

# 更高效率低成本的 Observability 2.0 時代即將來臨


Scott Liao

Sr. Solutions Architect  
AWS

Joey Wu

Solutions Architect  
AWS

# 這幾年來隨著 AI 的發展，Observability 也帶來更大需求



# 以 Netflix 為例：

## Metrics

### Atlas TSDB

17 TB Time Series Each Day  
With 2 week retention

## Logs Errors Events

### Insight Logs

3 PB of Logs Collected Each Day  
with 2 week-6 month retention

↑ 200% YoY

## Tracing

### NfDT


3 PB of Logs Collected Each Day  
with 2 week-6 month retention

↑ 230% YoY

# 現況大型企業的挑戰

- 必須監控的指標資料量爆炸性成長
- 成本管理逐漸失控
- 微服務架構下的系統複雜度與依賴關係
- 企業規模越大，監控延遲要求越低
- 從 API First 到 AI First 加劇 Observability 的問題

# Observability 1.0



# Observability

1.0

“Three pillars.”  
metrics, logs, traces



2.0

Single source of truth  
wide structured logs

# Observability 2.0

- Observability 1.0 傾向將 Logs, Metrics 以及 Traces 分開處理, 但 Observability 2.0 主張整合分析。
- 以事件為中心 (Event-Centric) 將單次 request/operation 看成事件, 追蹤生命週期
- 不再是「資源使用」監控, 而是以「用戶體驗」以及「使用者行為」的維度觀測
- 高維度資料, 例如 `user\_id`, `session\_id`, `feature\_flag` 等
- 即時性的探索, 不僅止於 Static Dashboard, 使用 GenAI 探索未知問題
- 地板價的儲存成本

# Observability 2.0 與 Data Lakehouse 架構





# 為什麼資料平台會演進成今日的 Lakehouse



2010 – 2015

## Data lake

- Fixed compute and storage capacity
- Mostly on-prem
- Harder to use and manage



2022 – ...

## Data lakehouse

- Data in open file and table formats
- No need to copy and move data
- Multiple best-of-breed processing engines



1980 – 2010

## Enterprise data warehouse

- Fixed compute and storage capacity
- Mostly on-prem
- Harder to use and manage



2015 – 2023

## Cloud data warehouse

- Scale storage and compute independently
- Must load data into proprietary system
- Limited to one processing engine
- Cost prohibitive



2023 – ...

## Data lakehouse 2.0/Data mesh

- Autonomous (AI) semantic layer
- GitHub for data: “Data as a Product”

# 資料量日益增長 需求改變了

Can we do all of this with the same data governance and open standards?

Can we decouple storage from compute and support diverse consumers?

DATA  
WAREHOUSE

Can we bring the open source flexibility to the data warehouse?

DATA  
LAKE


Can we deliver E2E Governance?

Can you deliver best price/performance?

Can we bring the performance and strong ACID properties of data warehouse to data lakes?


# 平台也在演進

Lakehouse 結合了 warehouse 和 data lake 兩邊的優勢



# Netflix 搭建的 Lakehouse 平台一覽

## Big data platform architecture



Ref: <https://www.youtube.com/watch?v=jMFMEk8jFu8>

# Open Table Formats

**Open table formats (OTFs)** provide transactional support and simplify data lake optimization and management



Apache Hudi

---



Apache Iceberg

---



Delta Lake

---

# Open Table Format Benefits



ACID  
Compliance

---



Scalable Meta Data Handling

---



Schema enforcement &  
Evolution

---




Time-Travel  
Support

---

# 現代的 Lakehouse 通常採用開放表格式(Open Table Format)


Application



Processing



Catalog




Open Table format



File format



Storage




Data Sources





# Why Lakehouse?



**Data lifecycle and management remains complex, especially for large organizations**  
Duplicative copies, “expert” ETL, “dark data,” governance complexity, not self-service



# 哪些企業正在導入 Lakehouse 解決問題？



# 使用場景分析

- Apple：大規模分析工作負載
- Netflix：內容推薦系統、可觀測系統、管理 exabyte 級數據湖
- Expedia：旅遊數據分析、客戶數據管理
- Tencent：手機 QQ 安全數據入湖（28億用戶維度表，日均百億級消息）、新聞文章索引系統
- eBay：電商數據分析、用戶行為分析

Uber's lakehouse stores GPS traces, ride events, driver behavior data, and operational metrics from millions of rides daily

- **Dynamic Pricing:** Real-time adjustment of ride prices based on weather, traffic, and demand
- **ETA Predictions:** Instant calculation of estimated arrival times using live traffic data
- **Fraud Detection:** Real-time identification of fraudulent activities across the platform

## Performance and Cost Savings

- With this approach, we are able to decrease the pipeline run time by 50% and also decrease the SLA by 60%.



Ref: <https://www.uber.com/en-TW/blog/ubers-lakehouse-architecture/>

# Migrates Hive legacy HDFS to Iceberg on S3

- Event data ingestion benefited particularly from Iceberg's flexible partitioning configurations
- Internal optimizations for CDC ingestion and performant physical deletes for selective deletion.

## Performance and Cost Savings

- All in all, Airbnb experienced a 50% compute resource-saving and 40% job elapsed time reduction in its data ingestion framework with Iceberg and other open source technologies.



Ref: <https://medium.com/airbnb-engineering/upgrading-data-warehouse-infrastructure-at-airbnb-a4e18f09b6d5>

# Tools for Building Lakehouses



GA


Dec. 3, 2024

# S3 Tables

Fully Managed Apache Iceberg Tables in S3



**Improved query performance** based on optimized data layout



**Simplified table security** controls




**Automated storage cost optimization** based on compaction, snapshot management and unreferenced file removal

# Amazon S3 Tables architecture





NEW

## S3 Tables

A new type of S3 Bucket specifically designed to store data in **Parquet** files and be used with **Iceberg** format



## Glue Data Catalog



# AWS-managed 開源服務 observability 1.0 架構





# Observability 2.0 架構



Athena



Databricks



Redshift



Clickhouse



# Summary

1. 統一存儲與成本優化
2. 打破數據孤島，實現統一分析
3. 避免鎖定廠商，增強可擴展性

以更低的成本獲得更強的分析能力的同時，為未來的技術演進保持開放性

AWS

# 亞太（台北）區域 正式推出



AWS SUMMIT  
TAIPEI

# 2025 AWS 台灣雲端高峰會

深入探討最新推出的 AWS 亞太（台北）區域相關應用

7月29日（二）- 7月30日（三）| 南港展覽館 2 館

Scott Liao

Facebook @shazi.liao

Linkedin @shazi7804

Joey Wu

Linkedin @ [joey-wu-0208aa60](#)

