

The Modern Data Paradigm:
Moving from Competition
to Synergy With
Microsoft Fabric and Databricks

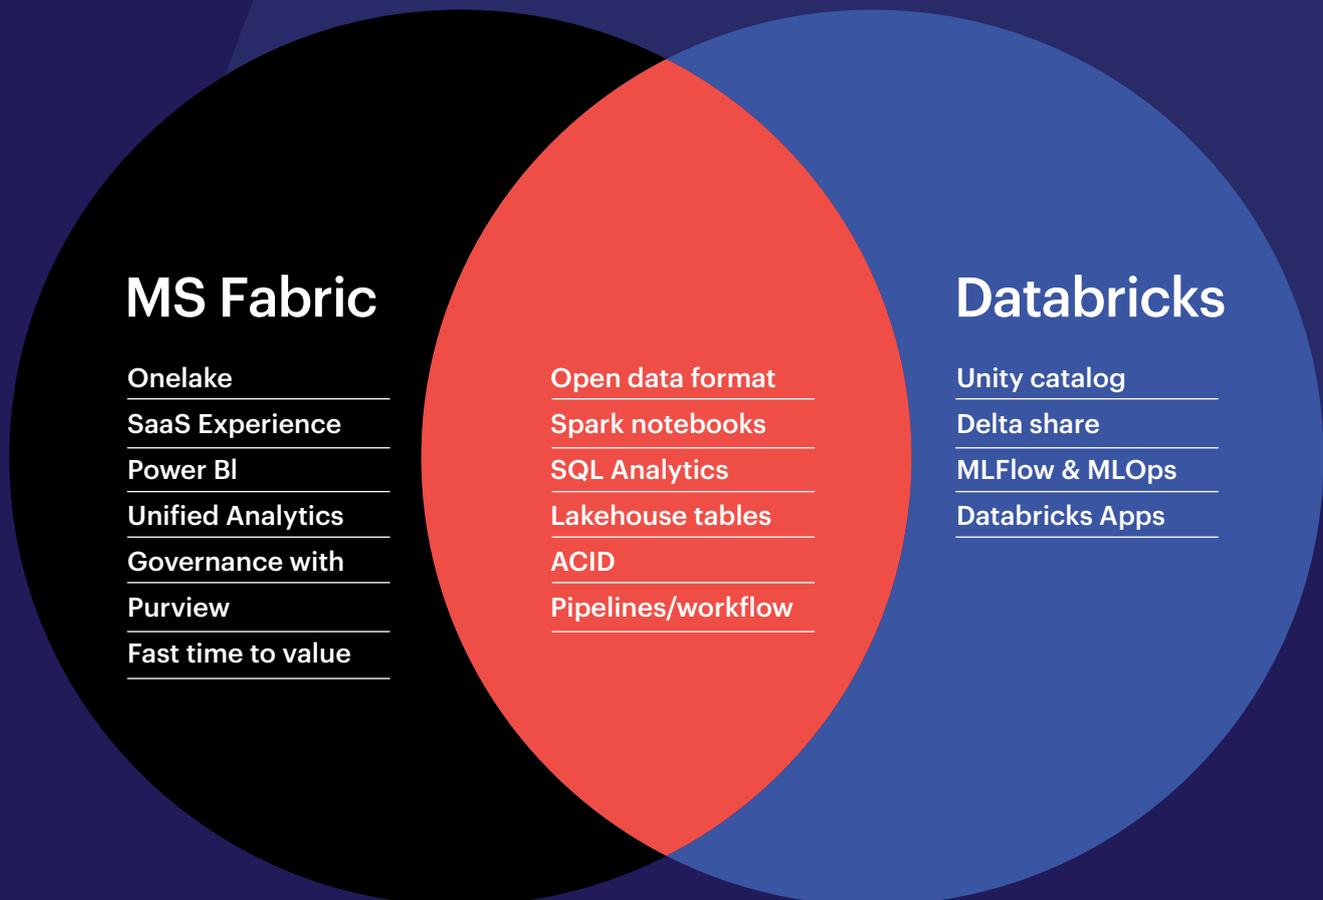
 PoV

The shift:

Why enterprises are prioritizing Microsoft Fabric

The analytics paradigm is shifting. MS Fabric has introduced a "SaaS-first" approach to the data lifecycle, removing the traditional burden of infrastructure management. By unifying data engineering, business intelligence (BI), and AI into a single, seamless environment, MS Fabric offers a predictable pricing model and a fast track to insights.

Conversely, Databricks remains the gold standard for distributed big data processing and complex machine learning (ML) workflows. With mature features like MLflow and highly customizable compute environments, it is the indispensable engine for heavy AI workloads and large-scale data science projects.



Microsoft Fabric vs. Databricks: Platform capabilities and overlap

The current industry reality

As the landscape evolves, MS Fabric is gaining rapid traction as a unified foundation for governance and analytics. This leaves enterprises at a crossroad: Should they double down on Databricks, migrate to MS Fabric, or adopt a hybrid strategy?

The real question: Integration over competition

The "MS Fabric vs. Databricks" debate is a false dichotomy. In a modern ecosystem, they are complementary pillars:

Databricks (PaaS): Optimized for high-performance data engineering, advanced ML, and multi-cloud flexibility.

MS Fabric (SaaS): Optimized for centralized storage (OneLake), seamless orchestration, and democratized analytics.

The strategic question is no longer "which one?" but rather: "How do we integrate both to build a scalable, governed, and AI-ready architecture?"

Better together A synergistic data landscape

By leveraging both platforms, organizations can play to the unique strengths of each:

MS Fabric: The primary interface for unified analytics, Power BI, and enterprise governance.

Databricks: The engine for advanced engineering, streaming, and deep AI/ML experimentation.

The shared foundation: Delta Lake serves as the universal language, ensuring total interoperability between the two environments.

Where Databricks Excel?

- Robust, highly available cloud infrastructure
- Scalable data engineering and AI/ML pipelines for large workloads
- High-volume data needing powerful compute
- Granular governance of data and AI via Unity Catalog
- Secure sharing with external parties on unknown infrastructure
- Strict egress control using traditional firewalls

In a MS Fabric Environment, Databricks Adds Value

For:

- Handle large-scale data engineering with complex transformations and orchestration
- Seamless data sharing with third parties
- Large Scale ML/Gen AI experimentation (notebooks, distributed training, model deployment & serving)
- Granular Network & Egress Security

Where MS Fabric Excels?

- Quick deployment
- Limited cloud expertise
- Predictable cost structure
- Curated workloads (semantic models, dashboards, reports)
- Enhances AI Foundry & Microsoft Purview for governance and AI
- Citizen developer friendly

In a Databricks Environment, MS Fabric Adds Value

For:

- Enterprise BI with governed semantic models and deep Excel/Teams integration
- Empower Business self-service with certified datasets, + Power BI and quick report delivery
- End user analytics enablement (row-level security, deployment pipelines, Data Hub)
- Streaming data via Event Hub

Choosing Databricks or MS Fabric — or both?

Achieving seamless interoperability

Integration is achieved through a "zero-copy" mindset:

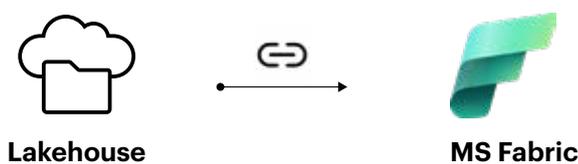
Data access: Databricks processes large-scale data and writes to Delta tables in ADLS Gen2. MS Fabric consumes this data via OneLake Shortcuts, enabling instant access for Power BI Direct Lake models without moving or duplicating data.

Orchestration: MS Fabric Data Pipelines or Azure Data

Factory can trigger and manage Databricks jobs within a larger workflow.

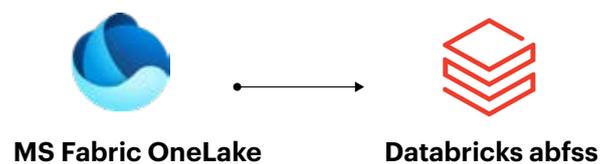
Governance: While Databricks utilizes Unity Catalog, Microsoft Purview provides an overarching governance layer across the entire Azure estate.

Databricks -> MS Fabric Shortcut



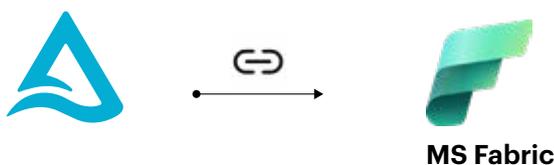
Data resides in ADLS/S3; MS Fabric creates shortcuts to source location

MS Fabric OneLake -> Databricks Access



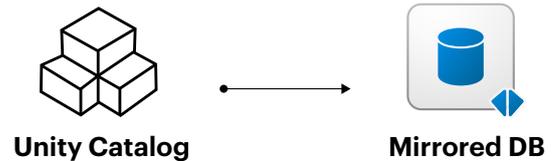
Data stored in OneLake; Databricks reads/writes via abfss

Data Virtualization in MS Fabric



Databricks delta tables exposed on MS Fabric using shortcuts

Mirrored Databases (Unity Catalog)



MS Fabric Mirrors selected Databricks UC

Engineering, Streaming, ML



Databricks

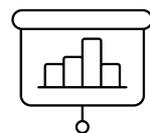
BI, Semantic Models, Governance



MS Fabric

purview

Reporting



Power BI with import, Direct Lake, Direct Query

Interoperability patterns: Databricks and MS Fabric

Zensar's point of view:

Synergy as the standard

At Zensar, we view MS Fabric not as a replacement for Databricks, but as its natural extension.

We recommend a dual-adoption model:

Maturity matters: For AI and ML, Databricks (specifically via MLflow and MosaicML) offers a level of maturity and scale that is currently unmatched. Databricks remains the system of record for model development and experimentation.

MLflow for experiment tracking, reproducibility, registry, and lifecycle management.

MosaicML (and broader Databricks AI capabilities) for scalable training, fine-tuning, and efficient LLM workflows.

Strong support for feature engineering patterns, distributed computing, and complex notebook-driven R&D. Why this matters: When the use case is high-scale training, rapid iteration, or sophisticated MLOps, Databricks reduces time-to-production and model risk.

Architectural best practice: Execute heavy-duty engineering and ML experiments within Databricks,

and use MS Fabric as the secure, user-friendly gateway for BI and reporting.

Databricks: Heavy ETL/ELT, streaming, data quality frameworks, advanced optimization, ML pipelines.

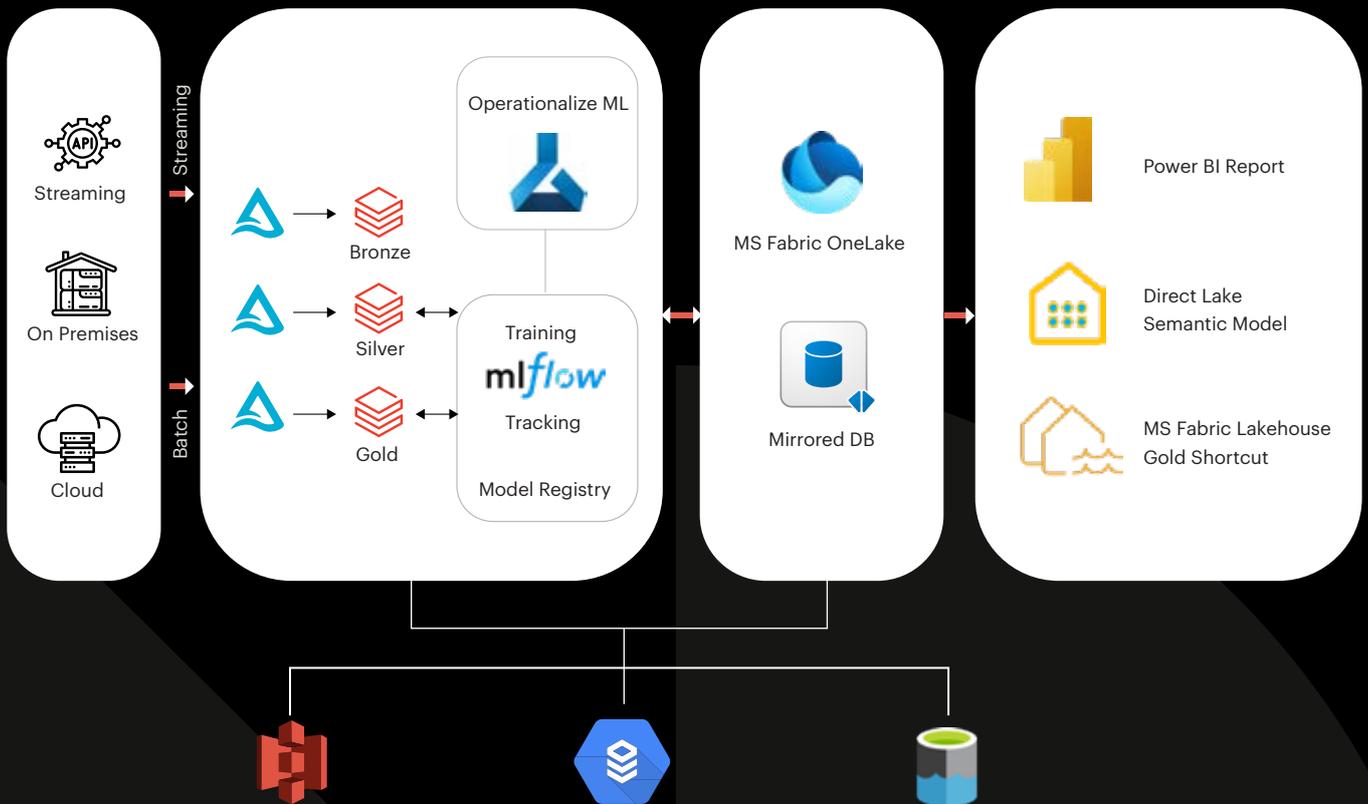
MS Fabric: Governed semantic models, self-service analytics, Power BI, business-friendly experiences, and workflows. Outcome: Fewer fragile BI pipelines, less duplication of curated datasets, and cleaner ownership boundaries (platform team builds; business consumes).

Databricks can produce excellent curated tables, but organizations often still need a strong enterprise semantic model. MS Fabric + Power BI semantic models allow you to standardize KPIs ("one definition of churn/yield")

Native integration: Co-engineered with Microsoft, Azure Databricks shares foundational platform alignment with MS Fabric, enabling increasingly seamless interoperability and smooth workload transitions across the modern Lakehouse ecosystem.

Databricks (Engineering & ML) PaaS-based

MS Fabric (Business Insights) SaaS-based



Challenges and future solutions

At Zensar, we have successfully operationalized this approach across several engagements in the banking and financial services domain, adopting a synergistic model in which Azure Databricks is used for scalable data engineering, transformation, and data quality; MS Fabric is used for analytics; and Azure Data Lake serves as the delta lake. This approach helps enterprises implement best practices for integration, performance, and governance in a scalable and sustainable manner.

Conclusion

Microsoft Fabric is not replacing Databricks—it is redefining how we deliver value. By combining the raw engineering power of Databricks with the intuitive, analytics-first SaaS model of MS Fabric, organizations unlock a complete end-to-end pipeline:

Raw Ingestion → Advanced Transformation → ML Experimentation → Governed Insights.

This "Better Together" architecture allows you to keep advanced engineering within your technical control (PaaS) while providing a simplified, high-speed lane for business intelligence (SaaS).

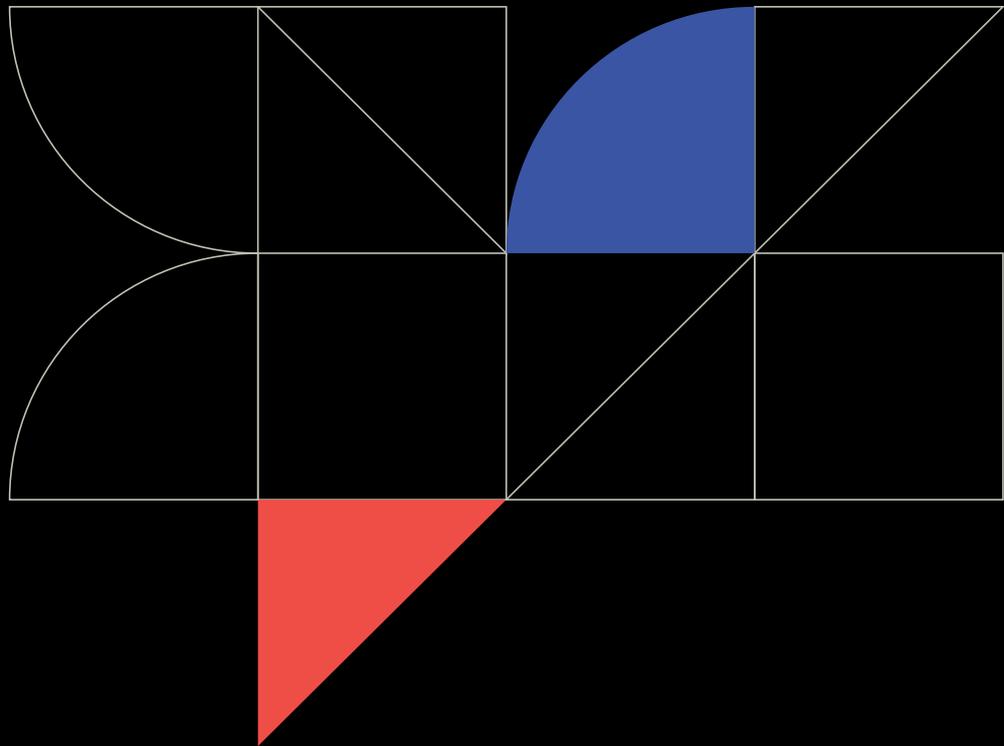
Author:

Pranav Ghatpande

Project Manager,
Data Engineering & Analytics, Presales

Pallavi Routaray

Solution Architect, Data Engineering & Analytics
Microsoft Fabric & Azure Data Platform



zensar
An  RPG Company

At Zensar, we're 'experience-led everything.' We are committed to conceptualizing, designing, engineering, marketing, and managing digital solutions and experiences for over 145 leading enterprises. Using our 3Es of experience, engineering, and engagement, we harness the power of technology, creativity, and insight to deliver impact.

Part of the \$4.8 billion RPG Group, we are headquartered in Pune, India. Our 10,000+ employees work across 30+ locations worldwide, including Milpitas, Seattle, Princeton, Cape Town, London, Zurich, Singapore, and Mexico City.

For more information, please contact: info@zensar.com | www.zensar.com