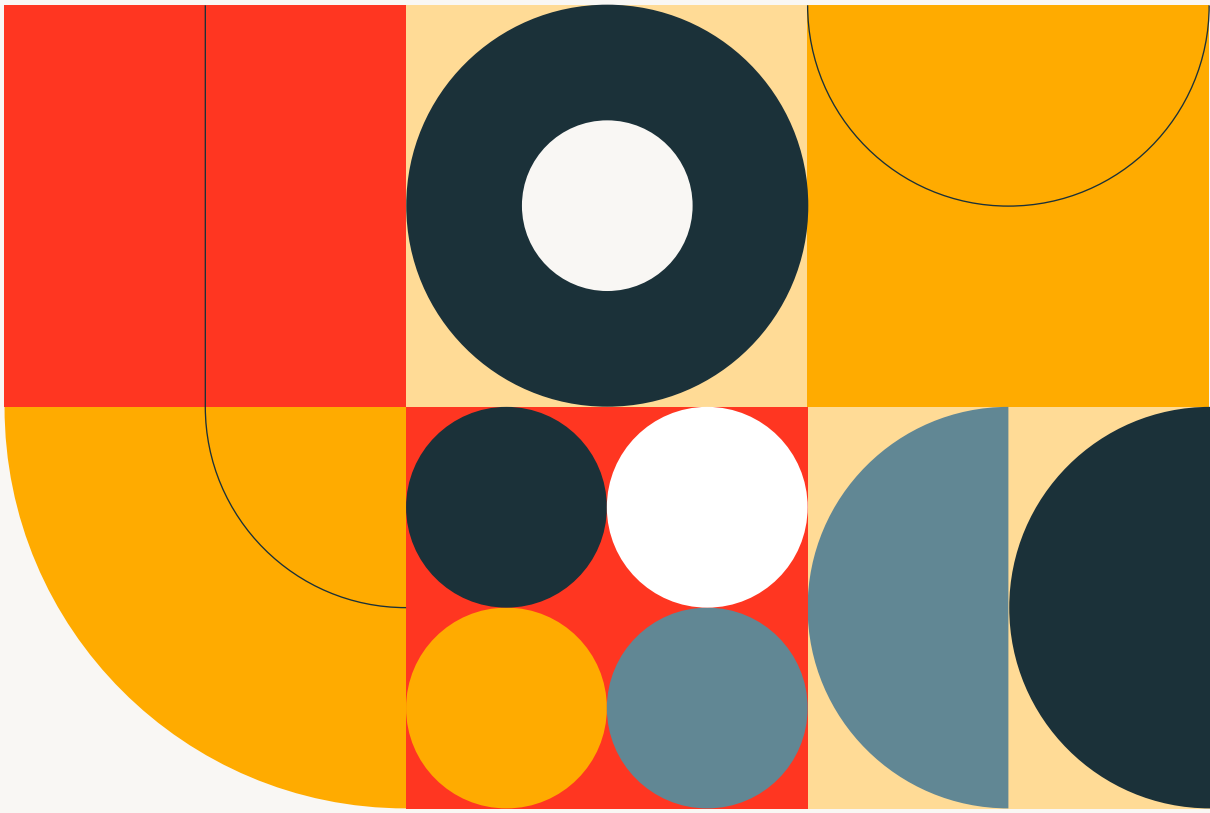

The Official Value Guide for Partners

Data Intelligence Platform

FY 2026

GTM METHODOLOGY

databricks

Welcome to the **Databricks Value Guide for Partners**

At Databricks, our partners are a vital extension of our go-to-market engine. For the first time ever, we are sharing how our own internal field teams learn & position the Databricks platform. Built on ***customer obsession***, these guides are designed to help you advise with impact, influence decisions, and drive real outcomes.

You'll find a simple framework woven throughout to help answer key questions:

1. 

   What problems is the customer trying to solve?

2. 

   What will success look like?

3. 

   Why is Databricks the right partner?

4. 

   Where have we done it before?

Together, let's raise the bar for our customers.

The Value Based Conversation

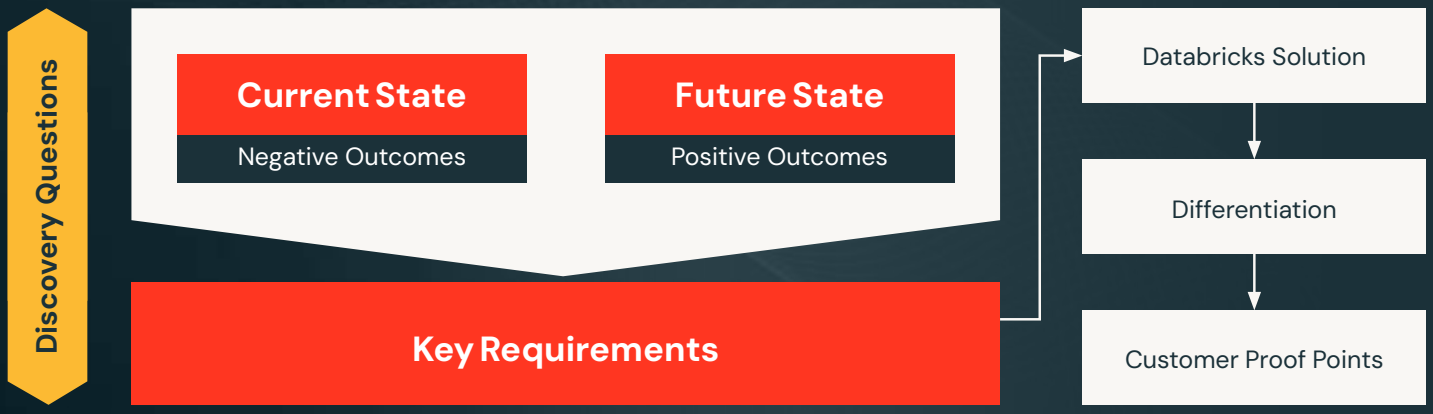

Customer Value Drivers

What's top of mind for our customers

1 Reduce Cost & Complexity

Unification allows all Data & AI use cases to be performed directly on one copy of the data with up to 8x better price/performance. Eliminate duplicative costs in tools, storage, redundant processes, etc.

A unified approach to governance mitigates risk across all Data & AI, bolstering business resilience & fostering customer trust.

AI-powered capabilities within the platform simplify maintenance / reduce costs even further.

2 Accelerate AI-Driven Innovation

CxOs are looking to drive business transformation with Data & AI; every company must accelerate AI-driven innovation or risk being disrupted by faster-moving competitors.

Building on a unified platform breaks down silos, allowing the right teams to collaborate on the right data in real-time, accelerating projects from experimentation to production.

Data & AI are democratized to non-technical users via natural language enabling greater productivity across the org.

3 Mitigate Risk Across Data & AI

As data grows in value and AI adoption accelerates, so do the risks. Customers face rising concerns around privacy, compliance, and the safe use of AI.

Databricks mitigates these risks by embedding security, governance, and responsible AI into every layer of our unified Data and AI platform.

With built-in access controls, encryption, monitoring, and auditability, customers can protect sensitive data, meet regulatory requirements, and ensure trustworthy AI—without slowing innovation.

The Discovery Conversation

What to look for & listen for

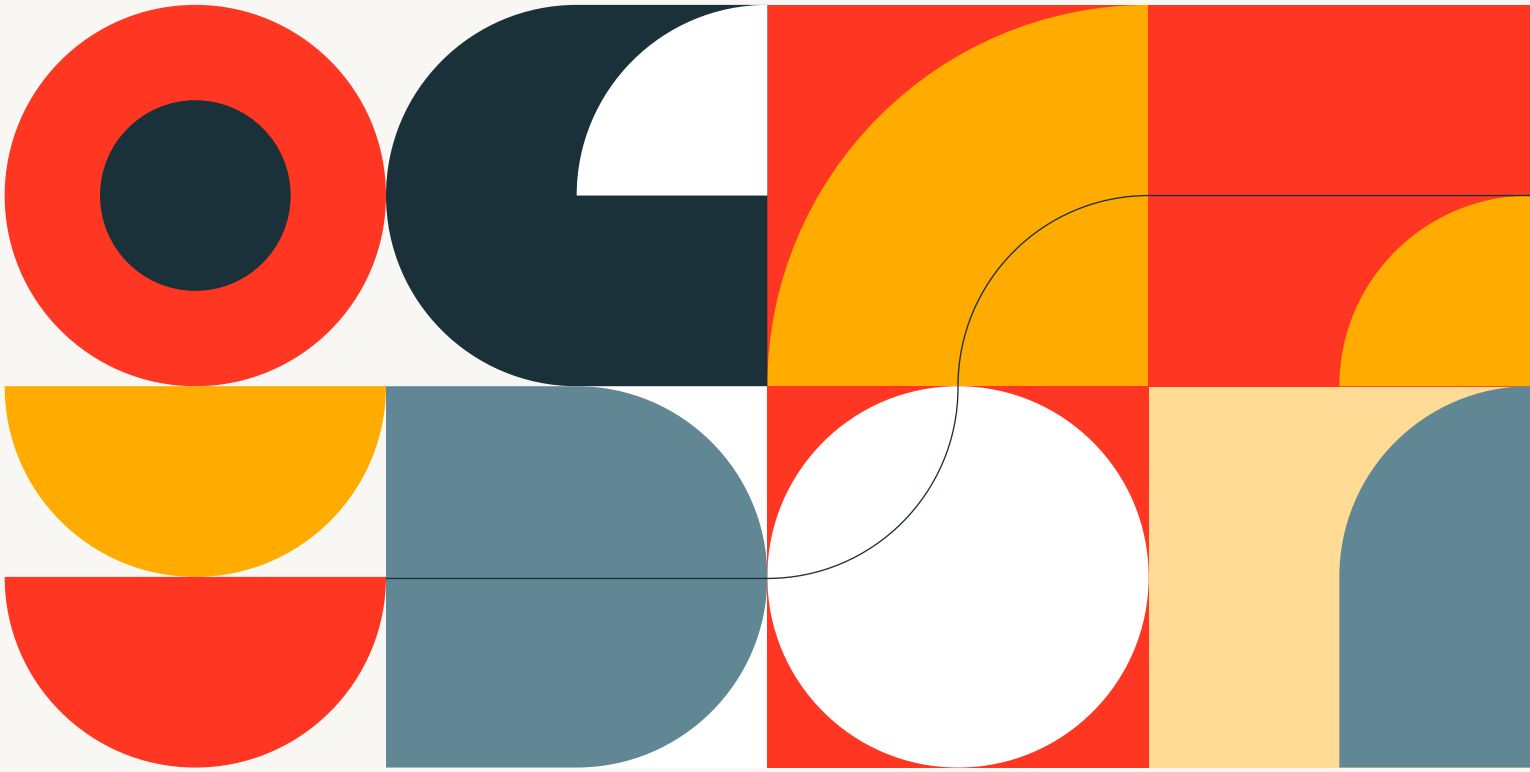

Current State	Negative Outcomes
Multiple data platforms are deployed across different workloads (data lake, warehouse, ETL, data science, AI & BI).	The environment is **unnecessarily complex** leading to **runaway costs** at scale and **longer time to value.**
Governance and lineage are fragmented across many systems.	**Security vulnerabilities and risks** in how data is used / changed throughout its lifecycle.
Data is stored in proprietary formats.	Proprietary investments lead to **slower innovation cycles** and **lock-in-related cost pressures.**
Small number of technical staff to produce data products.	Technical bar too high for most employees; **limited access to Data & AI.**
Difficult to build high quality AI agents in production.	Limited integration between the Data platform and AI platform leads to **development complexity** & **production risk.**

Future State	Positive Outcomes
The Data & AI platform maintains a single copy of data.	**Lower Cost:** Data assets available in real-time; workloads scale without runaway costs; increased production efficiency
Enterprises have a single view of governance.	**Lower Risk:** Perfect upstream / downstream lineage to reduce risk from changes / misuse. **Lower Cost:** Single source of truth that persists across workloads.
Data is stored in open formats; easily portable across platforms and clouds.	**More Innovation:** Easily integrate with tools, adopt new tech tomorrow, and reliably share data with stakeholders. **Lower Cost:** Maintain negotiation power over platform decisions.
The Data & AI platform inherently understands semantics / usage patterns of the enterprise's data.	**More Innovation:** All employees can create value from Data & AI. **Lower Cost:** AI-powered optimization improves data discovery / quality / platform performance.
The Data platform & AI platform are intrinsically linked & support high quality AI agents.	**More Innovation:** Productionize AI efforts at scale; massive efficiencies in building & serving models. **Lower Risk:** AI agents that produce domain specific and accurate outputs; leading to lower risk and higher trust in the application.

Remember GOAT

Influence the customer conversation towards our differentiated value proposition

Key Requirements	How We Do It
G — UNIFIED GOVERNANCE	Unity Catalog, Lakehouse Federation, Lakehouse Monitoring
O — OPEN SOURCE & OPEN STANDARDS	Delta Lake UniForm, Apache Spark, MLflow Delta Sharing, Apache Iceburg, Unity Catalog
A — AI POWERED	AI/BI Genie, AI/BI Dashboards, Databricks Assistant, Predictive Optimization
A — AI AGENTS	Mosaic AI, Unity Catalog, Lakehouse Monitoring
T — EFFICIENT TCO	Databricks SQL, Serverless, Photon Lakeflow, SAP

Discovery Questions

What to ask

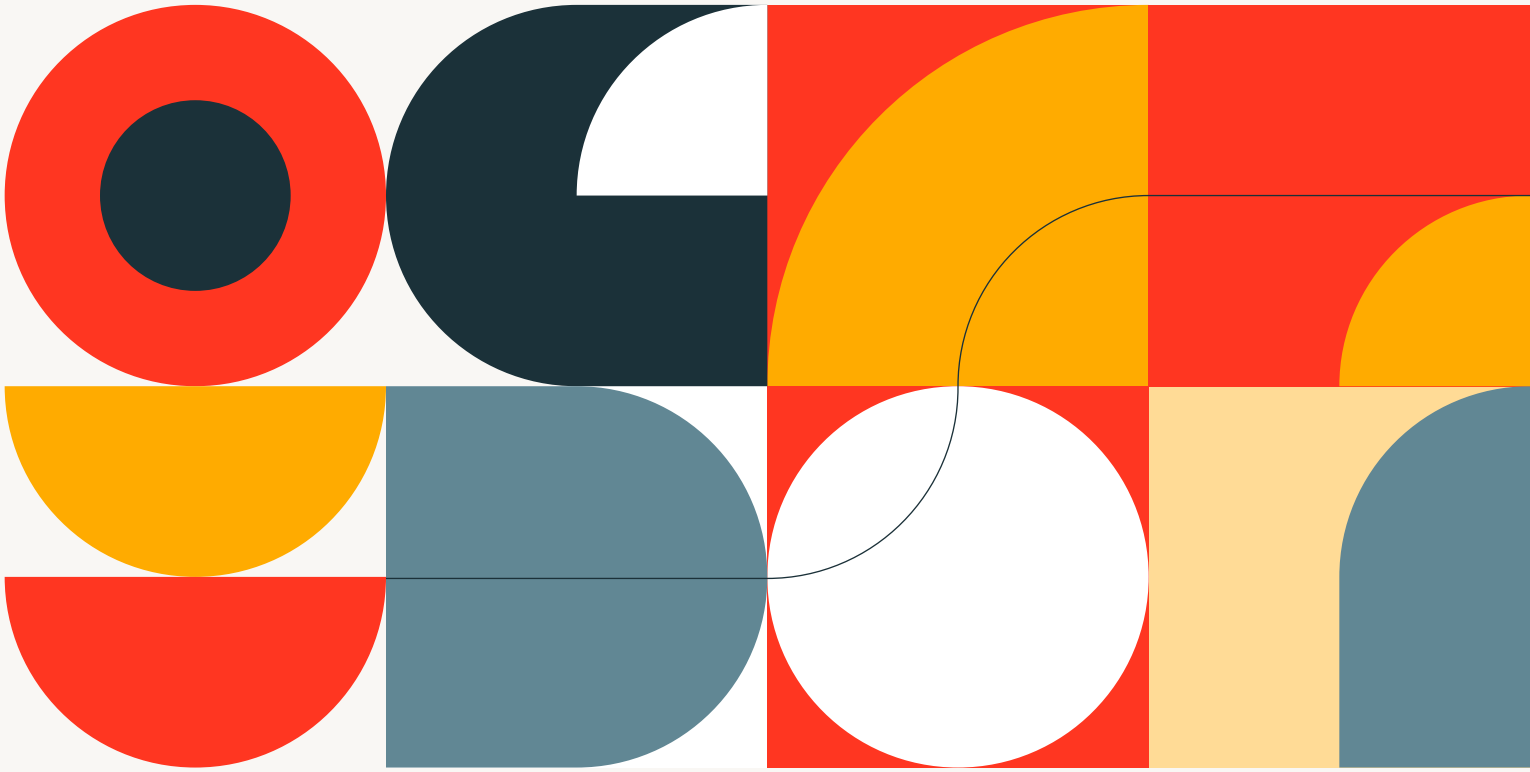

Broad Discovery

Focus on attaching to business value

Data & AI Innovation

What positive business outcomes are you looking to unlock?

Where are you in your journey of becoming a data-driven company (data maturity curve)?

What are the key components of your data strategy? How are you accelerating time to value? Biggest bottlenecks?

Team Productivity

Across the various data teams (DE, DA, DS) what collaborative challenges are you experiencing?

How much time are your teams spending on maintenance vs new initiatives? Goals to increase efficiency?

What are the key components of your data strategy? How are you accelerating time to value? Biggest bottlenecks?

In today's highly competitive market, how are you attracting / retaining top talent?

Overall Complexity

Describe your cloud strategy? Do you have data in the cloud currently? Migration priorities & timelines?

What challenges are you experiencing with the technologies you're using today (ex: Hadoop, legacy DW, EMR, etc)?

What challenges do you have managing various data sources and types? Downstream analytics (BI, SQL, ML)? Data teams / stakeholders?

What benefits do you see if all your teams operated on one copy of data across multiple use cases?

How does real-time processing support specific business cases and what SLA's are required?

Operational Cost & TCO

Many of our customers are experiencing cost pressures as data volumes explode; where are the biggest costs across your data workloads?

How much reduction in cloud costs are you targeting? Where would you reallocate those funds?

Operational Cost & TCO

How much would consolidating to a single data platform across multiple clouds reduce your cost?

How valuable would it be to run your analytics directly on your data lake (reduce dependency on your DW)?

Governance & Risk

Many of our customers experience higher cost / complexity / risk from siloed governance solutions; have you experienced similar challenges?

From a security standpoint, what are your biggest concerns or requirements?
What could happen if you don't meet those requirements?

How are you managing security / governance across your data lakes, DWs, streaming / ML systems in the cloud?

How do you currently collaborate and share data, analytics and AI with your partners?
Is the current solution open or does it lock participants in?

How would a common governance model on one copy of data give you more confidence in your security / governance posture?

Open Source & Sharing / Collaboration

How important is avoiding lock-in or integrating with a broad ecosystem of tools?

What types of data does your org want to share internally?
Acquire externally (from data marketplaces, vendors, apps, and/or partners)?

What current use cases benefit from sharing and collaboration at your org?

GenAI

From your perspective (CTO/CDO/CIO), what strategic business outcomes are you trying to drive with GenAI? Any concerns or caution?

How do you see AI informing strategic decisions? Creating new product / revenue growth opportunities?

What exciting new use cases for AI and GenAI have you identified? Any LLM use cases you think are interesting / promising?

Do you have an ML engineering team? Have you already tried to build a GenAI solution (with OpenAI, SageMaker, Open Source models)?

Do you use your enterprise data as part of the GenAI solution? What concerns do you have?

Databricks Solution

Articulate our value

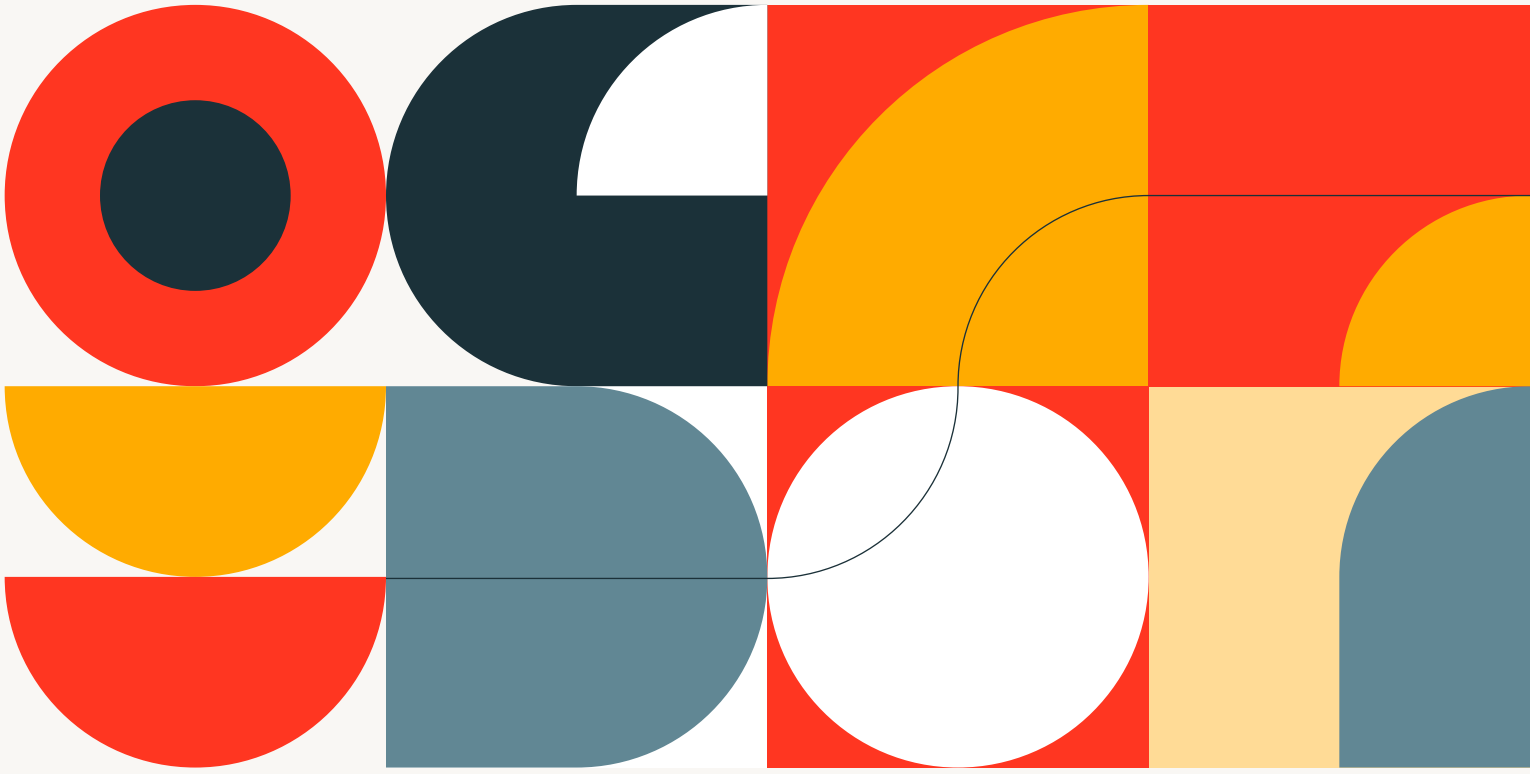

 | GTM Methodology | nabeel.ali@databricks.com

Databricks Data Intelligence Platform

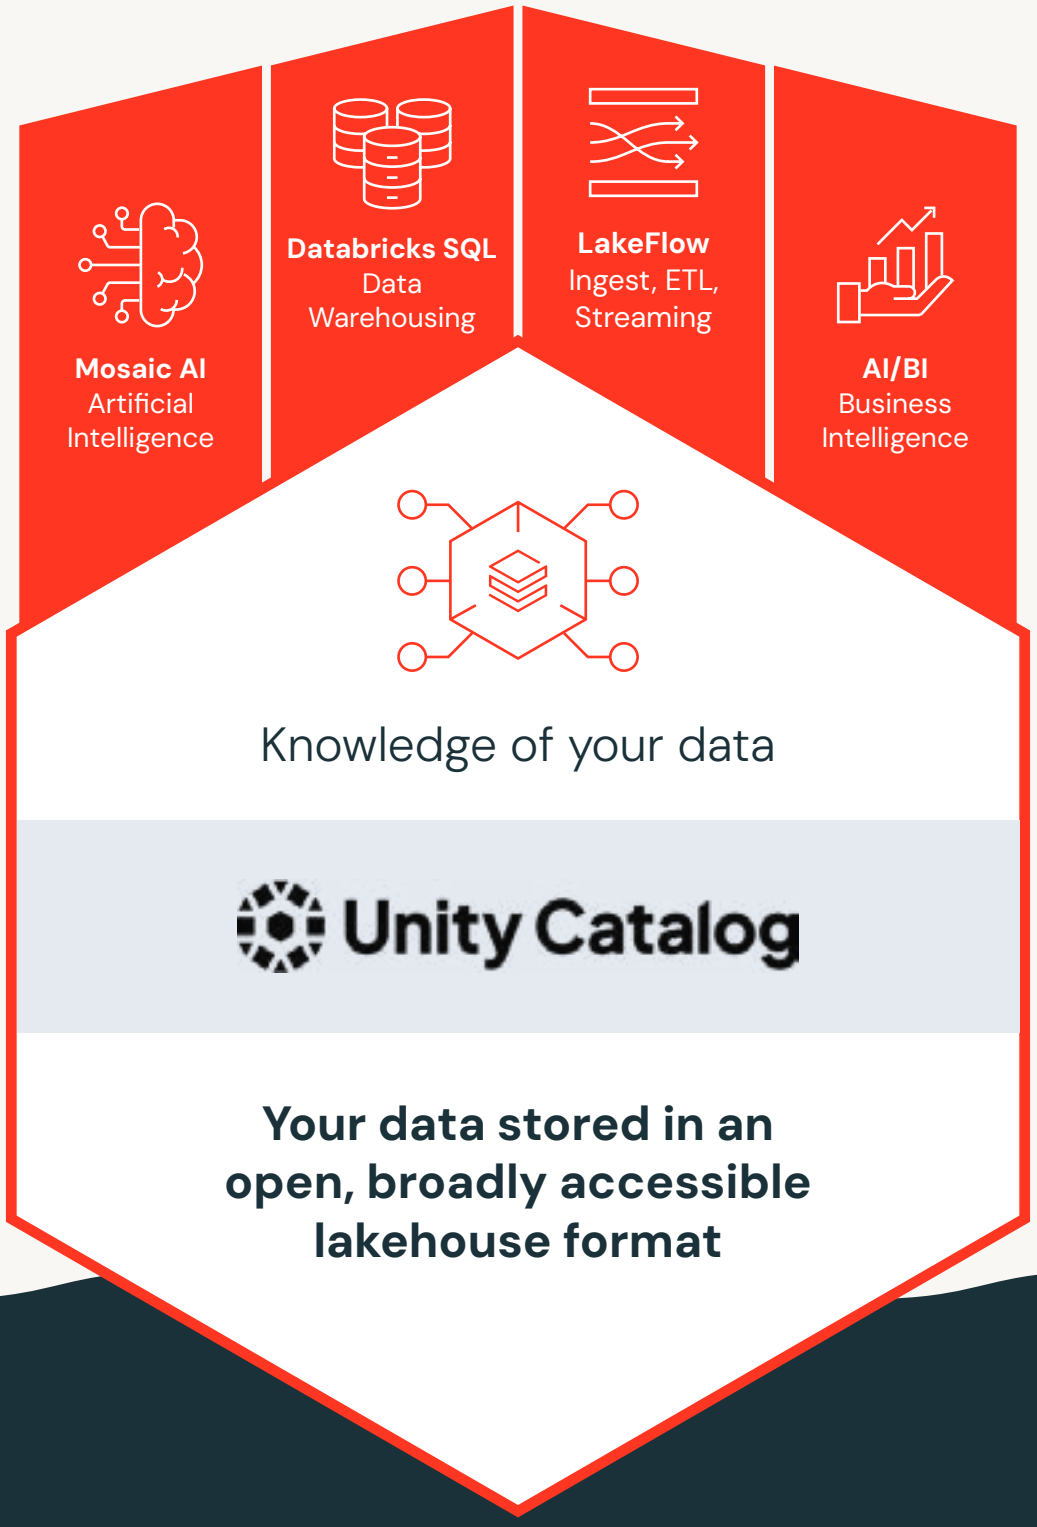

How We Do It

The Databricks Data Intelligence Platform unifies Data & AI on one open, governed, and scalable platform.

1. **It starts with an open Lakehouse architecture**
   Customers keep a single copy of data for all workloads across clouds and teams.

2. **At the core is Unity Catalog**
   The governance layer for data and AI. It provides built-in security, discovery, and lineage across data, models, and code. Now open source and Iceberg-native, it supports today's open data standards.

3. **With AI applied to the platform, we unlock "Data Intelligence"**
   AI that understands your data and enhances everything you do with it.

On top of this foundation, the platform supports all major Data & AI workloads:

Mosaic AI helps teams build and deploy AI agents with production-grade governance

Databricks SQL delivers fast, serverless analytics and interactive dashboards

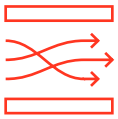

LakeFlow simplifies ingest, ETL, and streaming pipelines—with native SAP support

And customers can bring their own **AI or BI tools** to unlock insights across the business

We partner with customers to ensure rapid success through industry solution accelerators, migration support, customer training, and more.

Customer Proof Points

Where we've done it before

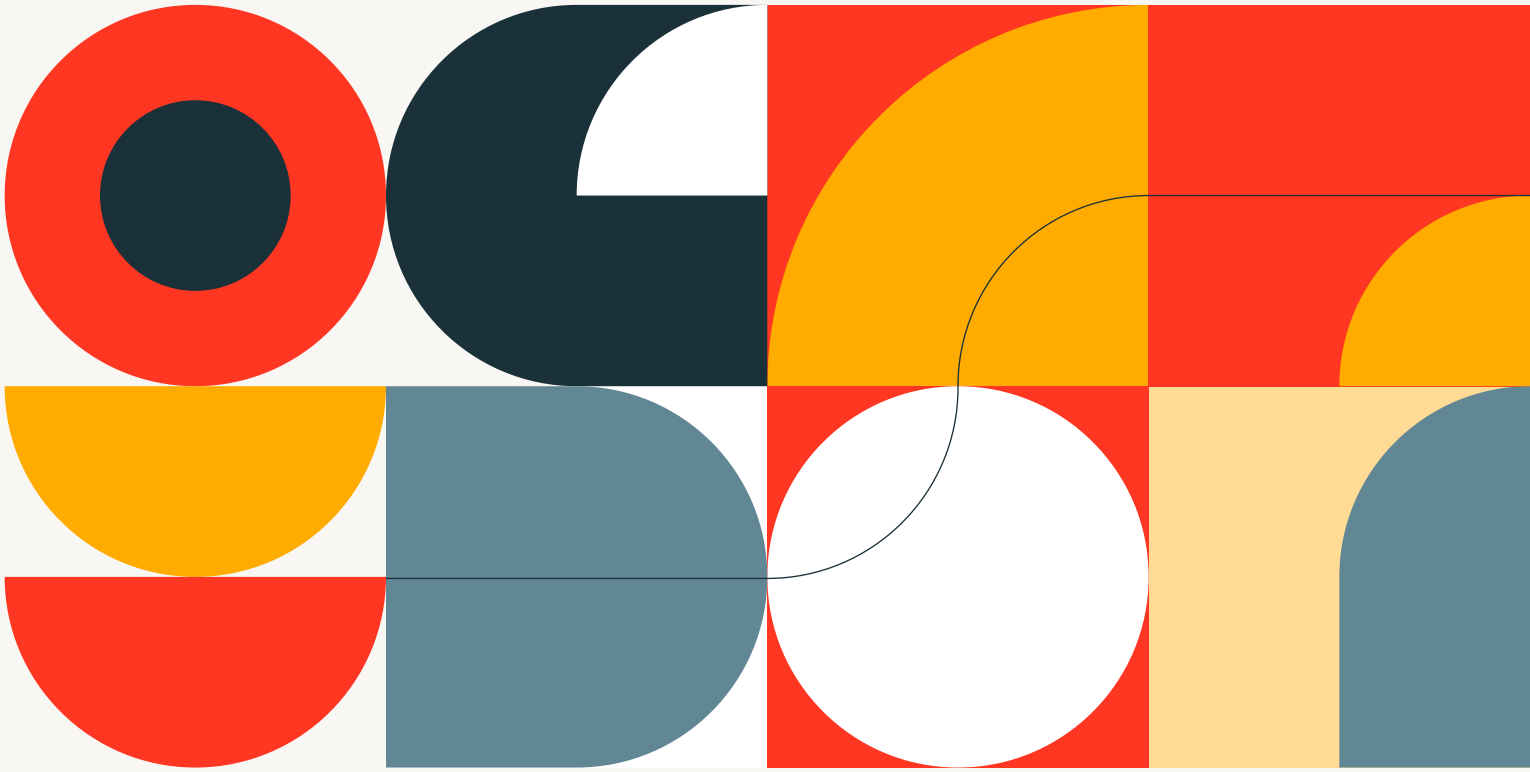

 | GTM Methodology | nabeel.ali@databricks.com

Customer References / Case Studies

AT&T

Problem

AT&T's 180 million wireless subscribers and 15 million broadband households generate 10+ petabytes of customer records and transactional data each day. This sheer volume of data created challenges for accessing and acquiring data due to AT&T's distributed and complex on-prem systems with 1500+ servers and 12000+ data sources.

Solution

AT&T migrated to Databricks to democratize data for over 2,000 users across 60+ business units, and create a unified view of over 100 million petabytes of data across their network while simplifying infrastructure management to increase agility and lower overall costs.

Result

Since implementing Databricks, AT&T has seen a 80% reduction in fraud, put 100+ ML models in production for fraud prevention, and is processing 100 million+ petabytes of data.

AT&T can now unlock value from its data at a much lower cost with $266M of value enabled by Databricks.

JOHN DEERE

Problem

John Deere is leading the industry in 'smart' agricultural solutions, leveraging data and analytics to drive precise outcomes in the field. But with sensors from 450,000+ 'connected 'machines streaming billions of rows of IoT real-time data, its complex system of legacy, on-premises data warehouses were unable to keep up-- increasing costs, decreasing performance and not meeting demands for ML.

Solution

Databricks ingests more than 8PB of structured and unstructured data and trillions of records (a 20.x increase over their data warehouse) to give data teams fast, reliable access to standardized data sets supporting more than 3,000 use cases across the organization--from proactive alerts that prevent failures, to precision agriculture that maximizes field output, to optimizing operations in the supply chain, finance and marketing.

Result

Shifting from a data warehouse to Databricks, Deere is saving more than $30M in infrastructure, more than $35M in productivity gains and $30M in currently deployed use cases.

jetBlue

Problem

Lack of online feature store hydration, inability to scale quickly, and high latency in its cloud data warehouse, JetBlue's data scientists were prevented from constructing scalable ML training and inference pipelines, hindering its ability to provide seamless customer experiences.

Solution

With Databricks, JetBlue is using LLMs built upon their own data to deliver better passenger experiences. Leveraging real-time streams of weather, IoT, and FAA data, JetBlue now operates the world's first digital-twin for efficient and safe operations, significantly minimizing delays.

Result

Increased Innovation in LLMs and generative AI, powering safer operations.

ATLASSIAN

Problem

Atlassian had a vision to improve speed to decision-making across its operations, democratizing data by centralizing and federating access to all users. Their first attempt did not succeed as siloed data warehouses were complex and costly to maintain and scale, were too proprietary, and did not support ML. They moved to a data lake architecture which improved scale and costs, but were slow for analytics workloads and had high barriers for self-service.

Solution

The best of both worlds, Databricks has brought together all analytics and ML on one open data platform, democratizing data securely to 3,000 users with minimal intervention from IT. More than half the company, including R&D, Support, HR, Marketing, Finance, are accessing analytics insights from the platform on a monthly basis.

Result

Databricks now powers virtually all business decisions at Atlassian and has delivered a 60% decrease in costs and a 30% reduction in delivery times versus their pre-Databricks infrastructure.

Awards & Benchmarks

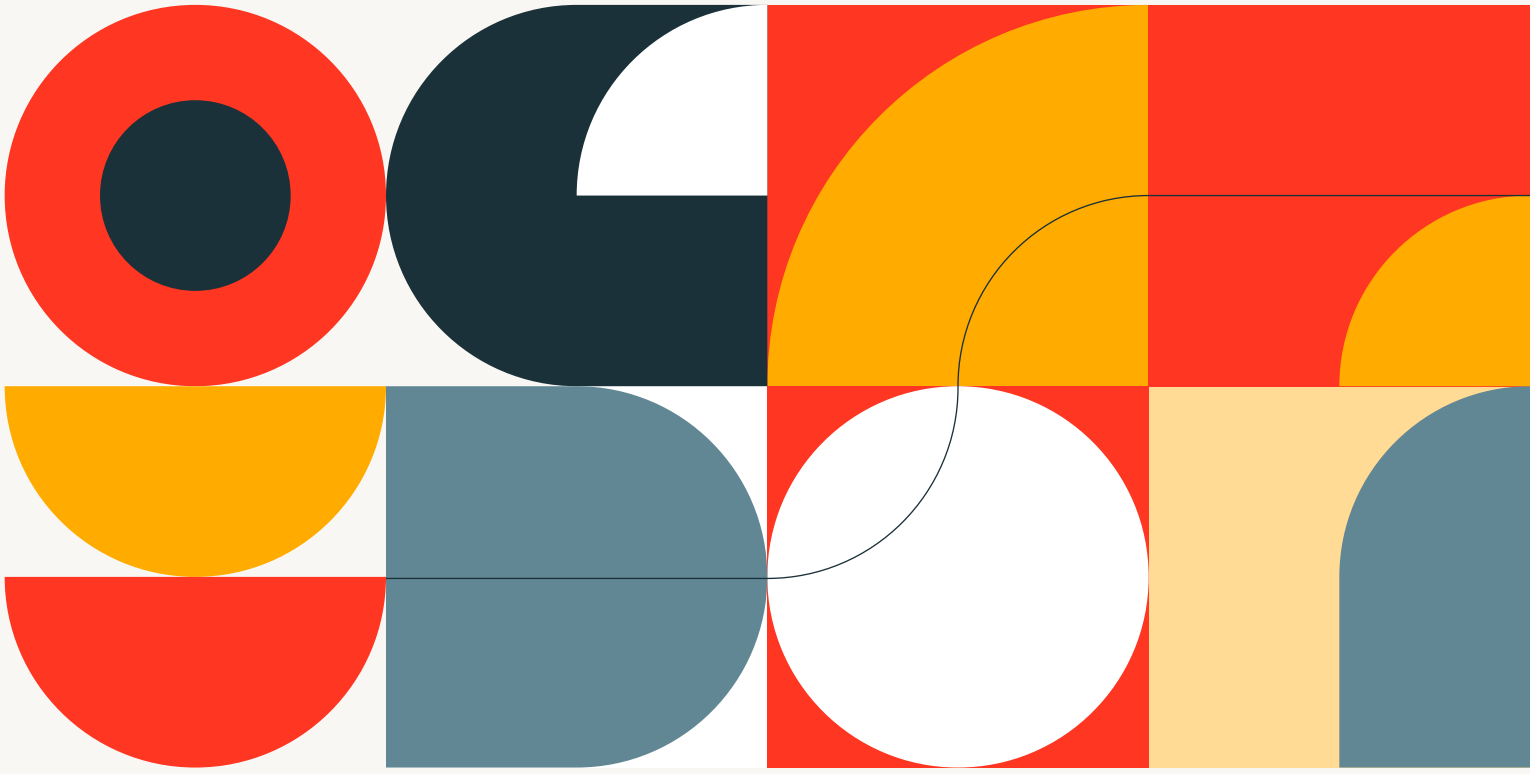

Analyst Reports

Gartner — Gartner Magic Quadrant

2023 Cloud Database Management Systems
(3rd consecutive year)

Learn More

2024 Data Science & Machine Learning Platforms

Learn More

Forrester — Forrester Wave

2024 Cloud Data Lakehouses

Learn More

2024 Cloud Data Lakehouses

Learn More

Databricks Reports/ Studies

CIDR Lakehouse Paper

"Lakehouse: A New Generation of Open Platforms that Unify Data Warehouse and Advanced Analytics" Seminal work on lakehouse paradigm authored by Databricks co-founders and engineers

Link to Paper

Databricks Reports/ Studies

Lakehouse Benchmark

"Databricks was 2.7x faster and 12x better in terms of price performance than Snowflake"

Link to Blog

MIT Technology Review: CIO Vision 2025

"72% of CIOs say that data is the biggest challenge for AI and 68% say unifying their data platform for analytics and AI is crucial"

Link to Paper

Technical Benchmarks

TPC-DS World Record 2021

Databricks sets official data warehousing performance record

Link to Blog

SIGMOD Photon Paper 2022

Photon: A Fast Query Engine for Lakehouse Systems

Link to Paper

Databricks Partner GTM Enablement

Driving success through end to end enablement to align to partner priorities

SCAN TO ACCESS TO PARTNER ACADEMY

LEARNING PATHS BY PRIORITY

- Getting Started
- Governance
- Generative AI
- EDW Migrations
- SAP
- Comp. Landscape

FOUNDATIONS	LEADERSHIP	SALES/PRESALES	TECH PRESALES
All customer facing roles	Executives, Practitioners, Managing Directors, Client Account Leads	Consultants, Sales, Field Engineering, Architect	Solution Architects, Sales Engineers, Field Engineers, Architects
SET THE **FOUNDATION** 1-3 hours each	STRATEGIC **INSIGHTS** 45 min each	DEEPEN **KNOWLEDGE** 2-3 hours each	VERIFIED **EXPERTISE** 2-3 hours each
Wondering with which course started your GTM enablement journey? You are in the right column.	Strategic Points of View that are forward-thinking and visionary.	Foundational Learning, Migrations Strategies, GTM Expertise and Industry.	PreSales Verified technical expertise needed to lead scoping and architectural calls with clients.
databricks GTM Foundations Pre-sales	databricks Strategic POV Leadership	databricks GTM Expertise Pre-sales; databricks Industry Expertise Pre-sales	databricks Tech Presales expertise Pre-Sales Verified

Databricks Partner Delivery Enablement

Driving success through end to end enablement to align to partner priorities

SCAN TO ACCESS TO PARTNER ACADEMY

LEARNING PATHS BY PRIORITY

- Getting Started
- Governance
- Generative AI
- EDW Migrations
- SAP
- Comp. Landscape

FOUNDATION	ASSOCIATE	PROFESSIONAL	SPECIALIZATION
Everyone who is hands on keyboard	Role-based tracks: GenAI Engineer, Data Engineer, ML/Data Scientist ML engineer, Data Analyst	Role-based tracks: GenAI Engineer, Data Engineer, ML/Data Scientist ML engineer, Data Analyst	Everyone who is hands on keyboard with a Databricks Delivery certification
DELIVERY **FOUNDATIONS**	BILLABLE **READINESS**	ADVANCED **EXECUTION**	STRATEGIC **SPECIALIZATIONS**
★	★★	★★★	
Establish core technical skills with Databricks Partner Foundations Training	Ensure readiness for billable engagements and gain competency in deploying Databricks solutions in production environments	Deliver advanced real-world projects and complex data and AI solutions	Master high-value specializations tailored for enterprise-scale Databricks solutions
Fundamentals Training Accreditations	Associate Assessment Certifications	Professional Assessment Certifications	Delivery Specialization Accreditation


databricks