL Solutions connects strategy, architecture, engineering, governance and operation through Frame, Prove, Build, and Run & Transfer. Each phase is separately scoped and priced in a written proposal before you commit. The first sixty-minute working session is free.

IMPLEMENTATION **The Reference Implementation**
The full design these patterns fit into.

MEANING **Context Before Query**
Business meaning before the call; policy at the data.

EVIDENCE **AI Governance Evidence**
What reviewers and regulators need to inspect.

RJL SOLUTIONS

WEB rjlsol.com

EMAIL rwhite@rjlsol.com

Products are named as examples of patterns, not endorsements. No vendor reviewed or paid for this paper. Statements reflect vendor documentation checked in October 2026.