



Implementing a Databricks Lakehouse



inFINITIVE



databricks

Table of Contents

1 Introduction and Background	3
2 Transitioning from Data Strategy to Data Lakehouse	5
3 Lakehouse Architecture in Databricks	9
4 Determining What Data Goes into the Lakehouse	11
5 Connecting the Lakehouse to Other Systems	13
6 Security and Data Governance in Databricks	16
7 Centralized vs. Distributed Approaches to Building the Lakehouse	18
8 Implementation Methodology and Project Plan	20
9 Testing, Validation, and Ongoing Maintenance	22
10 Advanced Considerations and Future Directions	24
11 Conclusion and Key Takeaways	26

1 Introduction and Background

1.1 Purpose of This eBook

The primary goal of this eBook is to guide organizations in planning and implementing a Databricks-based lakehouse. This work assumes an enterprise data strategy has already been established, outlining strategic goals, governance principles, and target outcomes. For the Infinitive eBook titled, “Building a Winning Data Strategy for the Modern Enterprise”, see [Building a Winning Data Strategy for the Modern Enterprise](#). Building upon that foundation, this eBook offers a prescriptive approach for translating high-level strategy into real-world data engineering, analytics, and governance practices.

In many enterprises, data initiatives often stall due to fragmented ownership, complex infrastructures, or misaligned priorities. By focusing on Databricks-specific features, readers can gain clarity on the practical steps required to stand up a lakehouse, integrate disparate, siloed data, set up security and role-based access controls, and manage ongoing governance. The guidance here is intended to reduce complexity and accelerate time to value, demonstrating how Databricks becomes a foundational platform for data-driven decision-making.

1.2 Audience and Prerequisites

This eBook is designed for data architects, data engineers, project managers, and IT stakeholders responsible for establishing or modernizing data platforms. Familiarity with basic data concepts, such as structured vs. unstructured data and ETL/ELT processes, will help readers grasp the technical content. Understanding Databricks fundamentals—like notebooks, clusters, and the Spark engine—ensures the ability to follow along with the recommended best practices and architectural guidelines.

Beyond having a high-level technical understanding, readers should be aware of their broader corporate data strategy, especially as it pertains to data governance and regulatory compliance. This context enables the alignment of data lakehouse efforts with strategic objectives, ensuring that a Databricks solution fits seamlessly within the enterprise data ecosystem.

1.3 Structure and Objectives

This eBook is organized into twelve main sections, each exploring a critical aspect of implementing a Databricks-based lakehouse. Early sections focus on background and strategy, bridging the gap between high-level corporate objectives and the Databricks platform. Subsequent sections detail architecture options, data selection and ingestion methodologies, security, governance, and project management. Approaches to centralized vs. distributed governance and best practices for ongoing maintenance form the backbone of the middle chapters. The final chapters outline advanced considerations, future directions, and a thorough appendix.

Readers will gain:

- A comprehensive understanding of Databricks as a unified analytics platform.
- Best practices for ingestion, transformation, and governance in a lakehouse environment.
- Insights into security, roles, and responsibilities for robust data protection, and to satisfy audit and compliance needs.
- An introduction to Infinitive's metadata-based pipeline generator as an accelerator for pipeline creation.

2 Transitioning from Data Strategy to Data Lakehouse


2.1 Recap of Key Data Strategy Outputs

Before diving into the specifics of implementing a lakehouse, it is important to revisit the outcomes of the enterprise data strategy.

Typical elements include:

- **Business goals:** measurable targets the organization seeks to achieve through data-driven initiatives.
- **Data domains:** categorization of the most critical data assets (customer data, financial transactions, operational metrics).
- **Data governance principles:** frameworks and policies to ensure data is consistent, accurate, and compliant.
- **Prioritized data sources and use cases:** an established roadmap that lists which data areas and analytics initiatives should receive immediate focus.

These outputs provide a roadmap that outlines which data sources to ingest first and how to align the technical implementation with the overarching governance and business priorities.



2.2 Why a Lakehouse Model?

A lakehouse blends the scalability and cost benefits of a data lake with the reliability and performance of a data warehouse. Traditionally, organizations have managed two separate environments: a data lake for raw, unstructured data and a data warehouse for structured, analytics-ready data. However, this bifurcated approach often leads to data duplication, increased operational overhead, and synchronization challenges. Moreover, this approach leads to duplicated efforts in applying security and governance controls across multiple systems, increasing the chances for errors and creating the risk that one system may fall out of sync with another until governance changes can be consistently implemented everywhere.

The lakehouse concept addresses these issues by enabling ACID transactions, schema enforcement and evolution, and fine-grained governance on the same platform that stores both raw and curated data. Databricks has emerged as a market leader in delivering these combined capabilities at scale.

This model can:

- Improve performance for queries and analytics.
- Support advanced analytics and machine learning on large datasets.
- Reduce complexities in data movement and schema synchronization.

Comparison of Modern Data Architectures: Data Warehouse vs. Data Lake vs. Data Lakehouse

Attribute	Data Warehouse	Data Lake	Data Lakehouse
Primary Purpose	Structured analytics, reporting, and business intelligence	Flexible storage of raw, unstructured, semi-structured, and structured data	Combines structured analytics and reporting with flexible, scalable storage of all data types
Type of Data Stored	Structured, cleaned, curated data	Raw, structured, semi-structured, unstructured data	All data types—structured, semi-structured, unstructured; combines curated data with raw data
Schema Approach	"Schema-on-write" (structured schema defined upfront before loading data)	"Schema-on-read" (schema applied later at query time)	Supports both "schema-on-read" and "schema-on-write," offering greater flexibility
Data Storage Technology	Relational databases (e.g., Snowflake, Amazon Redshift, Teradata, BigQuery)	Cloud storage, object storage (e.g., Amazon S3, Azure Data Lake Storage)	Cloud-based object storage combined with open, scalable data formats like Delta Lake, Apache Iceberg, or Apache Hudi
Performance	High-performance for structured queries (optimized through indexing, partitioning, and caching)	Can be slower due to raw, unoptimized data storage (requires processing during query time)	Optimized query performance similar to data warehouses with caching, indexing, and optimized file formats
Data Quality and Governance	High level of governance, security, and data quality controls	Often limited governance; lower-quality data without additional controls	Strong governance and data quality through built-in mechanisms (e.g., Delta Lake ACID transactions, schema enforcement)
Typical Users	Analysts, business intelligence (BI) users, executives	Data scientists, developers, data engineers	Analysts, BI users, data scientists, and engineers; combines usability across roles
Complexity and Cost	Higher complexity and cost due to structured transformations and maintenance	Lower initial cost and complexity but higher long-term management effort	Balances initial cost and complexity with long-term scalability and ease of maintenance

2.3 Core Capabilities of Databricks

The Databricks Data Intelligence Platform provides a unified analytics platform that integrates the Apache Spark engine with collaborative notebooks, automated cluster management, and built-in security controls.

For organizations seeking a modern data platform, it provides:

- **Collaborative development:** multiple data engineers and data scientists can work in the same notebooks, versioned through Databricks Repos.
- **Elastic compute:** autoscaling clusters allow for cost-effective handling of peak loads.
- **ML integration:** seamless integration with libraries like MLlib, TensorFlow, and PyTorch for machine learning projects.
- **Delta Lake:** ensures ACID transactions, scalability, and built-in performance optimizations such as Liquid Clustering.



3 Lakehouse Architecture in Databricks

3.1 Conceptual Architecture

At the core of a Databricks-based lakehouse is Delta Lake, an open-source metadata layer built on top of Parquet files.


Delta Lake provides:

- **ACID compliance:** reliable updates and deletes to data in the lake.
- **Time travel:** the ability to query historical snapshots of data for auditing or point-in-time analysis.
- **Schema enforcement and evolution:** structured metadata that maintains data quality without sacrificing flexibility.
- **Secure Data Sharing:** Supports secure data sharing across organizations via the open protocol Delta Sharing.
- **Security and Compliance:** Provides encryption, authentication, authorization, auditing, and governance features to meet industry standards and regulations such as PCI, GDPR and CCPA.

Building on Delta Lake, the widely adopted Medallion architecture structures data into Bronze, Silver, and Gold layers:

- **Bronze layer:** raw, unprocessed data stored in its native format.
- **Silver layer:** data that has gone through preliminary transformations (cleaning, normalization).
- **Gold layer:** curated, business-ready data sets with applied business logic, used by analysts and data scientists.

This layered approach ensures data cleanliness, reliability, and consistency at each stage, while maintaining the ability to trace data lineage back to the raw sources.



3.2 Storage and Compute Separation

A fundamental principle of modern cloud platforms is separating storage from compute. In a Databricks environment:

- Storage is typically managed via cloud object stores like Azure Blob Storage, Amazon S3, or Google Cloud Storage, stored in Delta Lake format.
- Compute is delivered through Databricks clusters that can auto-scale to meet workload demands.

This architecture provides elasticity and cost optimization. Data is stored in a low-cost, open source Parquet format and remains accessible even when compute clusters are shut down, eliminating the need for always-on hardware resources. Organizations thus pay only for the processing time needed for transformations, analytics, and machine learning tasks.

3.3 Databricks Ecosystem Integration

Databricks connects with a wide range of partner technologies, spanning:

- BI Tools (Power BI, Tableau, Looker) for dashboards and reporting.
- Machine Learning Frameworks (MLflow, MLlib, scikit-learn) for advanced analytics and experimentation tracking.
- ETL/ELT Solutions (dbt, Azure Data Factors, Prophecy) for data ingestion and transformation.

For collaboration and version control, Databricks Repos enables teams to manage notebooks and other code artifacts through Git services such as GitHub and Azure DevOps. This approach encourages best practices in code reviews, testing, and continuous integration.

4 Determining What Data Goes into the Lakehouse

4.1 Data Prioritization Strategy

Choosing which data sets to move into the lakehouse first is key to achieving early wins.

Criteria may include:

- **Business value:** data that drives revenue, reduces risk, or informs key operational decisions.
- **Regulatory requirements:** data that must be securely managed for compliance and audit initiatives (GDPR, CCPA, HIPAA, PCI-DSS).
- **Data quality:** reliable, well-documented sources are easier to integrate and maintain.

Align these priorities with the enterprise data strategy's documented roadmap, focusing on immediate pain points and high-value opportunities.

4.2 Data Profiling and Assessment

Before ingestion, data engineers should perform an assessment to check for data quality and consistency. Databricks provides quick data exploration via notebooks, allowing teams to sample rows, compute summary statistics, and detect anomalies. For deeper data quality validations, Lakehouse Monitoring is built into the Databricks platform helping teams proactively discover quality issues before downstream processes are impacted

In some cases, data may be too low in quality or too redundant to justify immediate migration. Making that determination early prevents unnecessary complexity in the lakehouse environment.

4.3 Scope Management and Phased Onboarding

Given the potential size and complexity of enterprise data, tackling ingestion all at once can overwhelm technical teams and the lines of business they are supporting.

- **Phase 1:** Focus on high-impact and well-understood data sources.
- **Phase 2:** Expand to reference data and complementary sources that enhance analytics.
- **Phase 3:** Incorporate historical or streaming data sources for near-real-time analytics.

Adopting a phased approach reduces risks, ensures visible progress, and accommodates evolving organizational priorities.

5 Connecting the Lakehouse to Other Systems

5.1 Data Ingestion Approaches

Databricks supports multiple ingestion modes:

- **Batch ingestion:** scheduling notebooks or jobs to extract data periodically. This is often used for structured data from relational databases or files that are updated on a daily or weekly basis.
- **Streaming ingestion:** leveraging Spark Structured Streaming with Kafka, Azure Event Hubs, or AWS Kinesis to process data in near real time. This approach is essential for use cases such as IoT data, fraud detection, or high-frequency analytics.

When standardizing ingestion tasks, Infinitive's Metadata-Based Pipeline Generator proves invaluable-recently helping one of our insurance industry clients reduce development time for new data ingest pipelines by as much as 75%. It automates repetitive pipeline creation steps by referencing metadata that defines source schemas, target tables, and transformations.

This automation:


- ✓ Reduces manual coding and the risk of inconsistencies.
- ✓ Enforces predefined quality checks and governance standards.
- ✓ Speeds up onboarding of new data sources.

5.2 Building and Scheduling Data Pipelines

Once ingestion methods are defined, building robust data pipelines in Databricks can be done through:

- **Databricks Jobs:** native scheduling of notebooks, JARs, or Python files.
- **Delta Live Tables:** a declarative approach that simplifies pipeline creation and automated testing, enabling incremental updates on datasets.
- **External orchestration:** solutions like Airflow, dbt Cloud, Azure Data Factory, or AWS Step Functions can coordinate multi-step processes across various services.

Best practices for pipelines include implementing retry logic, maintaining clear logging, and setting up alerts for failures or performance bottlenecks. Monitoring pipeline status is essential to ensure both reliability and data freshness.



5.3 Integration with External Data Warehouses and BI Tools

In many enterprises, existing data warehouses like Snowflake, Amazon Redshift, or on-premises solutions remain integral.

Integration strategies may include:

- Direct synchronization of curated Gold data back into external warehouses.
- Federated queries that bridge cloud data warehouses and the Databricks lakehouse.
- Publishing curated datasets to BI platforms like Power BI or Tableau for broader business consumption.

6 Security and Data Governance in Databricks

6.1 Security Foundations

Strong security in a lakehouse environment is paramount. It typically begins with:

- **Authentication and access control:** leveraging Azure Active Directory, AWS IAM, or equivalent systems to manage user identities and group memberships.
- **Network security:** using VNet or VPC integrations, private endpoints, and firewall rules to limit exposure to the public internet.
- **Encryption:** at-rest encryption on cloud storage, plus TLS in transit for data movement within and outside of Databricks clusters.

Security baselines vary by cloud provider and organizational policy, but a zero-trust approach helps ensure only authorized entities can access resources.

6.2 Unity Catalog for Governance

Unity Catalog is Databricks' unified governance solution that manages data permissions and lineage across multiple workspaces and data assets. It allows organizations to:

- Build a centralized catalog encompassing unstructured data, structured data tables, machine learning models, dashboards, and various other data assets to streamline and simplify data discovery.
- Enforce row- and column-level security, ensuring users can only see the data they are entitled to.
- Trace data lineage to understand data transformations and meet compliance regulations.

By adopting Unity Catalog early in the project, enterprises establish a robust framework for governance that spans data engineering and analytics use cases. Infnitive can help every step of the way and has a robust [Unity Catalog Migration](#) methodology.

6.3 Managing Permissions and Roles

Admins can assign privileges through a combination of group-level and individual permissions, mapped to roles like Data Steward, Data Engineer, or Business Analyst. A best practice is to implement a role-based access control (RBAC) system:

- Data Stewards manage data quality and definitions.
- Data Engineers create and maintain pipelines.
- Analysts consume curated datasets.

This approach reduces administrative overhead and aligns privileges with business functions.

6.4 Auditing and Monitoring

In regulated industries, comprehensive auditing is crucial for demonstrating compliance. Unity Catalog generates audit logs detailing user activity, data reads and writes, and modifications to datasets. These logs can be integrated into SIEM solutions for anomaly detection and cyber threat analysis. Proactive monitoring ensures timely responses to unusual patterns, unauthorized data access, or other security alerts. Databricks has achieved ISO certification and compliance with more than 15 standards including HIPAA, GDPR and CCPA, PCI-DSS and FedRAMP, among others.

7 Centralized vs. Distributed Approaches to Building the Lakehouse

7.1 Centralized Approach

In a centralized model, the organization forms a single program office or team that spearheads all aspects of the lakehouse:

- + **Pros:** standardized governance, consistent architecture, minimized duplication of efforts.
- **Cons:** potential delays due to lengthy approvals, risk of “big bang” failures if the scope is too large.

This approach can be advantageous for highly regulated industries or companies with a strong culture of central control, ensuring consistent data definitions and compliance across the board.

7.2 Distributed Approach

By contrast, a distributed model allows different functional areas or domains (finance, marketing, operations) to develop their portions of the lakehouse with relative autonomy:

- + **Pros:** rapid delivery of solutions tailored to local needs, domain expertise in each business unit, faster agile cycles.
- **Cons:** risk of duplicative work, inconsistent standards, governance fragmentation.

Communication across distributed teams and robust governance guardrails (such as a global Unity Catalog policy) can help maintain alignment while still enabling local innovation.

CENTRALIZED VS. DISTRIBUTED DATA TRANSFORMATION

Centralized Data Transformation

Consistent strategy and governance across enterprise.

Single corporate data model and governance process understood by all.

Slow implementation due to complexity and cross functional coordination.

Coordination and logistics across a large organization can slow program down.

Distributed Data Transformation

Better alignment with local business needs and agility.

Units with greatest benefits from data transformation can move more aggressively.

Risk of duplicative and redundant efforts across divisions.

Lack of standardization can lead to data silos and inconsistency.

7.3 Hybrid Model

A hybrid approach blends centralized oversight with domain-level autonomy. Unity Catalog, for instance, can provide standardized security and governance, while individual domains can design their own pipelines and analytics workflows. This approach harnesses local domain knowledge without sacrificing essential enterprise controls, and may include things like central decisions about what tooling to use.

7.4 Choosing the Right Approach

Deciding between centralized, distributed, or hybrid approaches often depends on:

- Organizational culture and leadership style.
- Regulatory and compliance demands.
- Maturity of data governance frameworks.
- Availability of skilled resources in each functional area.

For example, a large global bank might lean more centralized, whereas a smaller, fast-moving tech firm might adopt a distributed model for speed.

8 Implementation Methodology and Project Plan

8.1 Project Phases

A structured methodology ensures a smooth rollout. At a high level:

- **Phase 1:** Foundation and Infrastructure. Stand up the Databricks workspace, configure networks and security baselines, and establish initial governance in Unity Catalog.
- **Phase 2:** Data Ingestion and Transformation. Ingest priority data sources into Bronze, define transformations into Silver, and establish basic quality checks.
- **Phase 3:** Data Enrichment and Consumption. Create Gold-level datasets that incorporate business logic and advanced transformations. Begin ML experimentation.
- **Phase 4:** Governance and Scalability. Expand Unity Catalog usage, ensure cross-functional adoption, refine performance, and start building solutions on cleansed, enriched data such as advanced analytics, MLOps pipelines or GenAI applications.

Each phase should end with a clear set of deliverables and quality checks, ensuring readiness for the next phase.

8.2 Key Roles and Responsibilities

Successful implementations require collaboration among multiple roles:

- **Data Architect:** defines data architectures aligned with regulatory standards and optimized for performance and scalability, including support for cross-domain data integration and governance.
- **Data Engineer:** builds and maintains pipelines, focusing on ingestion, transformation, and scheduling.
- **Data Governance Lead:** manages policies, ensures compliance, and handles data stewardship processes.
- **Business SMEs:** provide domain expertise to shape business logic in transformations and identify the most valuable use cases.

Effective coordination among these roles is crucial throughout design approvals, data validation, and security setup to ensure a successful implementation.

8.3 Timeline and Milestones

Timelines can vary significantly based on an organization's scale and complexity.

A rough estimate might be:

- **Phase 1 (Foundation):** 4 to 6 weeks.
- **Phase 2 (Ingestion/Transformation):** 8 to 12 weeks.
- **Phase 3 (Enrichment/Consumption):** 6 to 10 weeks.
- **Phase 4 (Governance/Scalability):** ongoing or iterative over months.

Regular checkpoints ensure that unforeseen challenges can be addressed without compromising the overall schedule.

8.4 Using Infinitive's Metadata-Based Pipeline Generator

Throughout the pipeline creation process, Infinitive's generator can automate many tedious tasks.

The workflow typically involves:

- **Metadata specification:** define source and target connection details, table schemas, and transformation rules in a metadata file.
- **Automated pipeline creation:** the generator produces Databricks workflows that handle ingestion, error handling, and data validation.
- **Deployment and testing:** generated pipelines are tested in lower environments (e.g., dev or staging) before being promoted to higher environments.

Common pitfalls include incomplete metadata, misaligned naming conventions, or failing to keep metadata files updated as schemas evolve. Establishing a governance process around pipeline metadata ensures smooth, ongoing use of the generator.

9 Testing, Validation, and Ongoing Maintenance

9.1 Data Quality and Validation


As data progresses from Bronze to Silver to Gold, validation steps confirm accuracy and completeness:

- **Unit tests:** apply validation logic to data transformations in notebooks.
- **Automated checks:** employ frameworks like Databricks Lakehouse Monitoring to define and track data quality checks at each layer.
- **Peer review:** encourage code reviews and pipeline walk-throughs to catch logical errors.

Continuous validation ensures that the data delivered to users meets both technical and business requirements.

9.2 Performance Tuning

Databricks offers flexible cluster sizing and auto-scaling for resource optimization, but performance also depends on how data is organized. Liquid Clustering delivers optimal performance without manual tuning or maintenance by automatically handling data layout decisions based on query patterns and usage. Additionally, it incrementally clusters and organizes data as it is ingested, maintaining consistent file sizes and reducing the need for frequent manual compaction operations.



9.3 Operational Monitoring

While Databricks supplies native logs and the Spark UI for debugging and performance insights and extensive system tables with operational information, many enterprises integrate with third-party observability tools such as Datadog or Azure Monitor. By setting up alerts on key metrics (like job failures, spikes in compute usage), teams can quickly address issues and maintain consistent service levels.

9.4 Change Management

A well-defined change management process involves:

- Version control through Databricks Repos or external Git repositories.
- CI/CD workflows with automated unit and regression tests in a staging environment.
- Scheduled release cycles for new or updated pipelines.

This discipline is crucial to prevent breaks in data availability or quality when multiple teams are updating code and configurations simultaneously.

10 Advanced Considerations and Future Directions

10.1 Machine Learning and MLOps

Beyond batch analytics, organizations often want to implement advanced analytics or AI initiatives. Databricks has robust and industry leading support for MLOps through:

- **MLflow:** tracks experiments, hyperparameters, and model versions.
- **Delta Lake:** storage for training data, ensuring consistent and reproducible training sets.
- **Model deployment:** hosting model serving endpoints within Databricks or exporting models to external services.

Integrating MLOps from the outset can accelerate the transition from proof-of-concept models to production-grade solutions.

10.2 Real-Time Analytics

Low-latency data processing can be a game-changer for use cases that demand real-time insights.

This might involve:

- **Structured Streaming:** continuous ingestion and processing of events from Kafka, Apache Pulsar, Event Hubs, and other streaming services.
- **Delta Live Tables:** automating real-time transformations and sink definitions.
- **Real-time dashboards:** connecting BI tools directly to Silver or Gold tables that refresh in near real time.

Such capabilities can be critical for fraud detection, IoT device monitoring, or dynamic pricing engines.

10.3 Scalability and Multi-Cloud

As data volumes grow or organizations expand globally, multi-region and multi-cloud deployments may be considered.

Key challenges include:

- **Consistent governance:** ensuring Unity Catalog policies and lineage tracking extend across environments.
- **Network latency:** replicating data across regions requires robust networking design.
- **Security compliance:** different clouds and regions may have varied security standards or data residency regulations.

Automated infrastructure provisioning (using Terraform or equivalent) and standardized governance processes are essential to maintaining coherence across multiple deployments.

10.4 Evolving Tool Landscape

Databricks continues to evolve, introducing new features that extend lakehouse capabilities.

Potential future enhancements might include:

- Expanded Unity Catalog functionalities such as lineage for notebooks and ephemeral views.
- Lakehouse AI features that simplify advanced analytics.
- Closer integration with partner solutions for streaming, BI, and data observability.

Monitoring these trends helps organizations anticipate upgrades and plan technology roadmaps that keep pace with business needs.

11 Conclusion and Key Takeaways

11.1 Recap of Major Lessons

This eBook has traced the implementation journey from an established data strategy to a fully operational lakehouse in Databricks.

Key points include:

- Applying a phased methodology to source, ingest, and curate data while ensuring security and governance.
- Weighing trade-offs between centralized and distributed approaches to data ownership.
- Leveraging Delta Lake for ACID transactions, schema enforcement, and time travel.
- Emphasizing continuous integration, continuous delivery, and best practices for performance tuning.

Organizations that carefully plan each stage—from data selection to advanced analytics—are more likely to see rapid returns on their lakehouse investment.

11.2 Checklist for Success

Teams can review these essentials before and during project execution:

- Establish governance early, including Unity Catalog configuration.
- Prioritize data sources based on business impact and regulatory requirements.
- Document pipelines, transformations, and data quality checks thoroughly.
- Integrate version control and CI/CD to manage continuous changes.
- Continuously evaluate performance and adapt infrastructure settings as data volume and complexity grow.

11.3 Recommendations and Next Steps

Enterprises new to the lakehouse architecture should start with a pilot. By focusing on a single domain or a well-defined set of data, teams can:

- Validate architectural decisions in a lower-risk environment.
- Collect feedback and refine processes before scaling out.
- Demonstrate the value of the lakehouse to stakeholders and secure further investment.

Governance, security, and continuous improvement must remain core tenets throughout the lifecycle. As the enterprise grows more comfortable with Databricks, expanding into real-time analytics or advanced ML capabilities can further unlock business value.

11.4 Further Resources

For deeper technical insight or best practices, please reach out to Infinitive (Chelsea.Woodring@Infinitive.com), a trusted Databricks partner, for help. There are also extensive Databricks documentation, community forums and publications that can be found through Databricks. Infinitive has extensive experience with and documentation of best practices to utilize the platform and get you started that we would love to discuss with you.