

# Leveraging Open Table Formats for Scalable and Flexible Data Infrastructure

Ravi Rane, Pooja Mulik

**Abstract:** *This article offers a thoughtful deep-dive into one of the most practical challenges in data engineering today: building scalable, reliable, and agile data systems without compromising affordability or openness. The author moves beyond the usual surface-level praise of Open Table Formats (OTFs) to reveal how their architectural finesse through ACID compliance, schema evolution, and metadata optimization solves long-standing limitations of traditional data lakes. What stands out is the way Apache Iceberg, Delta Lake and Apache Hudi are presented not as abstract tools, but as enablers of real-world, high-performance analytics, handling everything from Time Travel queries to streaming ingestion and decentralized data ownership. It is evident that the narrative doesn't just list features, but connects them meaningfully to broader patterns like the lakehouse model and data mesh compatibility. This suggests that the evolution of OTFs is not just technical but philosophical, redefining how organizations approach data reliability at scale. That said, the piece could have drawn stronger parallels to industry case studies, yet its detailed walkthrough of transactional mechanics and storage abstractions compensated well. Overall, the work reads like a well-informed practitioner's guide, blending conceptual clarity with architectural pragmatism.*

**Keywords:** Open Table Format, Query Optimization, Delta Lake, Apache Spark

## 1. Introduction

The construction of robust, adaptable, and affordable data infrastructures is crucial for contemporary organizations that handle substantial amounts of data. Although conventional data lakes provide considerable adaptability and economical storage solutions, they encounter fundamental issues concerning transactional reliability, structural changes, and efficient data retrieval for large datasets.

**Open Table Formats (OTFs)**—including **Apache Iceberg**, **Delta Lake**, and **Apache Hudi**—address these challenges by introducing a **transactional, metadata-rich abstraction layer** over distributed object stores. These formats effectively unify the scalability and openness of data lakes with the reliability and structure of data warehouse systems, thereby enabling **robust and extensible architectural patterns** for modern analytics and data engineering workflows.

### 1) Understanding Open Table Formats

**Open Table Formats** serve as storage layer standards, offering **ACID transactions**, **schema evolution**, **version control**, and **efficient metadata management**. These functionalities are built upon columnar file formats such as Parquet and ORC.

- **Apache Iceberg** was created at Netflix and donated to the Apache Foundation. It introduces features like hidden partitioning, snapshot isolation, and atomic commits.
- **Delta Lake**, developed by Databricks, brings reliable data engineering features to Spark-based architectures.
- **Apache Hudi** is designed for streaming and incremental processing use cases, providing capabilities like record-level upserts and file-level indexing.

### 2) Core Scalability Challenges in Modern Data Lakes

Before exploring how OTFs solve scalability, it's essential to understand the architectural bottlenecks that exist in traditional data lakes:

- **High latency and full table scans** due to lack of metadata indexing.
- **Eventual consistency issues** with concurrent writes and schema changes.

- **Inefficient file sizes** resulting from streaming ingestion or frequent updates.
- **Limited support for ACID semantics** across distributed systems.

## 2. How OTFs Enable Scalable Architecture

### 2.1 ACID Transactions for Consistency at Scale

Open Table Formats support **ACID (Atomicity, Consistency, Isolation, Durability)** transactions, which are critical in multi-user, multi-job environments. For example, Apache Iceberg uses **snapshot-based isolation**, enabling safe concurrent reads and writes without locking entire tables.

"The snapshot isolation design in Iceberg ensures that readers see a consistent view of the data while writers operate in parallel." – [Apache Iceberg documentation](#)

### 2.2 Efficient Metadata Management

Metadata bottlenecks are a major concern at scale. OTFs address this by maintaining metadata in optimized formats.

- **Iceberg** stores metadata in **manifest files**, separating file-level and table-level metadata.
- **Delta Lake** uses a **transaction log (\_delta\_log)** to track changes incrementally.
- **Hudi** implements **timeline-based metadata**, optimizing for incremental queries and write operations.

This design enables **partition pruning**, **schema validation**, and **snapshot reads**, significantly reducing query planning time.

### 2.3 Schema Evolution and Enforcement

Scalable systems must accommodate change. OTFs support **schema evolution** with backward and forward compatibility.

- Iceberg enforces **schema validation rules**, preventing incompatible changes [1].
- Delta Lake supports column additions, renames, and type changes with schema enforcement [2].

- Hudi enables flexible schema evolution using Avro schemas [3].

This allows systems to evolve data models without breaking consumers.

## 2.4 Time Travel and Data Versioning

OTFs allow querying past snapshots or changes using **time travel** features. This is useful for:

- Reproducibility in data science
- Auditing and compliance
- Rolling back bad writes

For instance:

***SELECT \* FROM sales VERSION AS OF 12345;***

Such features are foundational for robust, scalable data management.

## 2.5 Streaming and Incremental Ingestion

Ingesting and processing data in near real-time is critical for scalability. OTFs support **streaming upserts** and **incremental pull queries**.

- **Apache Hudi** is optimized for **merge-on-read** and **copy-on-write** strategies, supporting Kafka-based pipelines [3].
- **Delta Lake** provides **structured streaming** support with transactional guarantees [2].

This allows organizations to maintain **low-latency pipelines** while ensuring consistency and reliability.

## 3. Architectural Patterns with OTFs

### 3.1 Lakehouse Architecture

OTFs form the core of the **lakehouse paradigm**, which unifies data lakes and data warehouses. With OTFs, data lakes gain:

- SQL query ability via engines like Trino, Presto, and Spark
  - ACID guarantees and schema management
  - Governance and time travel
- “The Lakehouse architecture with Iceberg enables high-throughput, consistent, and multi-engine analytics over massive datasets.” – Netflix Tech Blog

### 3.2 Multi-Engine Access

An Iceberg or Delta Lake table can be simultaneously accessed by **Spark**, **Flink**, **Trino**, and **Presto**, enabling **decoupled computation** and promoting scalability in compute-heavy environments.

### 3.3 Data Mesh Compatibility

In decentralized architectures like **data mesh**, OTFs help create **domain-owned, interoperable data products** with version control and governed access.

## 4. Best Practices

To simplify partition design, Iceberg employs hidden partitioning. Both Delta and Hudi utilize scheduled compaction to ensure sustained performance. For cost-effectiveness and scalable storage, leverage object storage with transactional Optimized Row Columnar (ORC) files.

## 5. Conclusion

Crafting a scalable data architecture involves more than just setting up storage; it requires ensuring strong **data reliability, consistent transactions, and adaptable architecture**. **Open Table Formats** provide a logical abstraction and sophisticated tools that transform basic object storage into a durable and scalable data management system, ideal for contemporary analytical tasks.

Online Transactional Files (OTFs) play a vital role in creating scalable, high-performing data platforms that integrate various systems. They achieve this by providing ACID transaction guarantees, adaptable schema management, optimized metadata indexing, and interoperability across different processing engines.

## References

- [1] Apache Iceberg Documentation. <https://iceberg.apache.org/>
- [2] Delta Lake Documentation. <https://docs.delta.io/>
- [3] Apache Hudi Documentation. <https://hudi.apache.org/>
- [4] Venkataraman, S., Yang, J., et al. (2016). *Apache Spark's Structured Streaming: A declarative API for real-time applications*. ACM SIGMOD.
- [5] Armbrust, M., Xin, R. S., et al. (2021). *Delta Lake: High-Performance ACID Table Storage over Cloud Object Stores*. VLDB 2020.
- [6] Ghodsi, A., et al. (2021). *Lakehouse: A New Generation of Open Platforms that Unify Data Warehousing and Advanced Analytics*. Databricks.