atlan

The Ultimate Guide to Evaluating a Data Catalog for the AI Era

The 5 steps to finding the right modern data catalog for your organization.

- **STEP 1** Define organizational needs
- **STEP 2** Create a customized evaluation criteria
- **STEP 3** Understand the data catalog vendors and offerings in the market
- **STEP 4** Take demos from select vendors
- **STEP 5** Execute hands-on proof of concepts (POCs)

INTRODUCTION

The AI Era Demands a New Data Foundation

Enterprise AI has a scale problem.

Despite massive investments in generative AI and machine learning, 95% of AI pilots fail in production. The culprit isn't the models themselves – it's the missing foundation beneath them.

AI systems are only as good as the context they understand. Without accurate, governed, and machine-readable metadata, AI confidently responds with the first answer it can find (without confirming the ask), confidently hallucinates, and erodes trust faster than it creates value. Models trained on data without lineage, quality signals, or business meaning can't distinguish between a customer's account ID and their Social Security number – or know which one is safe to use in production.

> ““Ironic, isn't it? That which used to be easy is now hard. And that which was once hard, is now easy,” said Joe DosSantos, Vice President of Enterprise Data and Analytics at Workday, during Re:Govern. “Why? Because our beautiful governed data, while great for humans, isn't really particularly digestible for an AI.”

The gap between AI promise and AI reality is due to missing context. Context fills the space between what humans know and what AI needs to be taught, clarifying business definitions, relationships, and hierarchies that are critical for accuracy and reliability.

Embedding context in data workflows and AI systems can’t be done with passive documentation – it requires a modern catalog that serves as an active metadata control plane for both human teams and AI systems. It unifies meaning, governance, and lineage across the data and AI estate, turning tribal knowledge into machine-readable context and transforming metadata from static documentation into active intelligence. That intelligence powers discovery, enforces policies, and enables AI agents to operate safely and autonomously at scale.

The shift toward active metadata is already underway: According to Gartner, **51% of organizations have already implemented metadata management, and 45% say it’s a top priority for the next two to three years**. More critically, Gartner estimates that by 2027, companies that broadly leverage metadata analytics will deliver new data assets up to **70% faster.**

Now, the question isn’t whether to invest in a modern data catalog – it’s how to choose the right one for an AI-first world.

This guide will walk you through a proven evaluation framework: from defining organizational needs and creating custom criteria, to understanding the vendor landscape, taking demos, and running hands-on POCs. By the end, you’ll have the tools to select a platform that doesn’t just catalog data, but activates context across your entire data and AI ecosystem.

WHY NOW?

No Context? No AI.

Enterprises are racing to operationalize AI, but most initiatives stall before production. The disconnect isn't due to technical or model weaknesses. It's due to lack of **context readiness.** AI models need more than raw data. They need:

- **Semantic understanding:** What does this data actually mean?
- **Lineage transparency:** Where did this data come from, and how has it transformed?
- **Quality signals:** Is this data accurate, complete, and fresh enough to trust?
- **Governance guardrails:** Can this data be used for this purpose, by this system, in this jurisdiction?

Without these trust signals, AI systems are unreliable, risky, and unsuccessful in delivering business value. But internal research from Atlan AI Labs found that **metadata improved query response accuracy by 5x** – indicating that context is the difference between an AI system that works and one that fails.

Context Infrastructure for Enterprise AI

A modern data catalog should provide a context layer that bridges raw data and trustworthy AI. It does three things traditional catalogs cannot:

Unifies meaning, governance, and lineage across disconnected systems into a single, queryable metadata graph.

Activates metadata by writing it back to source systems, propagating policies, and enforcing controls where data lives.

Powers AI agents and applications by exposing governed, real-time context through APIs, MCP servers, and natural-language interfaces.

This **metadata control plane** continuously ingests, enriches, and distributes context to both humans and machines, with embedded feedback loops for ongoing improvement. For data leaders, this is a move from reactive cataloging to proactive governance. For AI systems, it unlocks access to the structured knowledge required to operate reliably and compliantly at scale.

What This Means for Your Evaluation

As you evaluate platforms, prioritize those that:

- Treat metadata as active infrastructure, not static documentation
- Support policy push-down and enforcement at source systems
- Provide real-time, governed context for AI agents via APIs and frameworks like MCP (Model Context Protocol)
- Demonstrate measurable improvements in AI accuracy, data time-to-trust, and governance automation

The rest of this guide will help you identify these capabilities and run a rigorous evaluation that sets your organization up for AI success.

STEP 1

Define organizational needs for a data catalog

1. Start by identifying the top three challenges that are causing data and AI initiatives to fail at your organization.

Conduct surveys and interviews across your organization to identify the top pains and failure points for data and AI projects.

Common Pain Points:

- **Data data everywhere – just not when you need it:** Long waiting periods for business users to access data lead to productivity losses.
- **Data is a black box:** Lack of transparency into data quality, lineage, and meaning creates trust issues.
- **Human tribal knowledge is siloed and often lost:** Context around data assets lives in people's heads and Slack threads, not systems.
- **Metadata creation is a second job:** Manual processes for documentation, classification, and policy enforcement don't scale.
- **What will this break?:** Unclear data lineage and impact analysis make changes risky – and concerning.
- **Trust breaks in production:** Data quality issues surface in production, eroding trust in analytics and AI.

Sample Survey Questionnaire

PART 1: Role & Priorities

CATEGORY	QUESTION	OPTIONS (TO BE CUSTOMIZED)
Role and User Persona	Which of these personas best fits your role?	Data Scientist, Data Analyst, Data Engineer, Business Analyst, Business Manager, Data Steward, IT Support, Cloud Team, Other
	What are your top three priorities this year?	[Insert text]
	What are your top three challenges that affect productivity? How many hours would you save per week by solving each?	[Insert text]

PART 2: Data & AI Activities

CATEGORY	QUESTION	OPTIONS (TO BE CUSTOMIZED)
Data-Related Activities	Which activities do you perform? How much time per week?	Getting access to data, Finding the right dataset, Understanding context (column meanings, lineage), Running quality checks, Exploring data & running queries, Preparing data for AI/ML
	Which activities create productivity challenges? How many hours would you save?	Same list as above
AI-Related Activities	Do you build, deploy, or govern AI systems? What challenges do you face?	Context availability for models, Data quality for training, Policy enforcement for AI, Lineage from data to model to output, Agent access to governed data

PART 3: Collaboration & Governance

CATEGORY	QUESTION	OPTIONS (TO BE CUSTOMIZED)
Collaboration	Who do you collaborate with on data weekly? Do you face challenges, and how much time would you save by solving them?	Data Scientist, Data Analyst, Data Engineer, Business Analyst, Business Manager, Data Steward, IT Support, Cloud Team, Other
Governance & AI Governance	Do you manage policies, access controls, or AI governance? What are your biggest challenges?	Manual policy enforcement, Lack of visibility into AI data usage, Policy drift across systems, Inability to audit AI access to sensitive data

2. Map organizational needs and challenges with the core capabilities of modern data catalogs.

Understanding modern platform capabilities is critical. Modern data catalogs go beyond traditional catalog functions, serving as the metadata control plane for data and AI governance. **Core capabilities:**

DISCOVERY

Clear and comprehensive view of all data assets within the organization, enhanced by AI-powered search, automated classification, and natural-language querying

KNOWLEDGE

All context, information, and business know-how around data assets, automated through AI enrichment, living documentation, and embedded collaboration

TRUST

Information on data quality, coverage, and usage, supplemented by AI readiness scores, automated quality monitoring, and lineage at the query/code level

COLLABORATION

Intuitive UI for diverse team members to effectively collaborate on data assets, plus AI co-pilots that answer questions in natural language

GOVERNANCE

Access rights and policies to ensure legal and regulatory compliance, with policy push-down, automated enforcement, and cross-system governance

SECURITY & AI GOVERNANCE

Ensure secure and compliant usage of data assets and AI systems, with AI readiness scorecards, model lineage, and agent access controls

Map capabilities to key organizational needs

Example Mapping:

PRIORITY	TOP ORG NEEDS/ PROBLEM STATEMENTS	PRIORITY PLATFORM CAPABILITIES
1.	Long waiting periods for data access lead to productivity losses.	**Discovery:** AI-powered search across attached metadata, glossary, classification, and lineage. **Governance:** Automated and customizable access policies, self-service provisioning.
2.	Lack of trust in data quality and accuracy.	**Trust:** Automated data quality scoring and alerts, AI readiness scorecards. **Knowledge:** Automated lineage at the query/code level, change impact analysis.
3.	No clear context around data assets; multiple versions create confusion.	**Knowledge:** AI-assisted documentation, automated tagging, living data dictionaries. **Collaboration:** Embedded Slack/Teams integrations, version control.
4.	AI systems lack context and produce unreliable outputs.	**AI Governance:** Context layer for AI agents via Metadata Lakehouse and MCP Server. **Trust:** Real-time metadata for RAG systems, governed data contracts for AI.
5.	Manual governance processes don't scale	**Governance:** Policy push-down and enforcement at source systems. **AI Governance:** Automated AI readiness scoring and compliance monitoring.

3. Evaluate non-functional considerations related to modern data catalogs.

Every organization is unique in terms of its tech stack, user population, maturity, and investment. That's why it's important to look at non-functional aspects of a platform – and determine how relevant they are to your team's needs. **Key considerations:**

Ability to integrate with other tools & technology

Standalone platforms that don't integrate become "just another tool." **Look for:**

- **Native integrations with BI tools** (Tableau, Power BI, Looker), notebooks (Jupyter,Databricks), orchestration (Airflow, dbt), and AI frameworks (LangChain, LlamaIndex).
- **Bidirectional metadata exchange:** Not just reading metadata, but writing it back to source systems.
- **Embedded experiences:** Metadata surfaces directly in the tools data teams already use.
- **MCP Server support:** Enables AI agents to query and act on governed metadata in real time.

Built to prevent technology lock-in

Modern catalogs with active metadata are foundational infrastructure. To maintain flexibility, **prioritize:**

- **Open standards:** Support for OpenLineage, OpenMetadata API, and industry-
- standard protocols.
- Metadata Lakehouse **architecture:** Stores metadata in an open, queryable format
- (e.g., Apache Iceberg) that you control.
- **API-first design:** Comprehensive REST and GraphQL APIs for programmatic access.
- **Export capabilities:** Ability to extract your metadata in portable formats if you ever need to migrate.

Interface optimized for business, not just IT

Metadata management is no longer just for IT teams. Business stakeholders actively use data and AI systems, so they need intuitive, accessible interfaces. **This includes:**

- **Low-code environment:** Business users can create policies, tag data, and manage glossaries without SQL or Python.
- **Natural-language interfaces:** AI co-pilots let users talk to data and ask questions, with no coding experience required.
- **Visual query builders:** Drag-and-drop interfaces for data exploration.
- **Design-first approach:** Consumer-grade UX, not enterprise software from 2010.

Ease of onboarding and setup

Traditional platforms require full-time engineering teams to set up, support, and maintain, not to mention paid consultants and months of delay. **Instead, look for:**

- **Self-serve setup:** Clear onboarding flows that don't require professional services.
- **Guided implementation:** Persona-based user journeys and in-app guidance.
- **Initial value in 2-6 weeks:** Not 6-12 months.
- **No hidden costs:** Transparent engineering time requirements and minimal ongoing maintenance.

Pay-as-you-go models

Pricing should scale with adoption, not hopeful commitments. **Look for:**

- **Low initial investment:** No massive upfront licensing fees that require executive approvals before you prove value.
- **Scale costs as adoption occurs:** Costs tied to usage (e.g., number of active users, data sources, or API calls) so pricing grows with value realization.
- **Transparent pricing:** No surprise fees for connectors, storage, or support tiers.

STEP 2

Create a customized evaluation criteria

Finalize a clear set of objective criteria that will guide the evaluation process:

✓ Core capabilities mapped to organizational needs
✓ Other high-priority considerations and how to evaluate them

Sample Evaluation Criteria – Core Functionalities

CATEGORY	FEATURE REQUIREMENT	PRIORITY
Discovery	AI-powered master search across metadata, glossary terms, classifications, lineage, usage, and more.	1
	Natural-language query interface (e.g., "Show me all PII tables updated in the last 30 days")	1
Knowledge	Automated documentation and enrichment via AI (descriptions, owners, tags, glossary links) with high expert acceptance rates (70%+)	1
	Comprehensive data dictionary with automated profiling, sampling, and statistics	1
	Query-level and code-level lineage: Track data transformations through SQL, Python, Spark, dbt, and more	1
	Automated impact analysis: Identify downstream dependencies before making changes	2
	Version control and change history for data assets and metadata	2
Trust	Automated data quality rules with quality scoring and monitoring	2
	AI readiness scorecards: Configurable metrics (representativeness, quality, risk) combining internal signals and external inputs	1
	Real-time data quality alerts and incident management workflows	2

CATEGORY	FEATURE REQUIREMENT	PRIORITY
Collaboration	Embedded collaboration: Comments, @mentions, and discussions directly on data assets	2
	Integration with Slack/Teams for notifications and workflows	2
	Saved queries and shared collections for team knowledge sharing	2
Governance	Policy push-down and enforcement: Active-metadata policies (e.g., PII tags) propagate to source systems and trigger masking/controls via open APIs	1
	Customizable access policies at table, row, column, and attribute levels	1
	Real-time monitoring of data access with audit logs	2
	Self-service access request workflows with approval routing	2
	Data contracts and data products: Define and enforce SLAs, ownership, and quality expectations	2
AI Governance	Control plane for AI agents: Provide real-time, governed context via Metadata Lakehouse and MCP Server for agentic AI, RAG, and enterprise apps	1
	AI readiness scoring across data, models, and applications for continuous governance	1
	Model lineage: Track data → feature → model → inference relationships	2
	Policy enforcement for AI systems: Ensure models and agents comply with data usage policies	1
Security	Integration with Active Directory, SSO, and RBAC systems	1
	Automated classification and tagging of sensitive data (PII, PCI, PHI, etc.)	1
	Automated compliance reporting (GDPR, CCPA, SOC 2, etc.)	2

CATEGORY	FEATURE REQUIREMENT	PRIORITY
Extensibility	Open API and SDK for custom integrations and automations	1
	App framework: Build custom metadata applications on top of the platform	1
	Webhook support for event-driven workflows	2

Sample Evaluation Criteria – Other Considerations

CATEGORY	KEY CONSIDERATIONS	HOW TO EVALUATE?
Ability to integrate with every user's favorite tools	**Test integrations with:** Jupyter, Microsoft Excel, Power BI, Tableau, Looker, Databricks, Airflow, dbt, Snowflake, Fivetran, LangChain, etc.	**Level 1:** Ask vendor if native integrations are available **Level 2:** Ask vendor to showcase integrations **Level 3:** Test integrations deeply in POC
Open architecture that avoids technology lock-in	Ensure platform uses open standards (OpenLineage, OpenMetadata), provides Metadata Lakehouse with portable storage (Iceberg), and offers comprehensive APIs	Review architecture documentation Ask about metadata export capabilities Evaluate API completeness
UI optimized for business, not just IT	Ensure catalog is easy for non-technical business users, with features like natural-language search, visual query builders, and AI co-pilots	**Level 1:** Include business users in evaluation and rate ease of use on a 1-5 scale **Level 2:** Include business users in POC and observe usage without support
Onboarding effort and time	Evaluate human hours needed for setup and time to initial value	**Level 1:** Ask for detailed setup documentation with time estimates **Level 2:** Track actual time spent and ad-hoc requests during POC

CATEGORY	FEATURE REQUIREMENT	PRIORITY
Product pricing that supports success	Evaluate flexibility and scalability of pricing for current and future usage	**Request detailed pricing:** upfront cost, POC cost, scaled-up cost, additional licenses, infrastructure requirements
Active metadata & policy enforcement	Evaluate whether the platform can write metadata back to source systems and enforce policies at the point of data access	Ask for demonstrations of policy push-down Test metadata write-back during POC Validate enforcement in tools like Snowflake, Databricks, etc.
Metadata Lakehouse and interoperability	Ensure metadata is stored in an open, queryable, and controllable format, and is accessible via SQL	Request architecture details on storage layer Ask about query capabilities and export options Validate portability claims
AI governance & agent support	Confirm that AI agents and applications can access governed context in real time via APIs or MCP Server	Test API response times and data freshness Validate MCP Server integration if relevant Ask for agent use case examples Test with a set of enterprise "golden questions"

STEP 3

Understand the vendors and offerings in the market

There are broadly three types of related offerings in the market, with a fourth emerging that's built for AI:

	WHO USES IT?	WHAT DO THEY USE IT FOR?	WHAT DATA DOES IT WORK ON?
Technical account	Uses of the technical tool (eg: AWS, Azure, Snowflake, Databricks)	• Discover data • Understand data	Data in that technical tool
Embedded catalog	Users of the data tool (eg: dbt, Tableau)	• Understand data • Trust data	Data connected to that tool
Universal catalog	Users of the catalog	• Discover data • Understand data • Trust data	Data connected to that catalog
The Future: Data and AI Control Plane	Users of any tool	• Discover data (catalog, marketplace) • Understand data (metadata, lineage) • Trust data (alerts, data products) • Govern data • Access & security • Quality & observability	Data or operations in any tool

1. Technical Catalogs

Examples:
AWS Glue, Azure Purview, Snowflake Polaris, Databricks Unity Catalog

STRENGTHS	LIMITATIONS
Solve for context by exposing metadata about data assets, products, lineage, tags, policies, etc. for a single data source Enable data discovery & understanding Some provide governance capabilities, such as access control enforcement	Tools are tightly coupled to respective platforms, increasing tool overhead and consistency gaps Limited cross-system lineage and policy propagation Built for technical users, not business users No built-in features for AI readiness, model lineage, or agent governance No collaboration or knowledge management layer Not designed for federated governance

2. Embedded Catalogs

Examples:
dbt Explorer, Tableau Catalog, Snowflake Horizon

STRENGTHS	LIMITATIONS
Context is surfaced within the data tool, so users don't have to open a separate tab Native integrations and real-time syncs make setup and updates easy Platform-specific governance and domain-focused lineage improve lineage and management Cost-effective and performant for single-platform teams	Views, governance, search, policies, and lineage don't scale across systems Built for technical users, not business users Limited collaboration features like discussions and knowledge sharing Potential for vendor lock-in and multiple sources of truth Weak support for federated architectures and AI governance

3. Universal Catalogs

Examples:
Alation, Collibra

STRENGTHS	LIMITATIONS
Cross-platform metadata aggregation from diverse sources Single pane of glass for search and discovery Centralized governance policies and business glossaries Comprehensive asset coverage across the data estate, with strong audit and compliance capabilities Cross-system lineage Dedicated UI for data discovery	Static, passive metadata with no write-back or enforcement Built primarily for technical users Disconnected from users workflows Substantial cost overhead and implementation timeline Limited support for AI systems, policy enforcement, and federated governance Collects but doesn't activate metadata

4. Data and AI Control Plane

Examples:
Atlan

STRENGTHS	LIMITATIONS
Cross-platform metadata aggregation with bidirectional metadata activation AI-first features, including Atlan AI, AI Governance Studio, Talk-to-data co-pilot Built for technical and business users, with context embedded everywhere users work Open and interoperable via Metadata Lakehouse architecture, open APIs, standards-based integrations	Newer in market compared to legacy vendors (though rapidly growing customer base) Limited support for on-premises data sources (optimized for cloud-first organizations)

STRENGTHS	LIMITATIONS
Extensible, with app framework for custom metadata applications Fast time-to-value: Initial value in 2-6 weeks, depending on scope Automated policy push-down and real-time monitoring for simplified compliance Real-time context for AI agents via APIs and MCP Server	

STEP 4

Take demos from modern data catalog vendors

After conducting secondary research on market offerings, reach out to select vendors to set up deep-dive demos.

Best practices to set up successful demos

1. Share finalized evaluation criteria document with vendors in advance

Why:

- Helps vendors understand your problem statements and organizational context upfront
- Serves as an alignment document to ensure all questions and priorities are covered during the demo
- Enables vendors to prepare relevant examples and use cases

2. Ensure stakeholders from different teams attend the product demos

Why:

- Brings together different user personas to gather the most comprehensive feedback possible and ensure needs are met across the board
- Ensures buy-in across the organization and surfaces diverse use cases

Who should attend:

- Data leaders (CDO, Head of Data Governance, Head of Data Engineering)
- Practitioners (data engineers, analysts, scientists)
- Business stakeholders (business analysts, product managers)
- IT/Security (if governance and compliance are critical)
- AI/ML teams (if AI governance is a priority)

3. Conduct a data architecture compatibility check

Why:

- Ensures that the platform is compatible with both your current data architecture and your vision/roadmap for the next 1-3 years

What to validate:

- Does the platform support your current data sources (cloud warehouses, lakes, databases, BI tools, orchestration, etc.)?
- Can it scale to support planned additions? (new data sources, AI systems, data products)
- Does it align with your target architecture (e.g., data mesh, lakehouse, federated governance)?

STEP 5

Execute a hands-on proof of concept (POC)

After taking demos, align with your internal stakeholders to narrow down your options. Then, reach out to selected vendors to set up hands-on POCs.

L1: PRE-POC ACTIONS

✓ CONDUCT A DESIGN SPRINT

- Identify top use cases and user flows to test (but don't boil the ocean by going too broad)
- Ensure key capabilities in the evaluation criteria are included
- Prioritize use cases that will demonstrate value to stakeholders

✓ SET UP POC ORG STRUCTURE

- Prepare your ideal tech architecture: sources, connectors, data samples
- Properly onboard business users who will be involved in POC execution
- Designate a POC lead and core team with clear responsibilities

L2: DURING POC ACTIONS

✓ ENSURE DEDICATED TIME SPENT

- 80% of the core team's time should be spent on execution
- Block calendars to avoid distractions and delays
- Capture detailed feedback at each step of POC

✓ STAGGERED IMPLEMENTATION

- Do not aim to conduct POCs with more than 2 vendors
- Do not run multiple POCs simultaneously – stagger them to ensure focus and compare results fairly

✓ HIGH-TOUCH COMMUNICATIONS

- Use a dedicated Slack or Microsoft Teams channel for continuous communication
- Adopt a partner mindset: create continuous feedback loops with the vendor team
- Success depends on people as much as technology – manage the POC team well

L3: POST-POC ACTIONS

✓ DEEP-DIVE FEEDBACK

- Set up calls with vendor teams to share detailed feedback on the POC experience and gauge response to feedback
- Evaluate how vendors respond to criticism – are they defensive or collaborative?

✓ QUICK SPRINT TO DECISION

- Bring together the core team and other decision-makers to make a quick decision on the platform
- Prevent lag and leverage current momentum – delays kill adoption

Create a foundation for a successful POC

High-touch communications through Slack or Microsoft Teams: Continuous feedback loops are critical. Treat the vendor as a partner, not just a product provider.

Adopt the partner mindset: Success depends on collaboration. Create open channels for questions, feedback, and iteration.

Success depends on people as much as technology: Manage the POC team well by ensuring they have the time, tools, and clarity to execute effectively.

Finally — choose a partner, not a vendor.

As you finalize your decision, consider:

- How do the vendors respond to feedback? Are they listening or only trying to sell?
- Are they upfront and honest?
- Do they share their limitations with you while acknowledging positives in other competitors?

The best partnerships are built on transparency, collaboration, and mutual success. Choose a vendor that will grow with you, not one that locks you in.

THANK YOU FOR READING THE ULTIMATE GUIDE TO

Evaluating a Modern Data Catalog for the AI Era

Check out our other resources for data teams:

EBOOK

The Data Catalog Primer – Everything you need to know to set up one

Download →

EBOOK

The Ultimate Guide to Building a Business Case for an Active Metadata Platform

Download →

EXPLORE ATLAN

The Active Metadata Platform Built for the AI Era

Atlan is recognized as **a Leader** in the 2025 Gartner® Magic Quadrant™ for Metadata Management Solutions.

Trusted by data teams around the world

Atlan powers data and AI governance for leading organizations across industries, from Fortune 500 companies to fast-growing startups.

Medtronic workday. mastercard Nasdaq CME Group Unilever

Ready to see Atlan in action?

Talk to Sales — See how Atlan accelerates data and AI governance for your organization.

Product Tour — Explore Atlan's features at your own pace.

atlan