Timon: A Timestamped Event Database for Efficient Telemetry Data Processing and Analytics

Wei Cao, **Yusong Gao**, Feifei Li, Sheng Wang, Bingchen Lin, Ke Xu,
Xiaojie Feng, Yucong Wang, Zhenjun Liu, Gejin Zhang
Alibaba Cloud Database Department

2020.06

Alibaba Cloud

Database services at Alibaba Cloud

Cloud Database Marketing:
1st, Asia Pacific

Database Products and Services:
26 Products or Services

Enterprise Users:
100 thousands

Databases Migrated:
400 thousands

As of July 2019 ©Gantner,inc

Expressway for running fully managed databases on Alibaba Cloud.
https://www.youtube.com/watch?v=5VkLDC_uIxM

Alibaba Cloud

Database systems and services at Alibaba and Alibaba Cloud

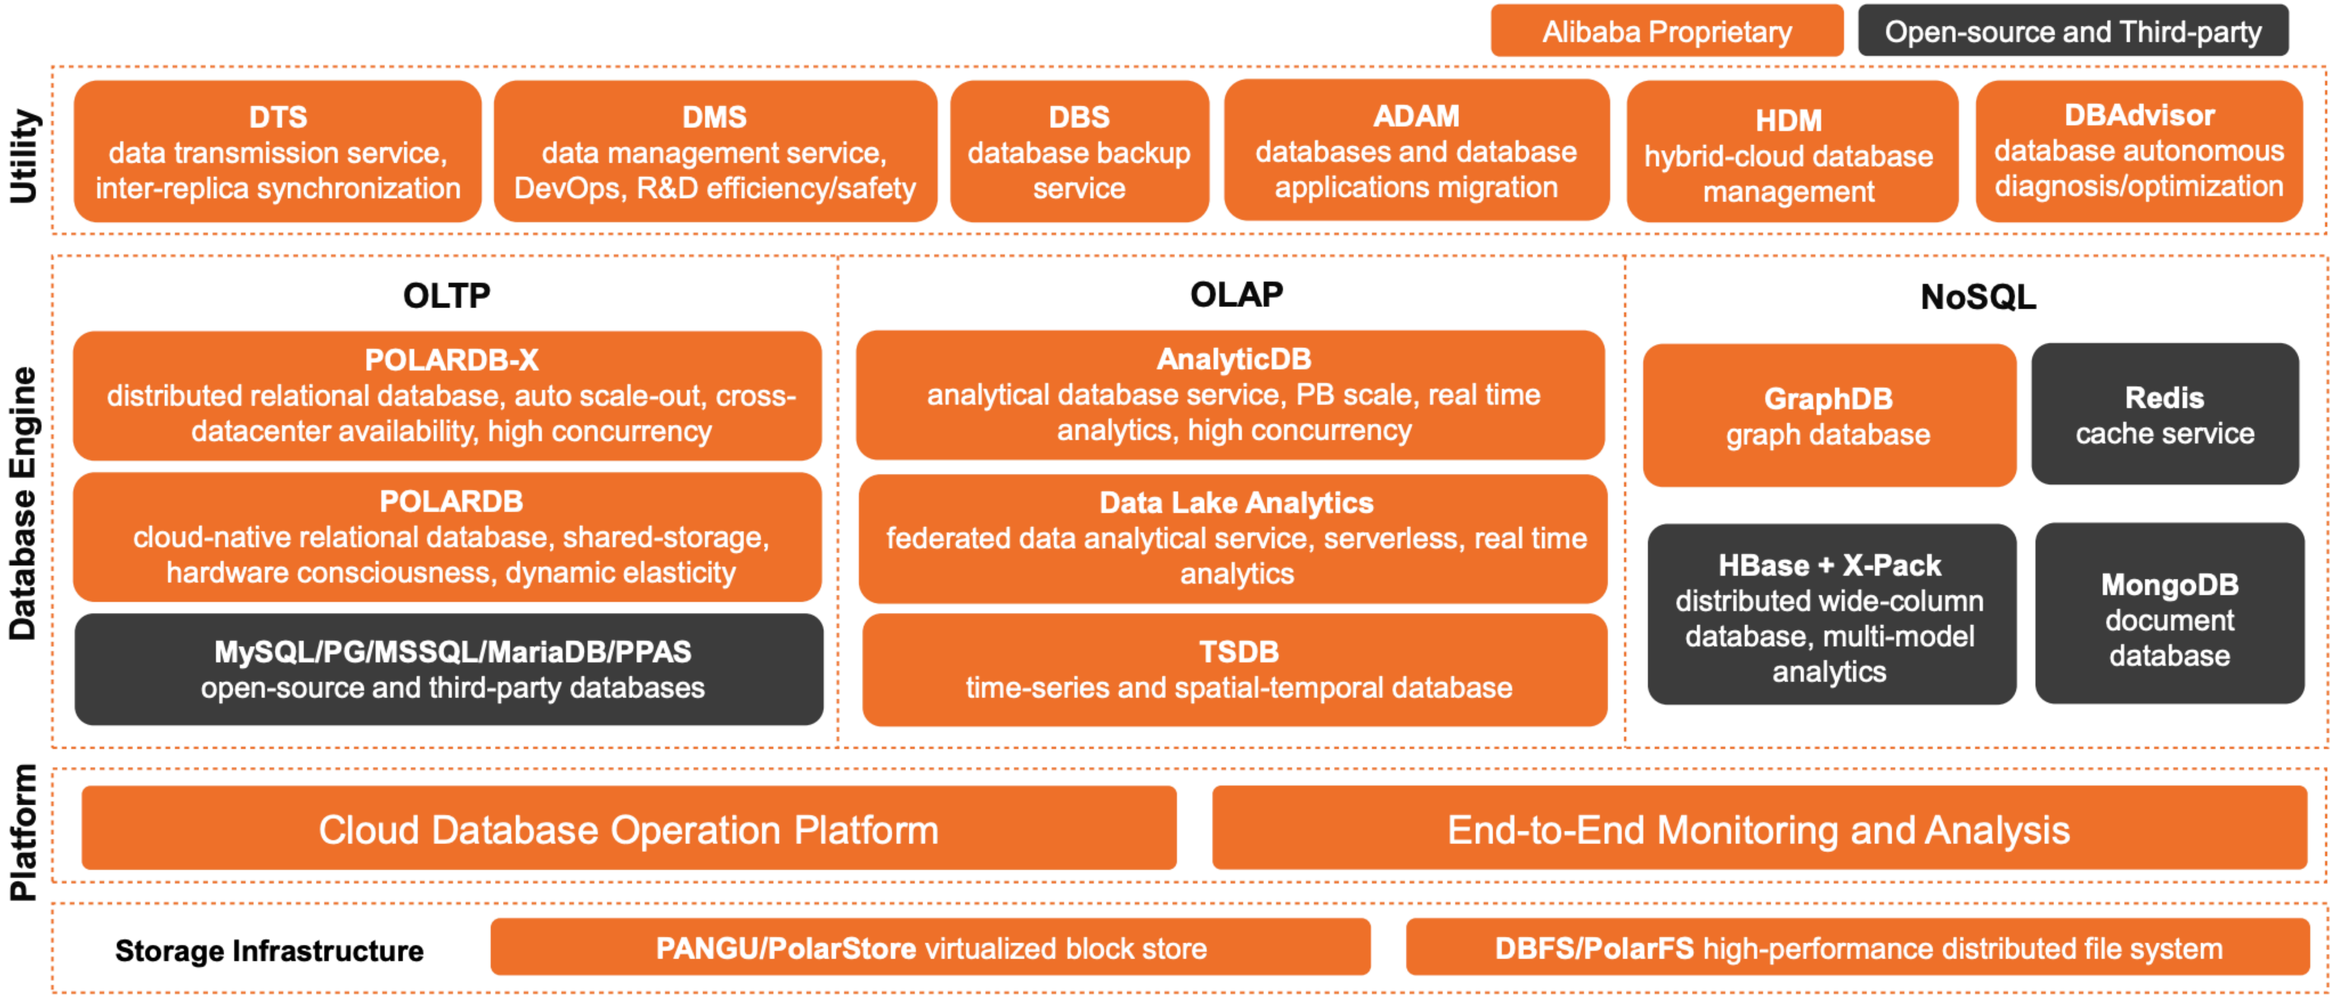

Feifei Li. Cloud-Native Database Systems at Alibaba: Opportunities and Challenges. PVLDB, 12(12): 2263 - 2272, 2019. DOI: https://doi.org/10.14778/3352063.3352141

The scenarios of using telemetry data

End-to-End Tracking System

SIGMOD'18 TcpRT: Instrument and Diagnostic Analysis System for Service Quality of Cloud Databases at Massive Scale in Real-time.

Database Autonomy Service (DAS)

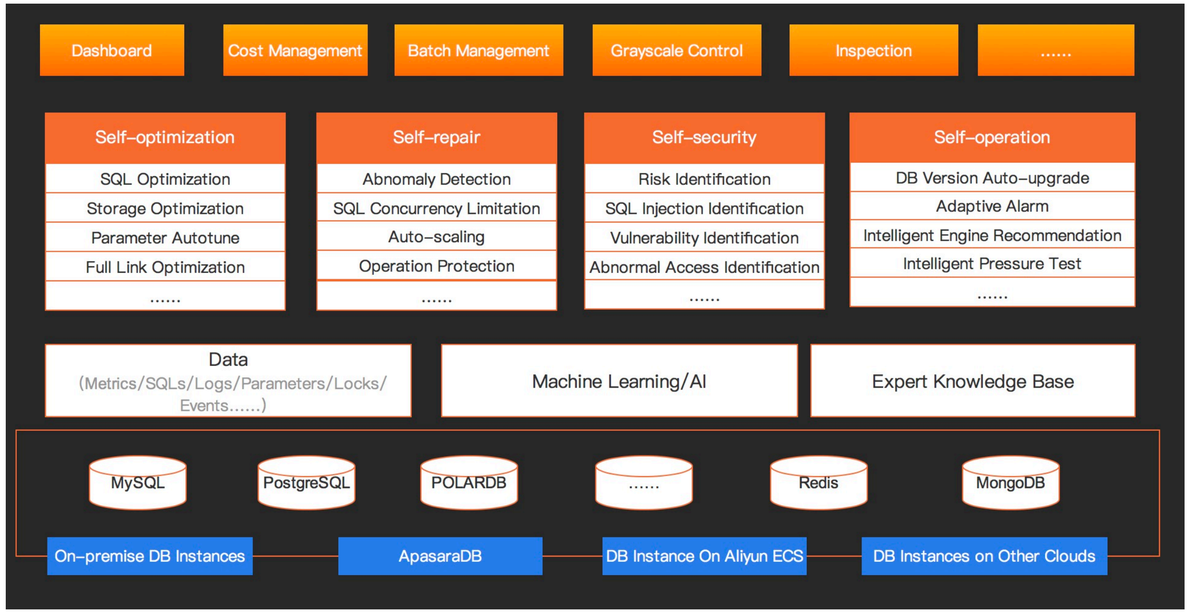

https://www.alibabacloud.com/help/doc-detail/64851.htm?spm=a2c63.p38356.b99.2.61d09bb0Xe1MPU

Alibaba Cloud

Challenge of processing telemetry data in cloud database services

- ~10 million objects
  - Cover database engine, network, operating system, and even each individual OS process.
- Support 1 second granularity and hundreds of millions data points per second
  - Find and explain peaks which last a short time.
- Support long-term & multiple granularities queries with low response latency
  - Find trend and periodicity and compare with historical data
- Strict SLA, realtime & accurate even when out-of-order events exist
  - Find and explain anomalies as soon as possible.
  - Not just for monitoring but also for autonomous optimization.

Alibaba Cloud

Sources of out-of-order events

- Distributed Computing Environment
  - Machine failures
  - Network failures
  - Clock skew
  - etc...

- Application Natures
  - For instance: TcpRT events can't be collected until requests complete. (53.90% out-of-order events)

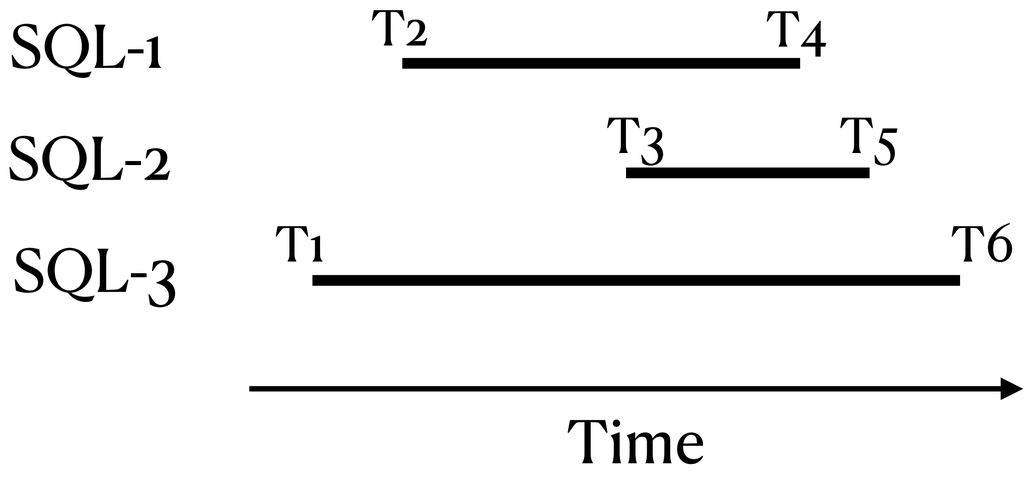

The order of events arrival will be T2, T3, T1

Alibaba Cloud

Our first generation data processing system

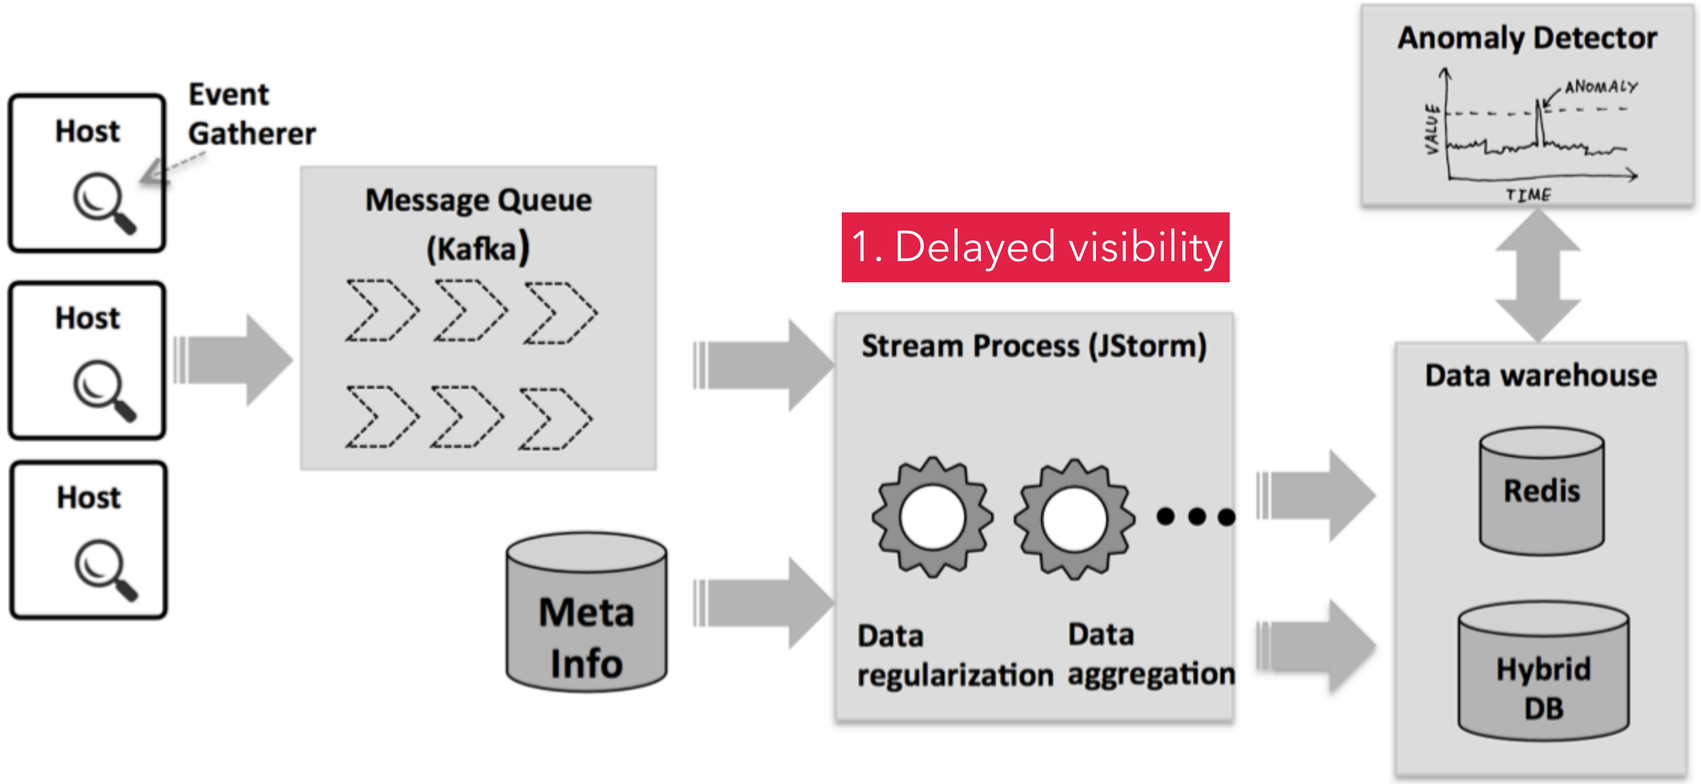

3. It's not good for analysis on massive and long-term events

2. When late data arrives, a Read-Modify-Write operation is required to read and update the previously aggregated value on external storage.

Alibaba Cloud

Solution: using blind-write and moving aggregate to storage

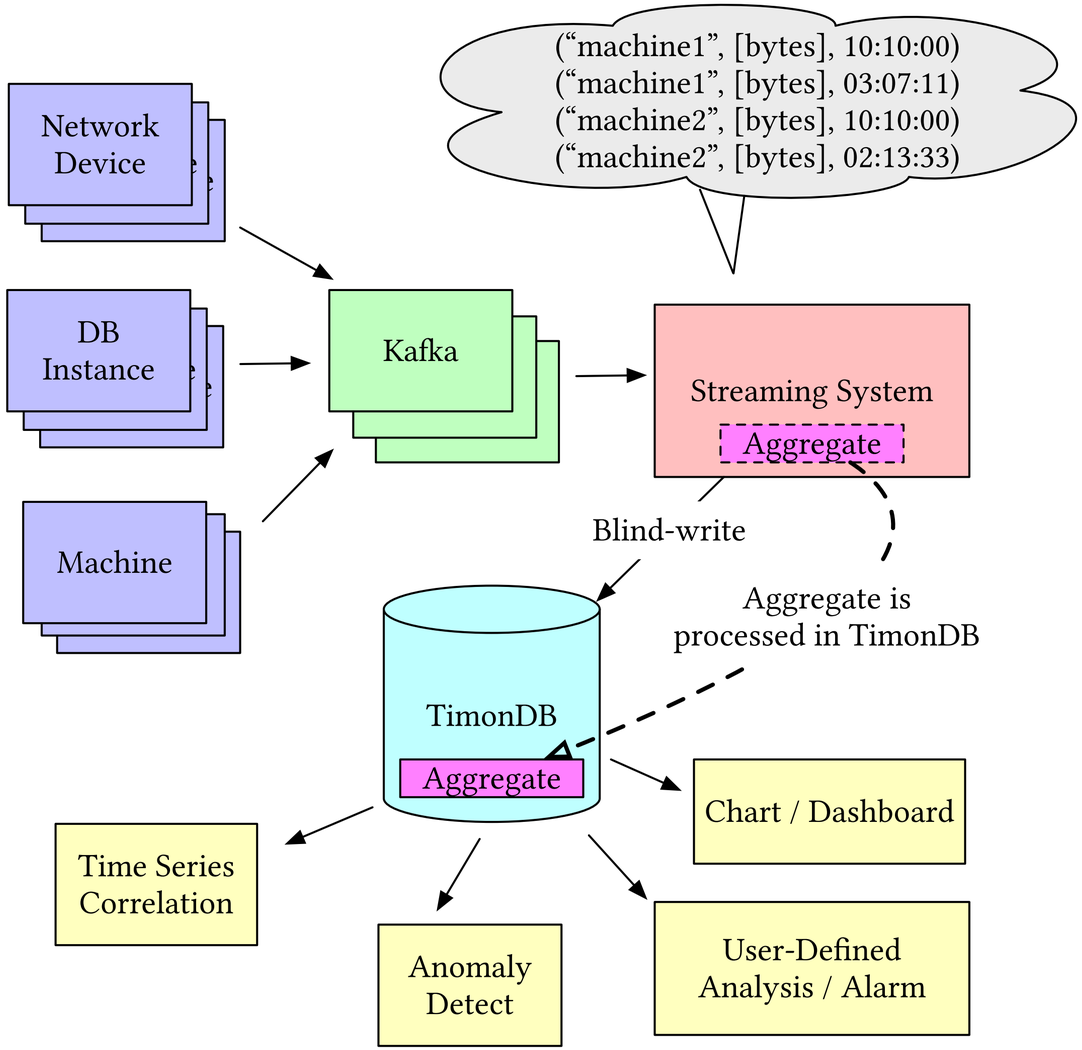

- The most common operators are associative and commutative.
  - **sum**, **max**, **min**, **avg**, **stddev**
  - **quantile**: histogram, t-digest*
  - **distinct**: HyperLogLog
- So it's a good idea to support incremental processing after a data point is written.

* t-digest: https://github.com/tdunning/t-digest

Alibaba Cloud

Incremental Processing Issue: Idempotence

- If failovers occur while the data points are being written, we have to remove the points that have been successfully written.

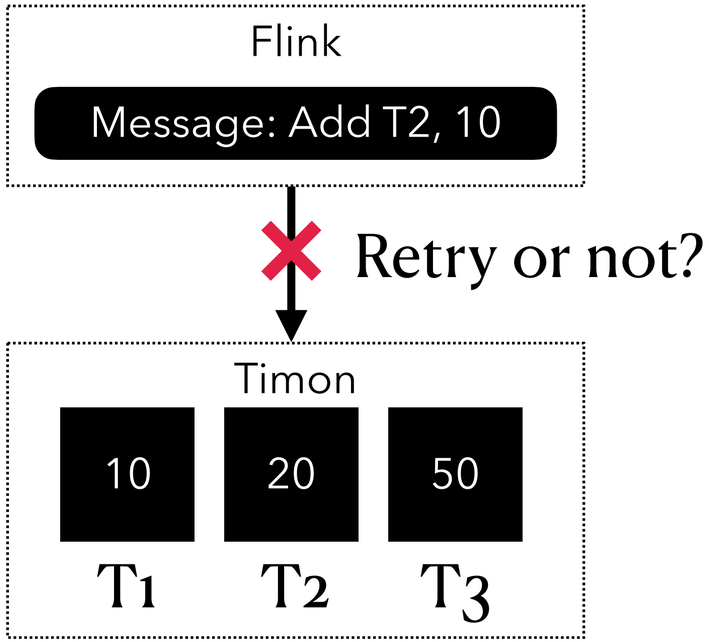

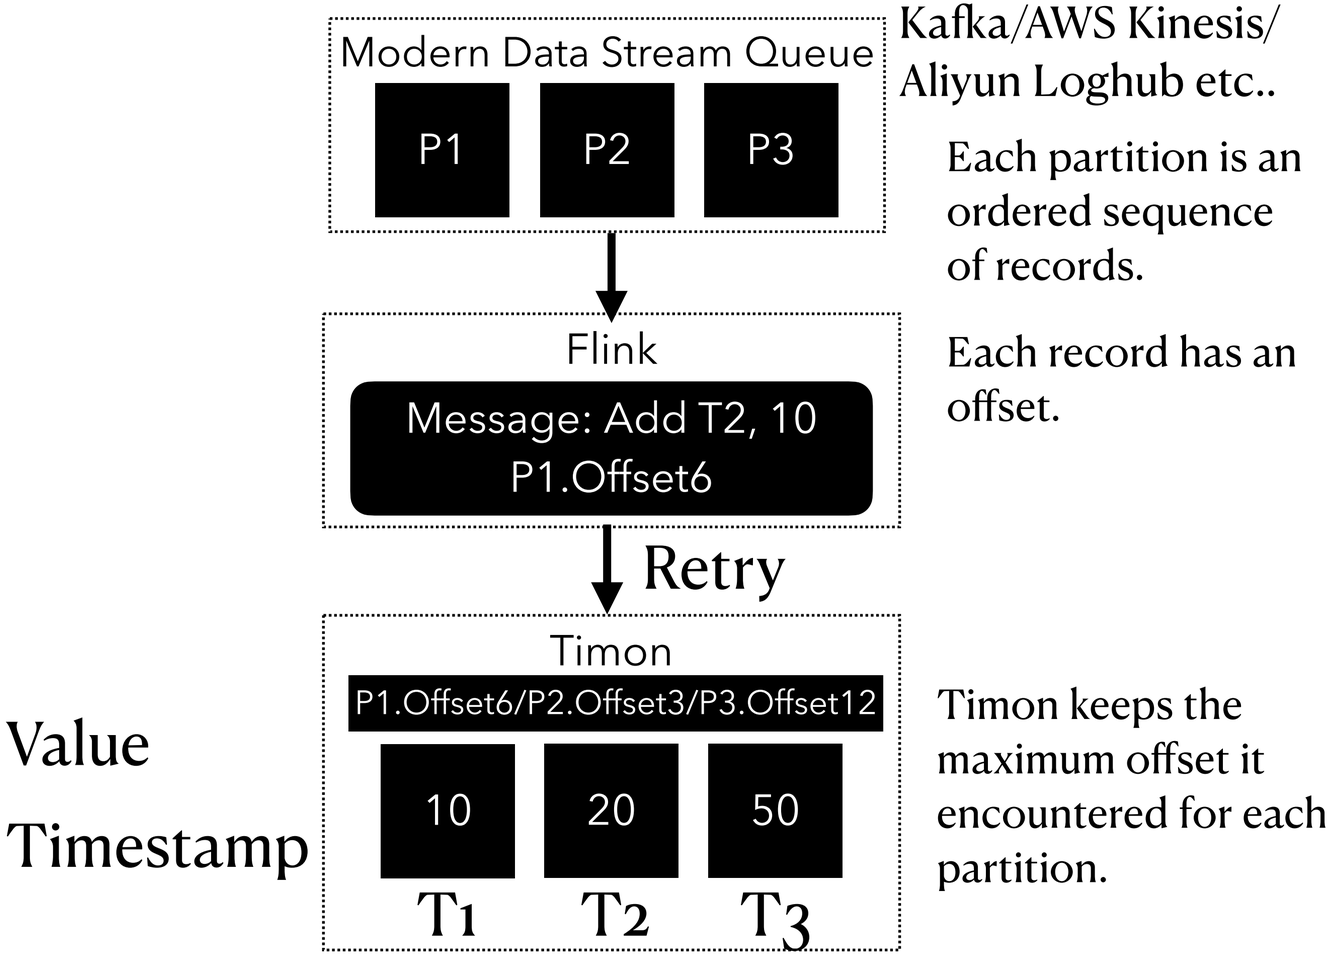

Alibaba Cloud

How to support blind-write and incremental processing efficiently

- LSM-Tree by P O'Neil - 1996

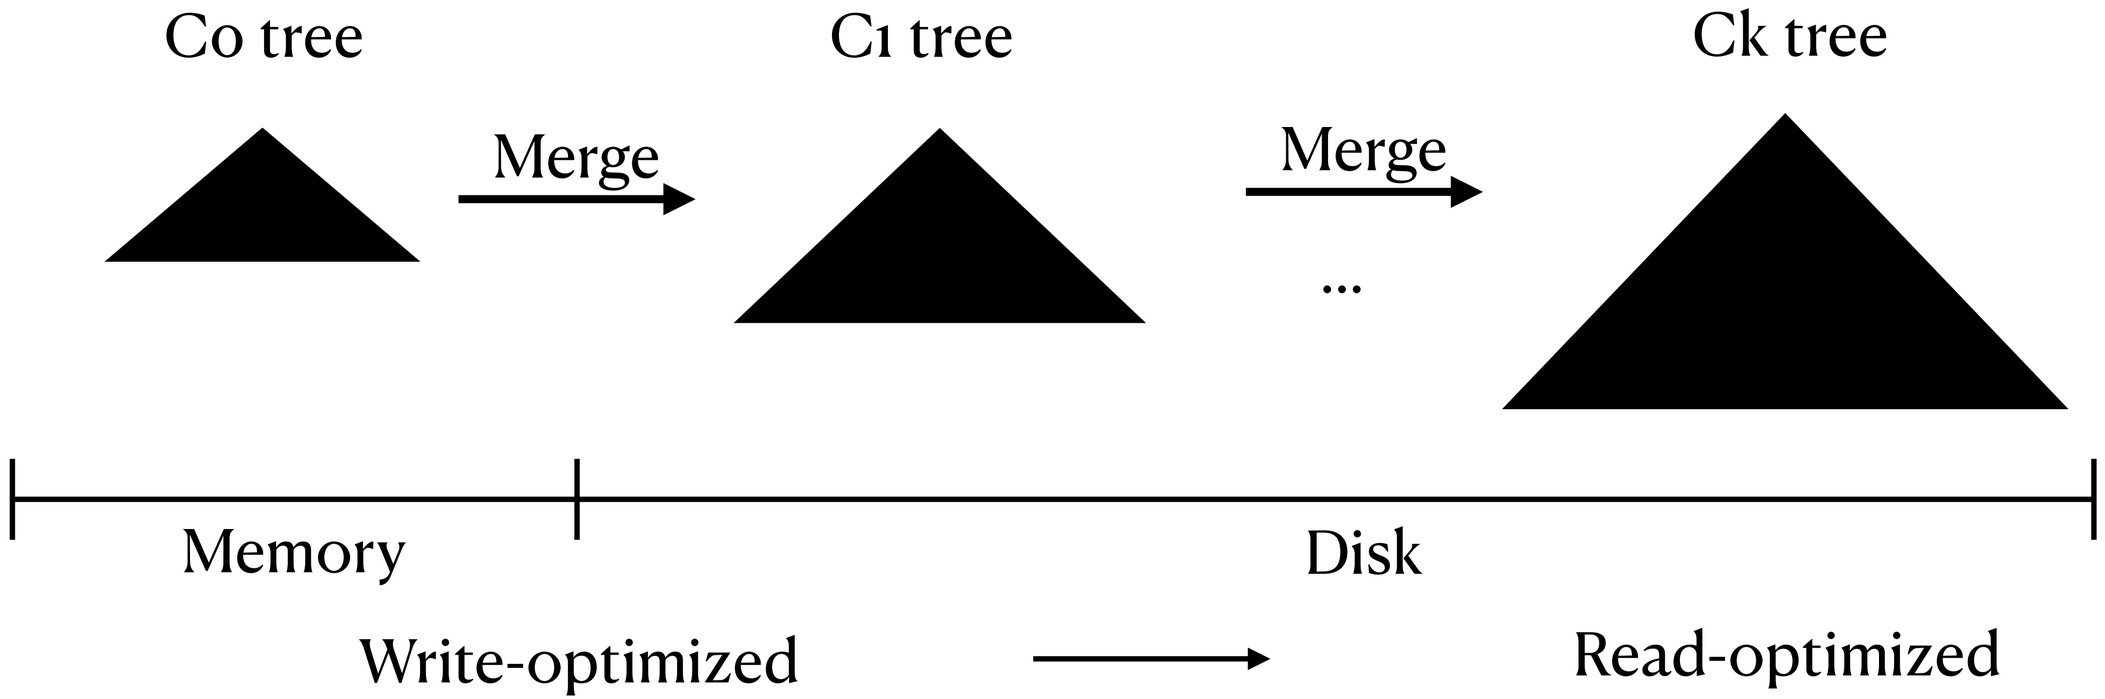

Append & merge to tolerate out-of-order events.

Build time-partitioning tree index for long-term queries.

Alibaba Cloud

Time-Segment Log-Structured Merge-Tree

Lazy Merge

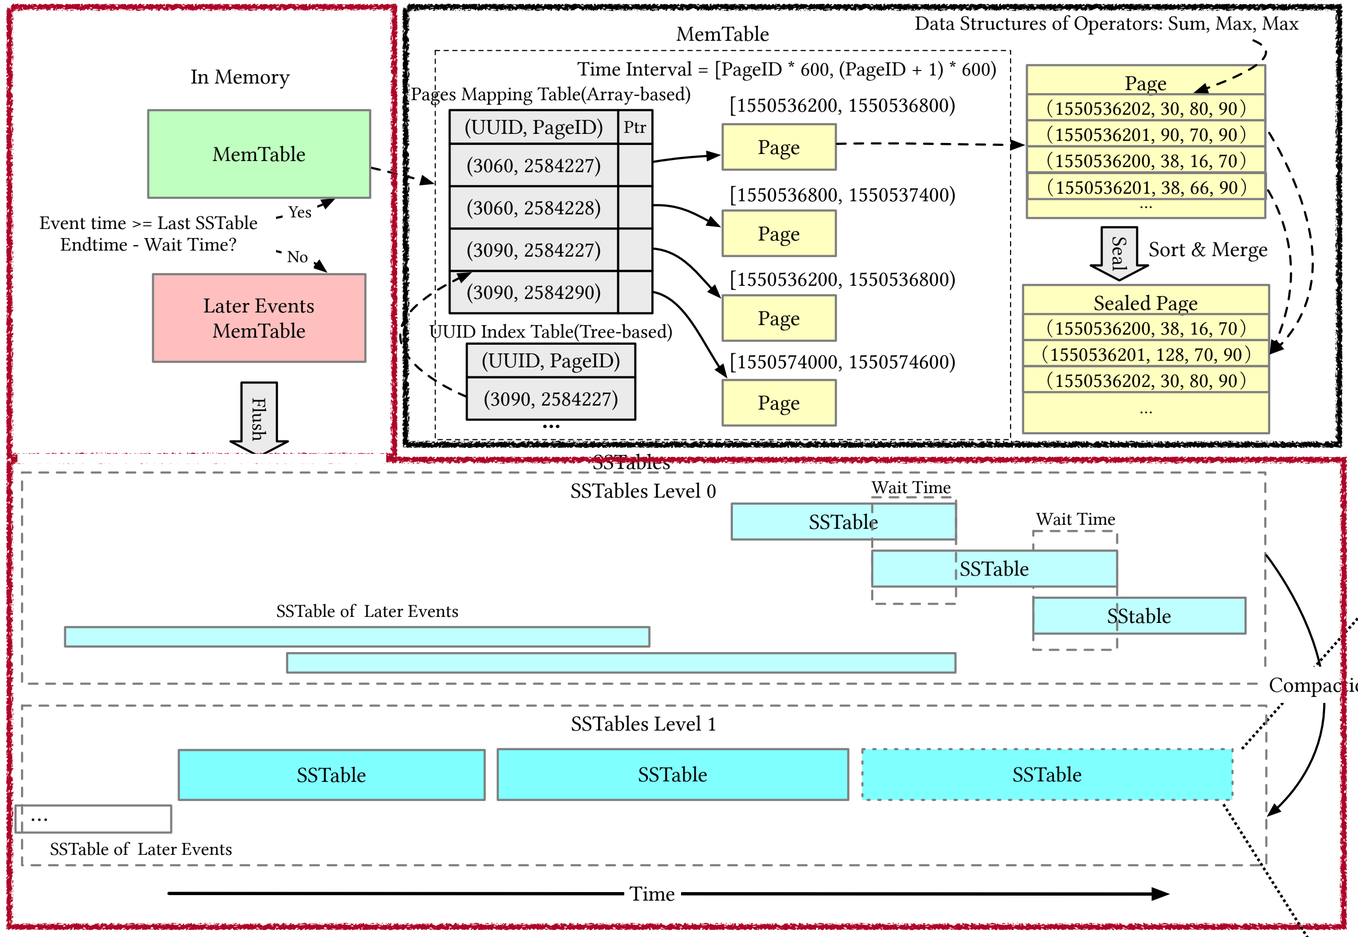

MemTable is optimized for time series data.

Build time-partitioning tree index when compacting.

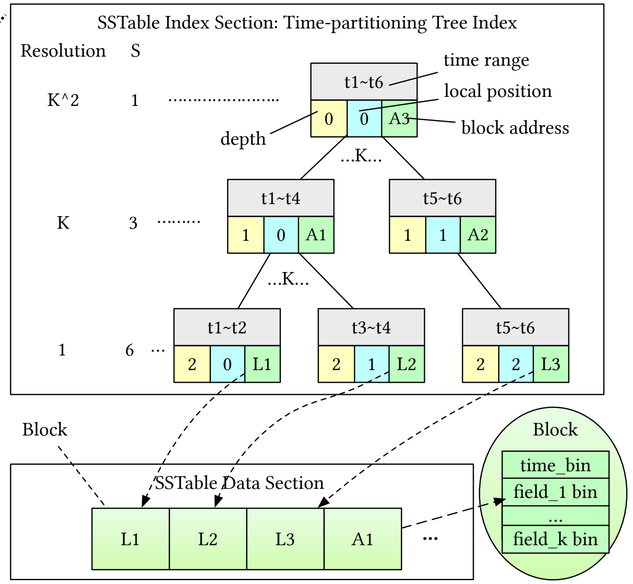

Alibaba Cloud

Optimized MemTable for Time Series

- Hybrid Tree-based Structure and Array-based Structure
- Pages are used to process sparse events.

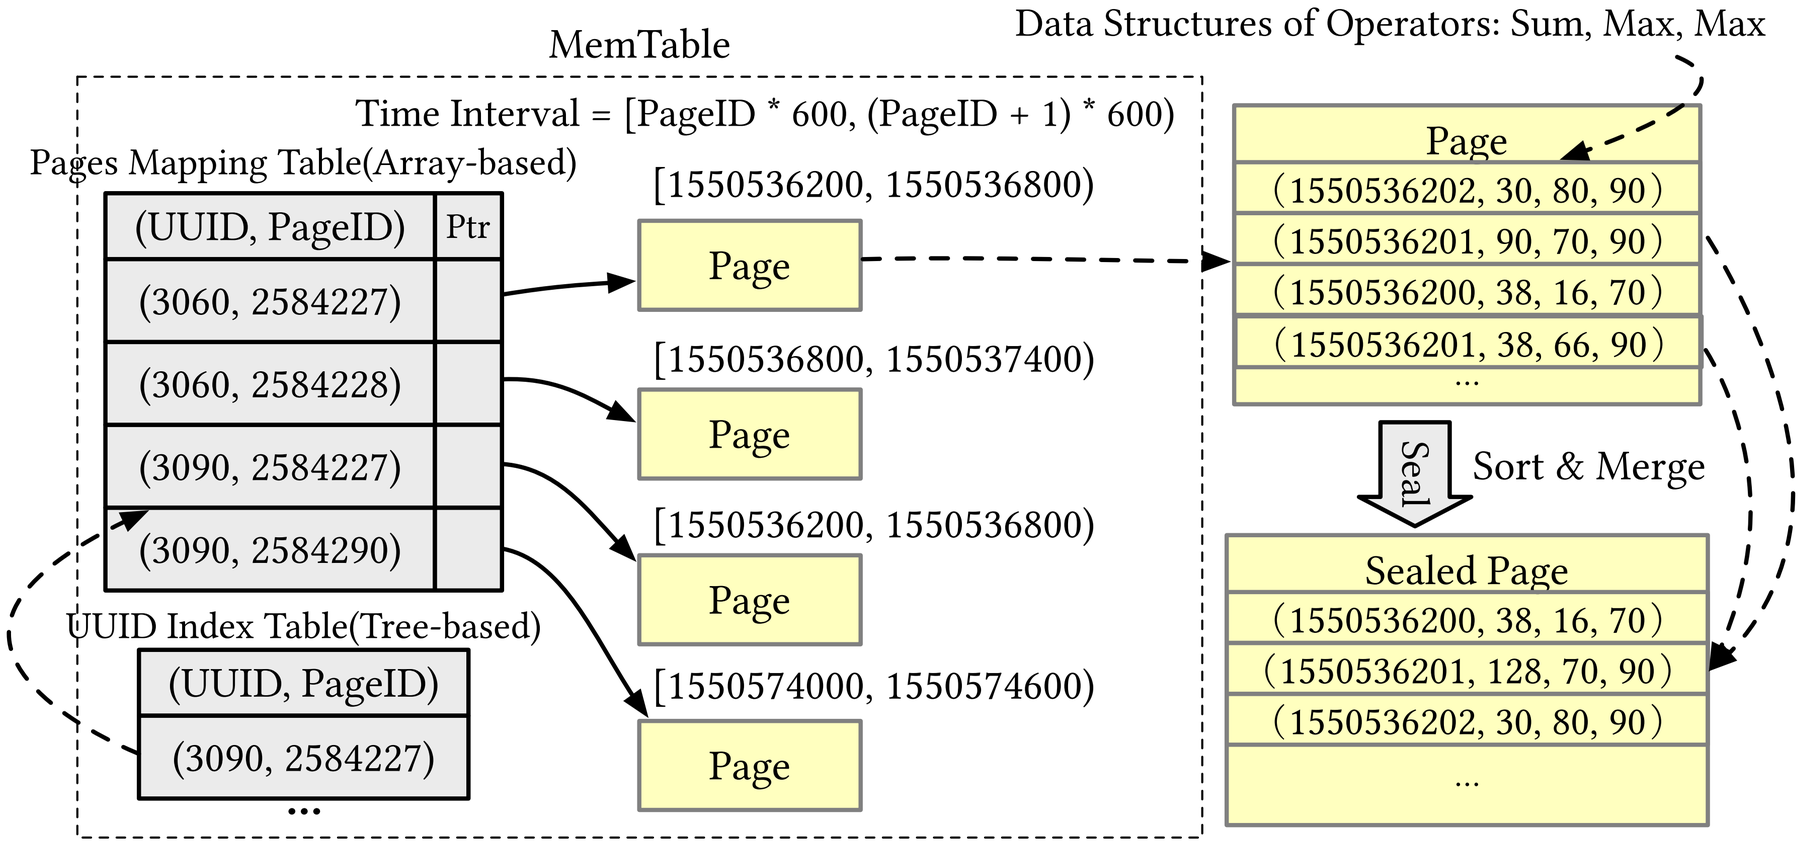

Tolerating out-of-order events (1)

- Most of out-of-order events arrive with a delay of less than 5 minutes.
- Some latecomers are small in proportion but wide in time range.

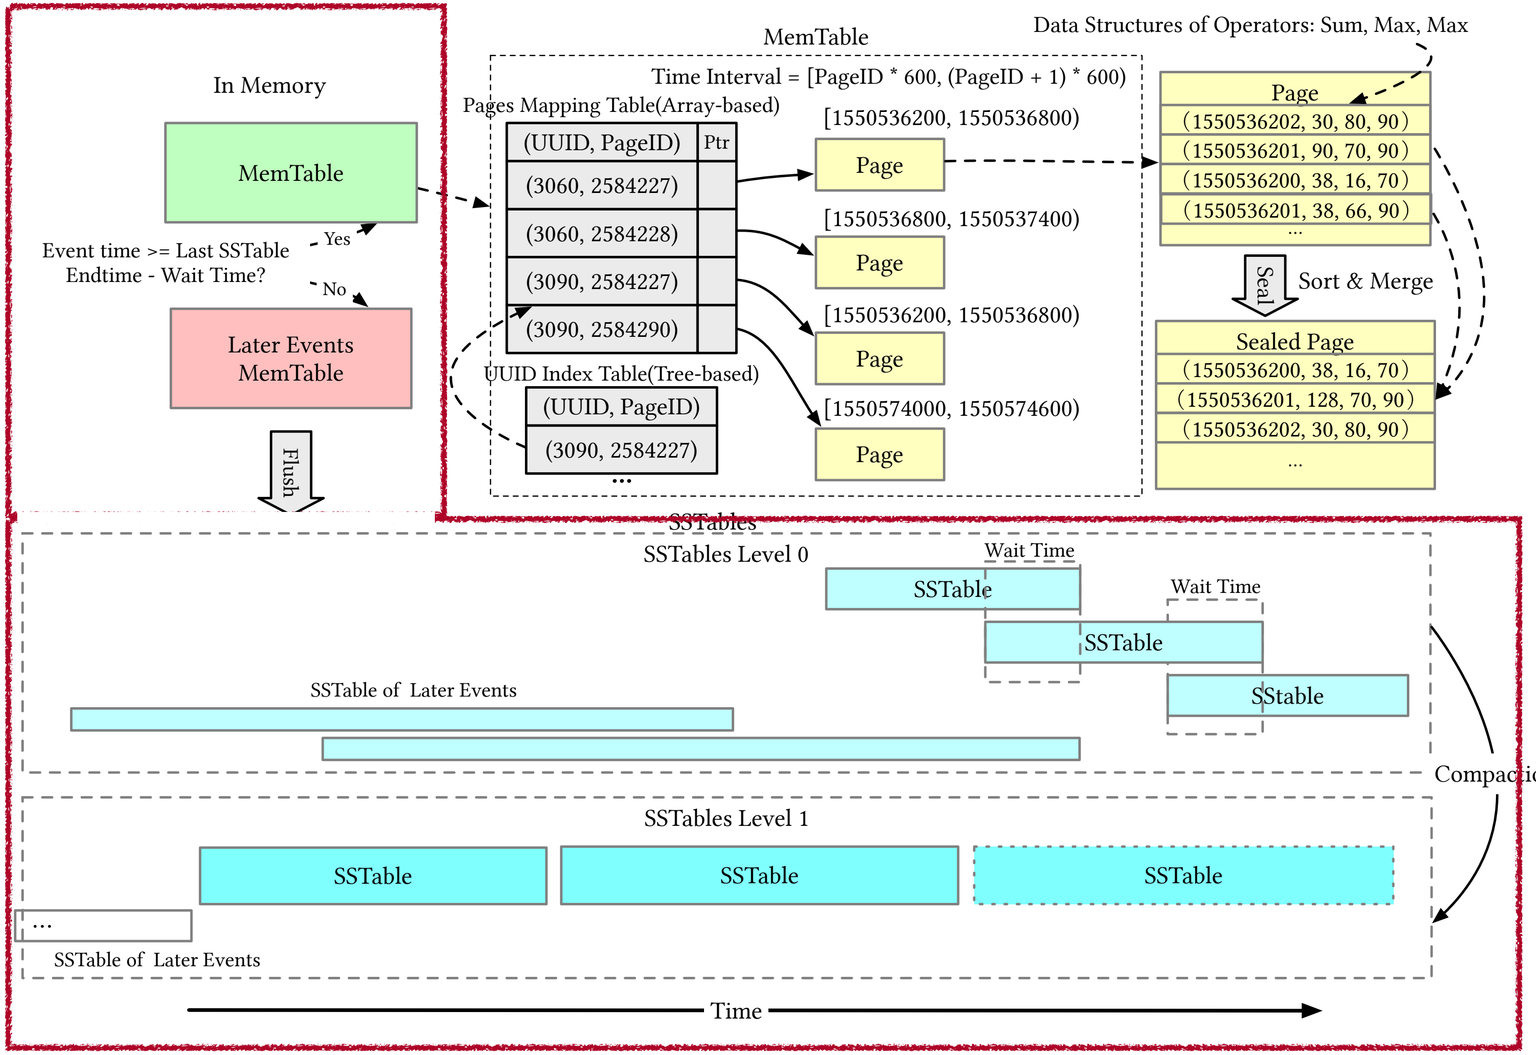

Alibaba Cloud

Tolerating out-of-order events (2)

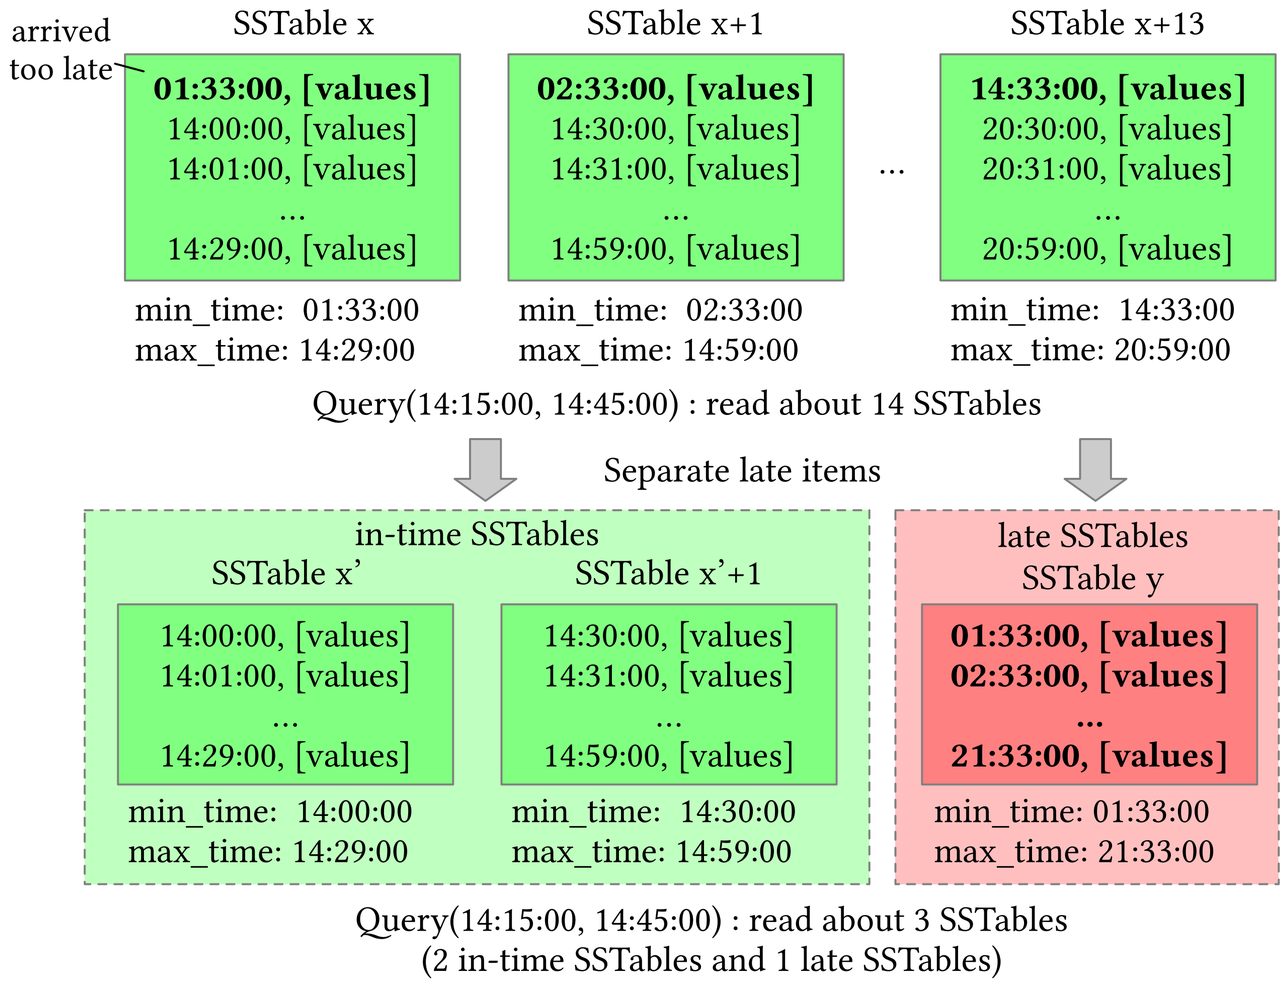

How to process aggregated queries efficiently

- Build time-partitioning tree index by compaction for fast exploration of long-term time-series.
- An aggregated query will scan recent data from MemTable and Lo SSTable, and historical data from time-partitioning tree.

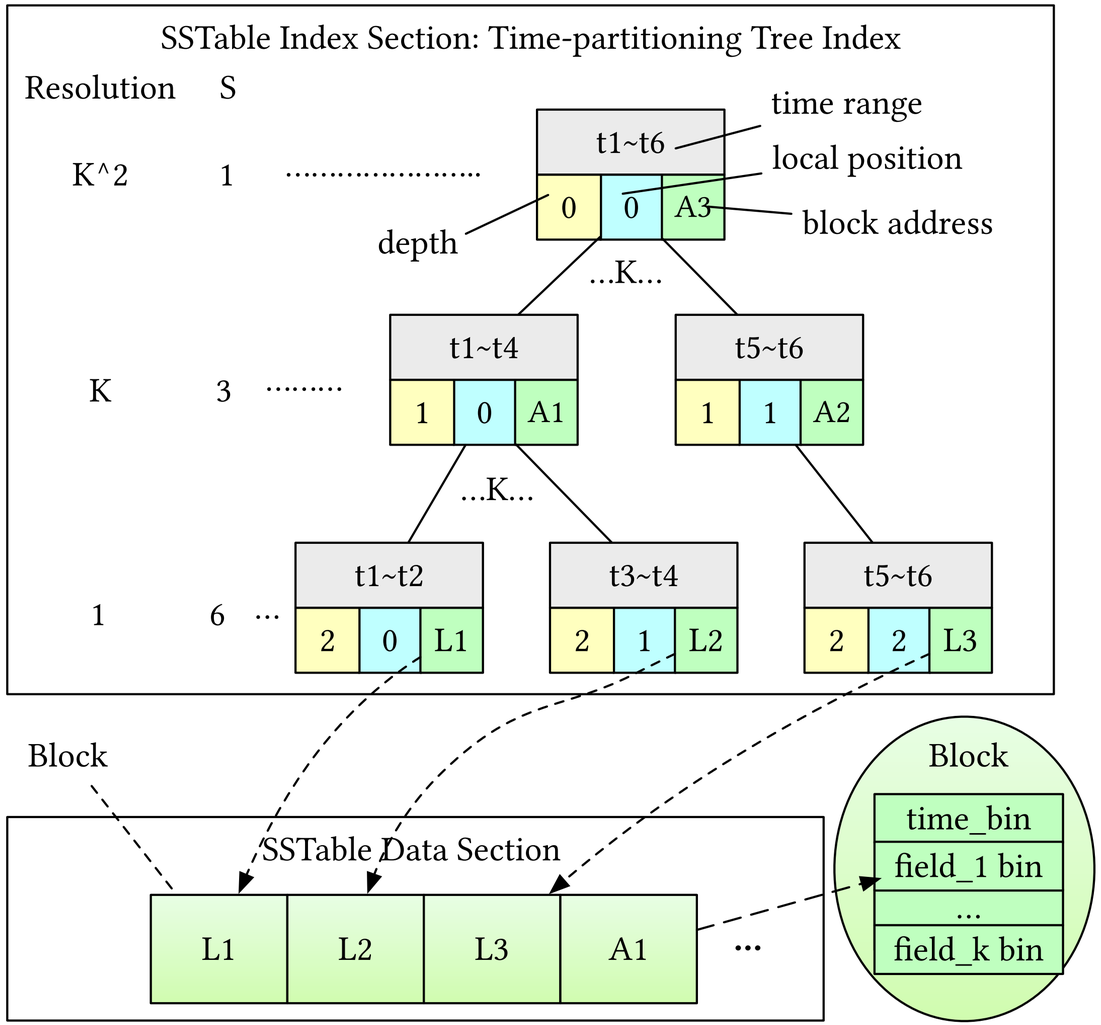

Alibaba Cloud

Timon with user-friendly tools and facilities

- We enhance Timon with user-friendly tools and facilities, such as metric set, materialized view and TQL.
  - Metric set
    - For record which contains dozens of metric values.
  - Materialized view
    - Aggregating data on a higher abstract level, e.g., the region level.
  - TQL
    - SQL-like query languages which allows users to retrieve and analyze the underlying timestamped event data with rich semantics.

Alibaba Cloud

Benchmark

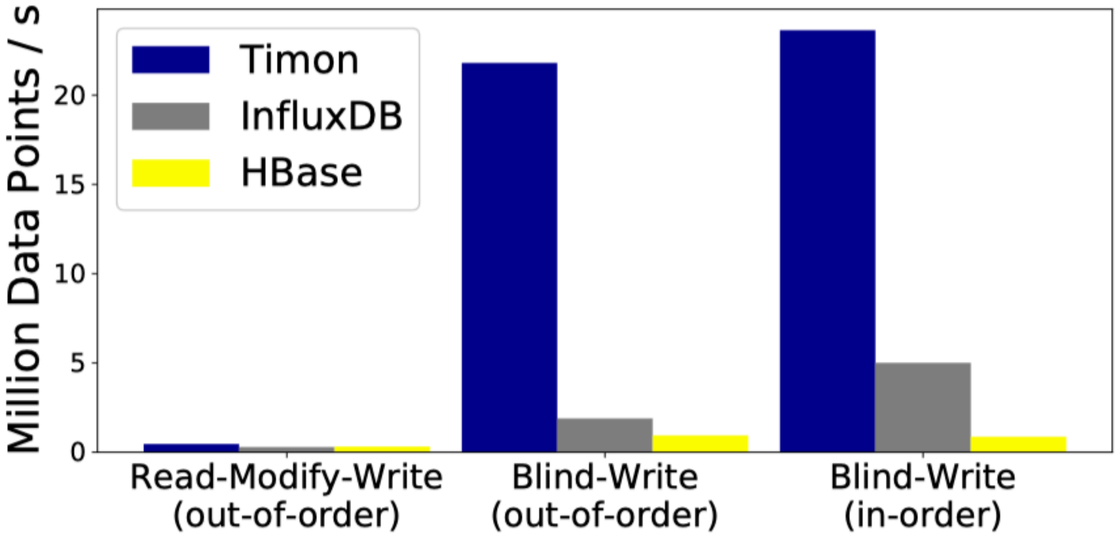

Throughput with different write mode

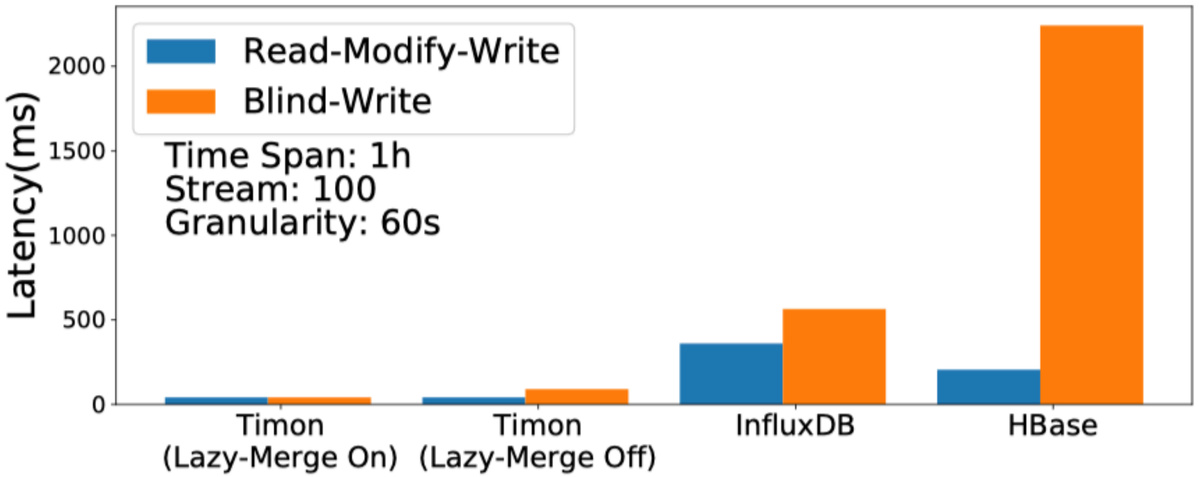

Query latency with different write mode.

Alibaba Cloud

Benchmark: Write Throughput as batch size grow or streams grow

- BtrDB (M. P. Andersen at FAST 2016) is a state-of-the-art TSDB with very high performance and long-term time-series exploration support.
- BtrDB is much better for ultra-high frequency data points (i.e., sub-microsecond precision timestamps).
- Timon maintains high write performance stably when the number of streams increases.
- The main difference is that Timon builds Time-partitioning Tree Indexes by a batch and async procedure, but BtrDB inserts records directly into its tree indexes.

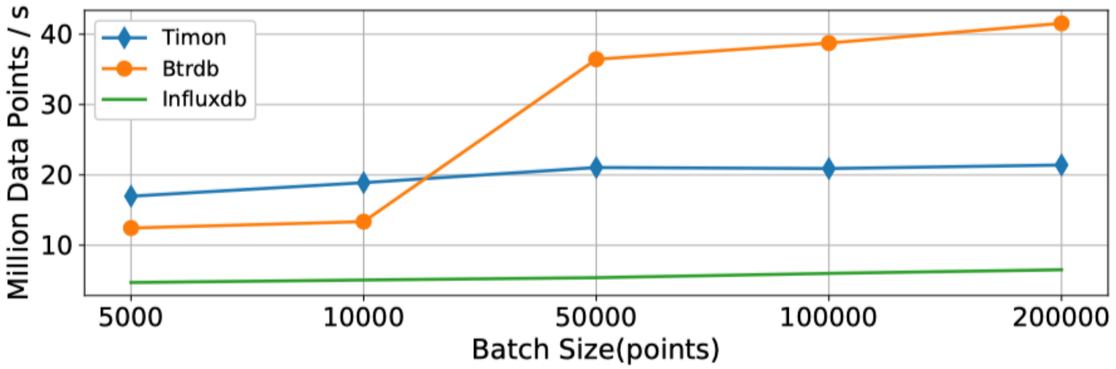

(a) Write Throughput as batch size grow

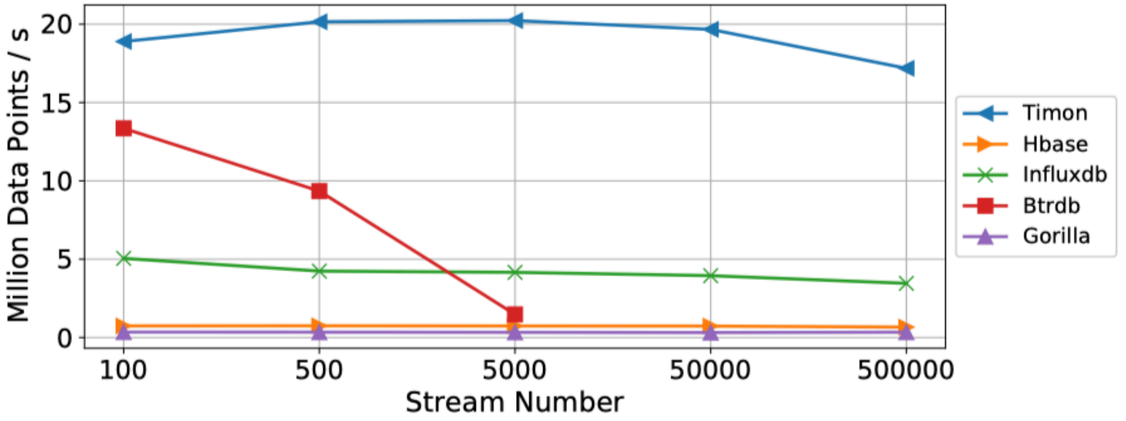

(b) Write Throughput as streams grow

Alibaba Cloud

Benchmark: Query Latency

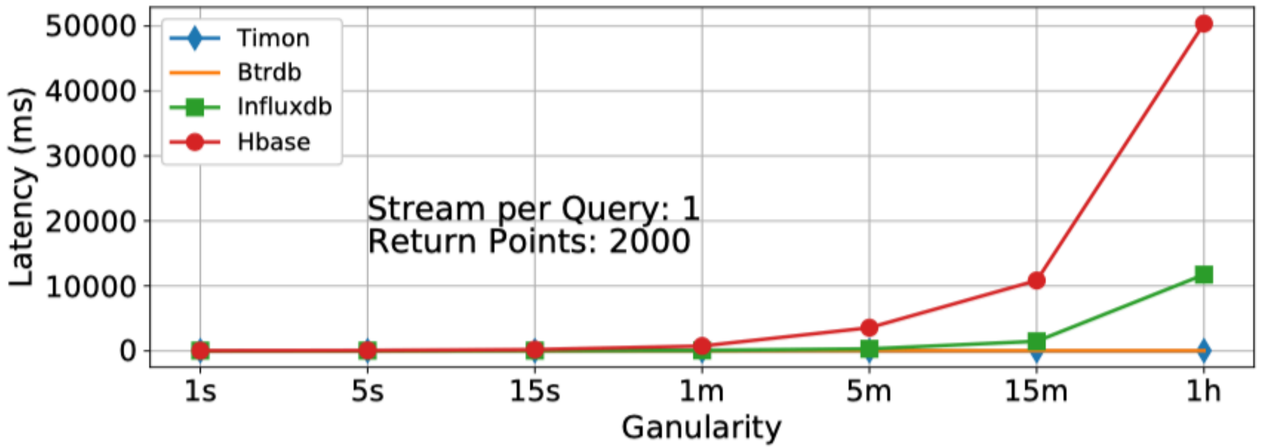

Figure 17: Query latencies between databases as time granularity changes.

Deployment

- Deployed in data centers distributed in 21 regions around the world.
- The biggest application cluster: 97 Timon nodes support about 500 million data points writing per second from the system, and the busiest node serves about 18 million data points per second.

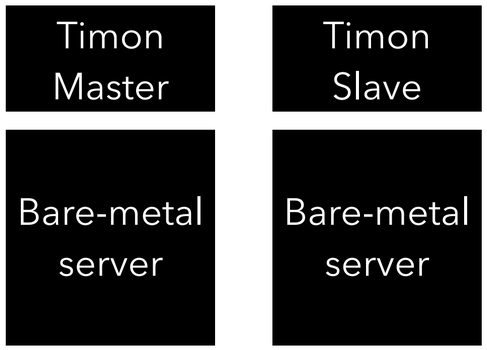

On Bare-metal Server

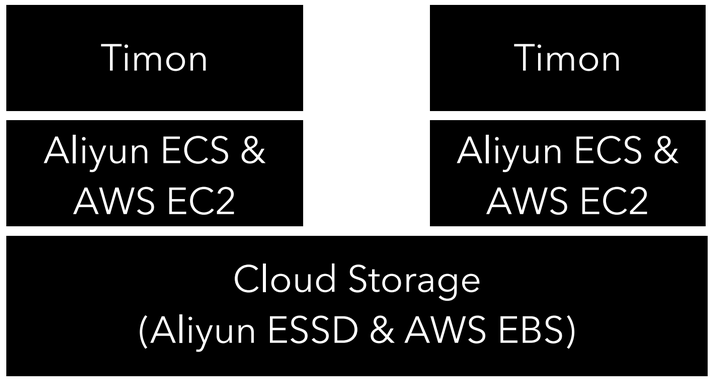

On Elastic Compute

(Each Timon node sees its own data)

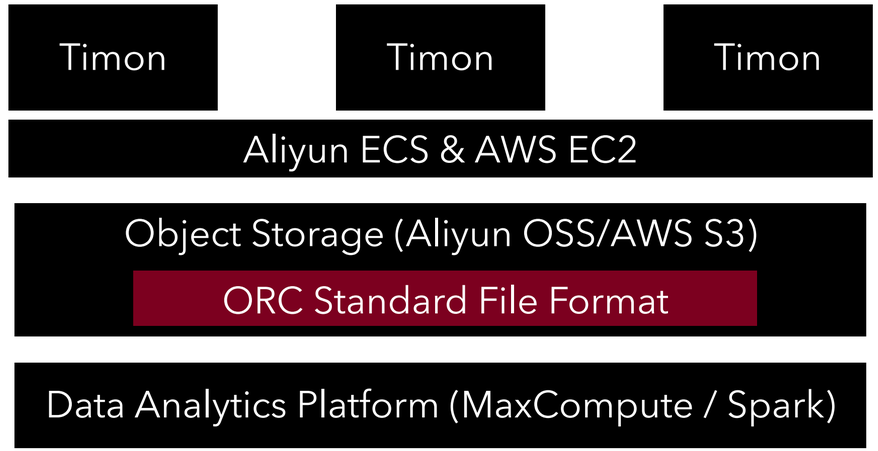

On Elastic Compute

And Using Object Storage

Alibaba Cloud

PolarDB

Thanks

Yusong Gao

jianchuan.gys@alibaba-inc.com