Alibaba Cloud

Towards Millions of Database Transmission Services in the Cloud

Hua Fan, Dachao Fu, Xu Wang, Jiachi Zhang, Chaoji Zuo, Zhengyi Wu, Miao Zhang, Kang Yuan, Xizi Ni, Guocheng Huo, Wenchao Zhou, Feifei Li, Jingren Zhou

Alibaba Cloud

Data Transmission Services (DTS)

Implements logical replication from a *source* to a *target* databases.

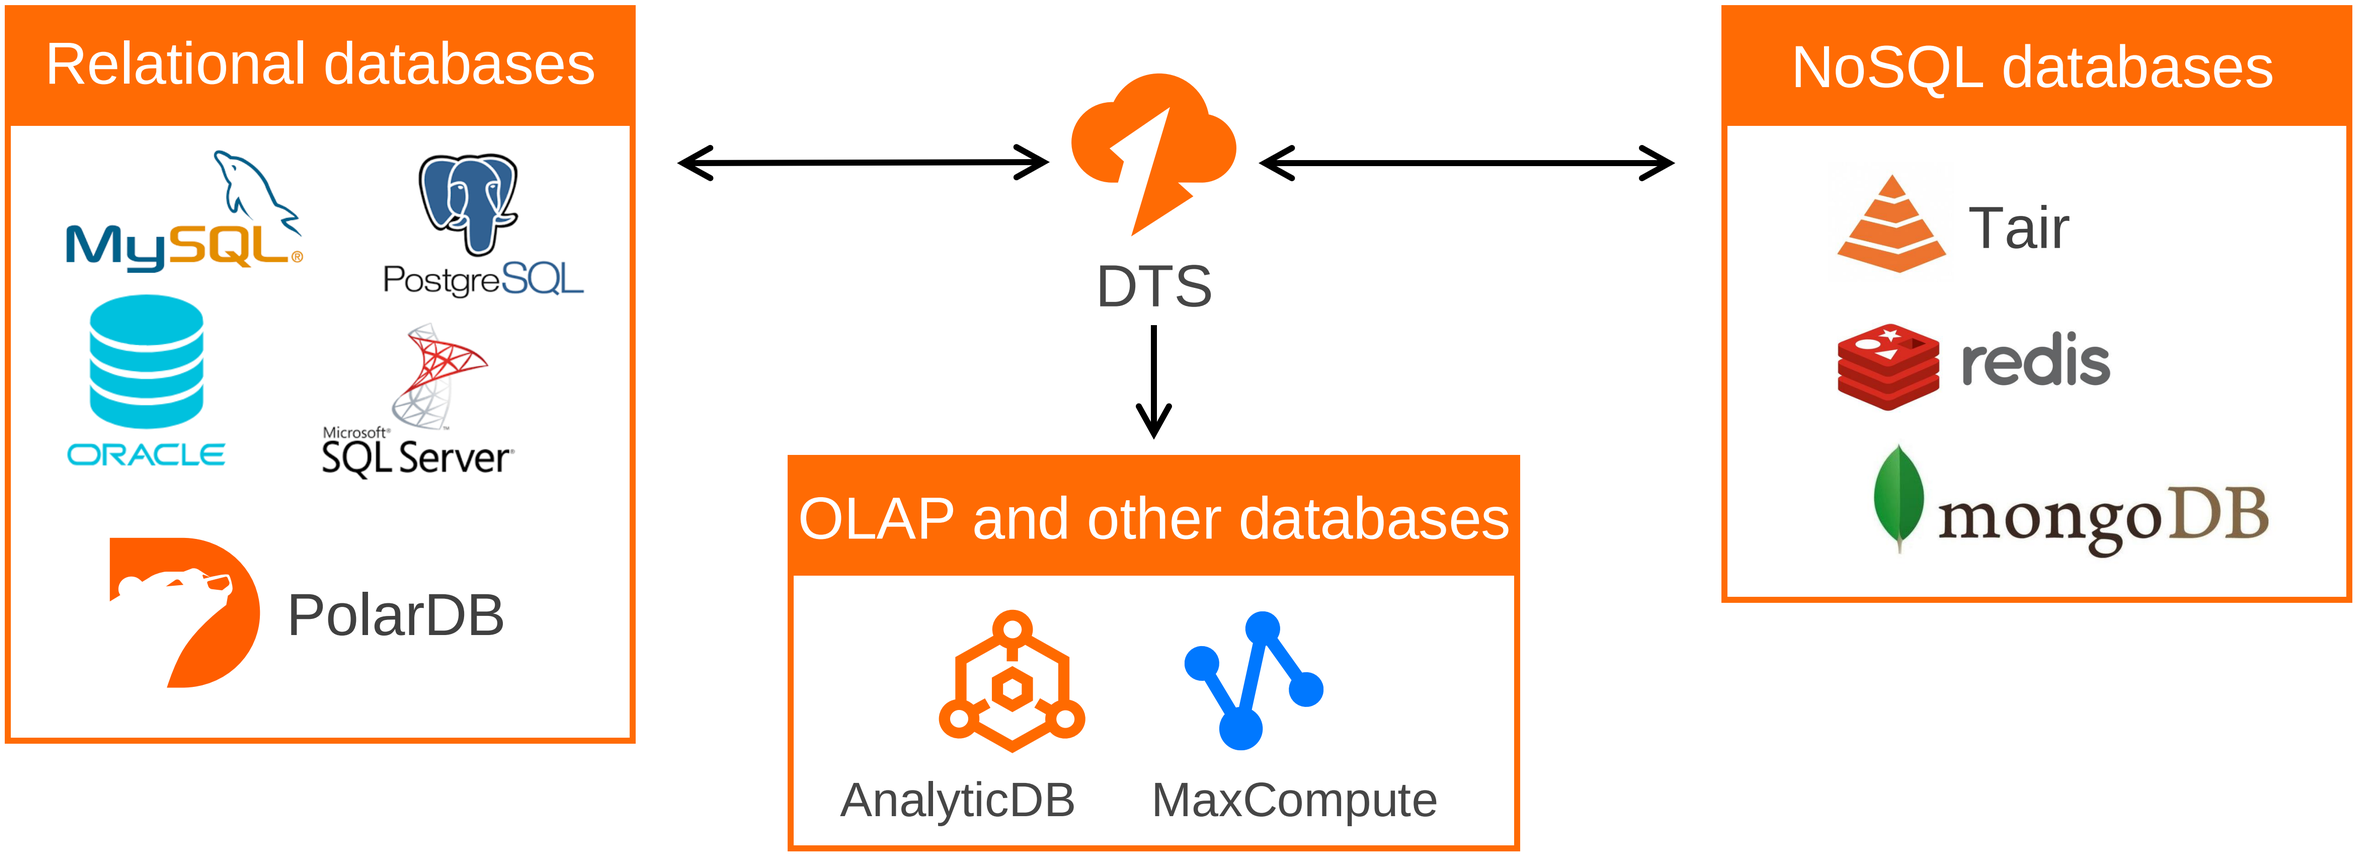

Alibaba Cloud

Data Transmission Services (DTS)

Provides a wide range of transmission services.

Data Migration

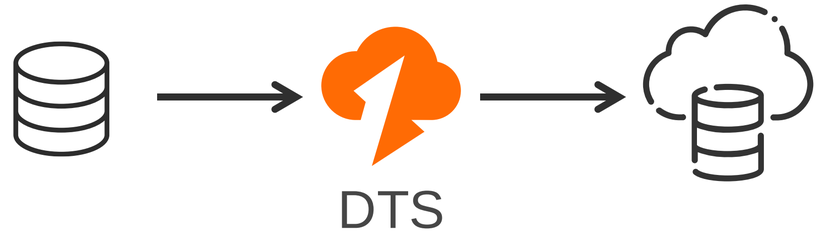

Data Synchronization

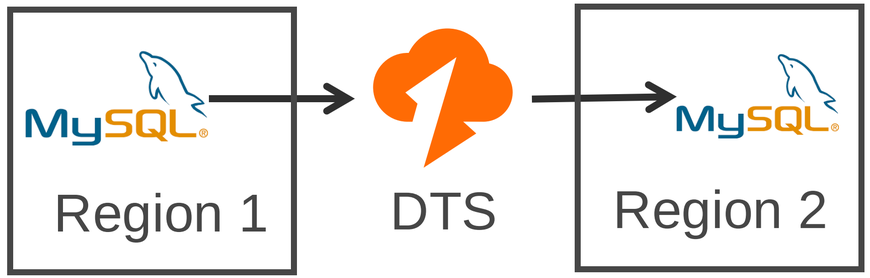

Extract-Transform-Load (ETL)

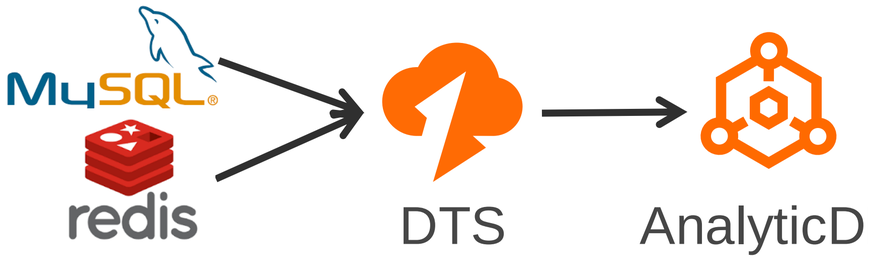

Alibaba Cloud

Towards millions of data transmission

Challenge C1: high database type diversity

M source database types

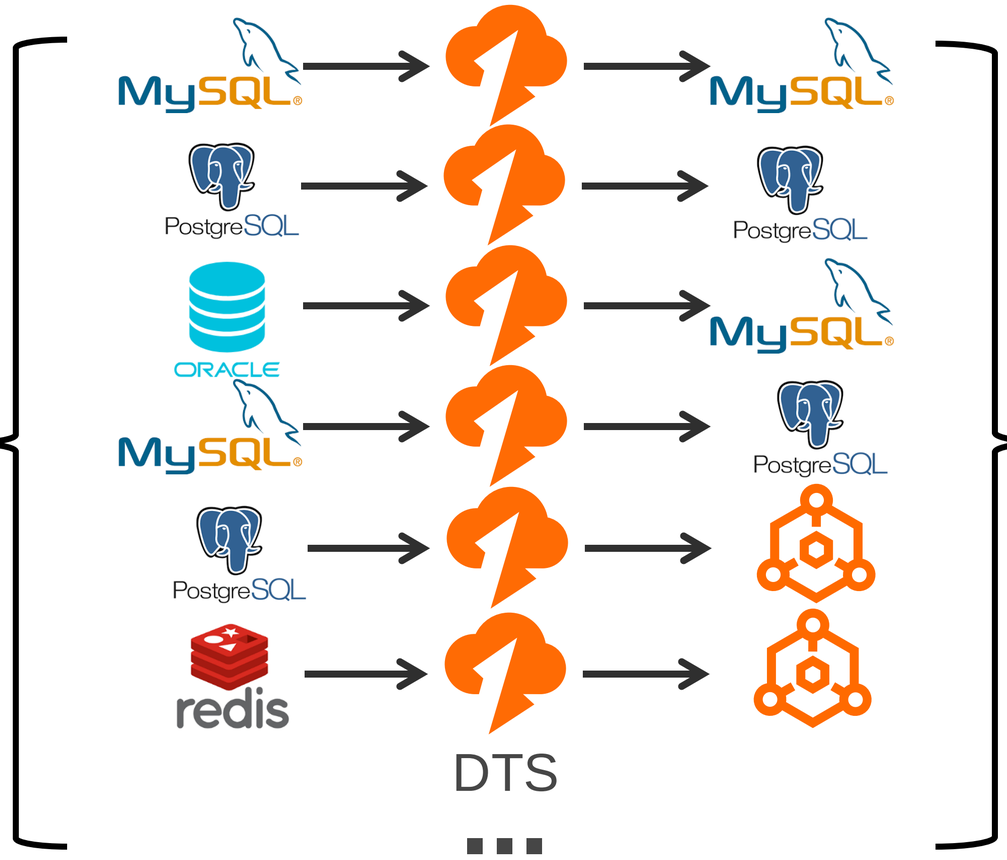

N target database types

Developing for every pair of source-target types requires $M * N$ efforts

Alibaba Cloud

Towards millions of data transmission

Challenge C2: large transmission task number

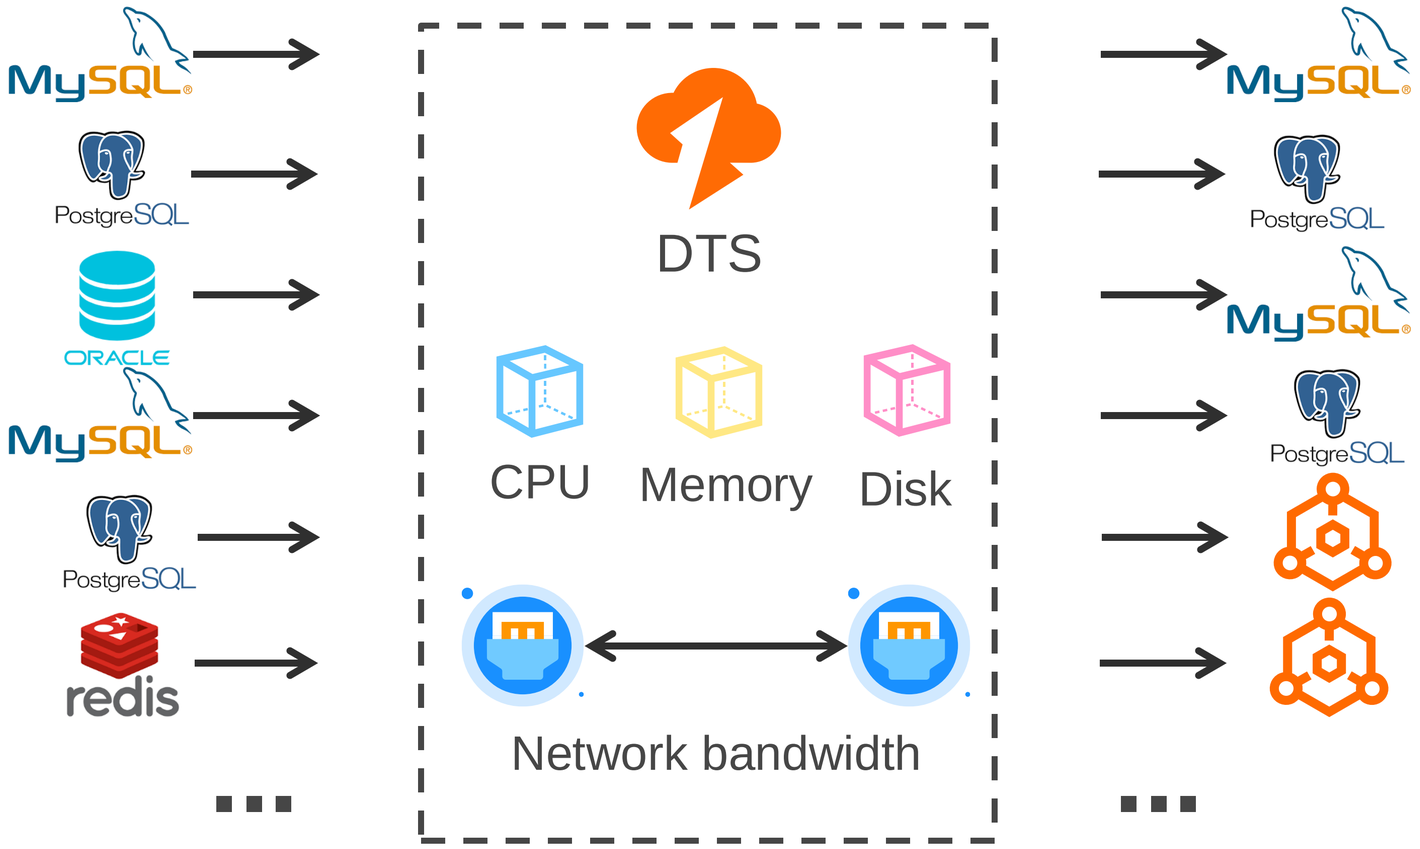

Efficient resource management in the cloud

Alibaba Cloud

Towards millions of data transmission

Challenge C3: High transmission velocity

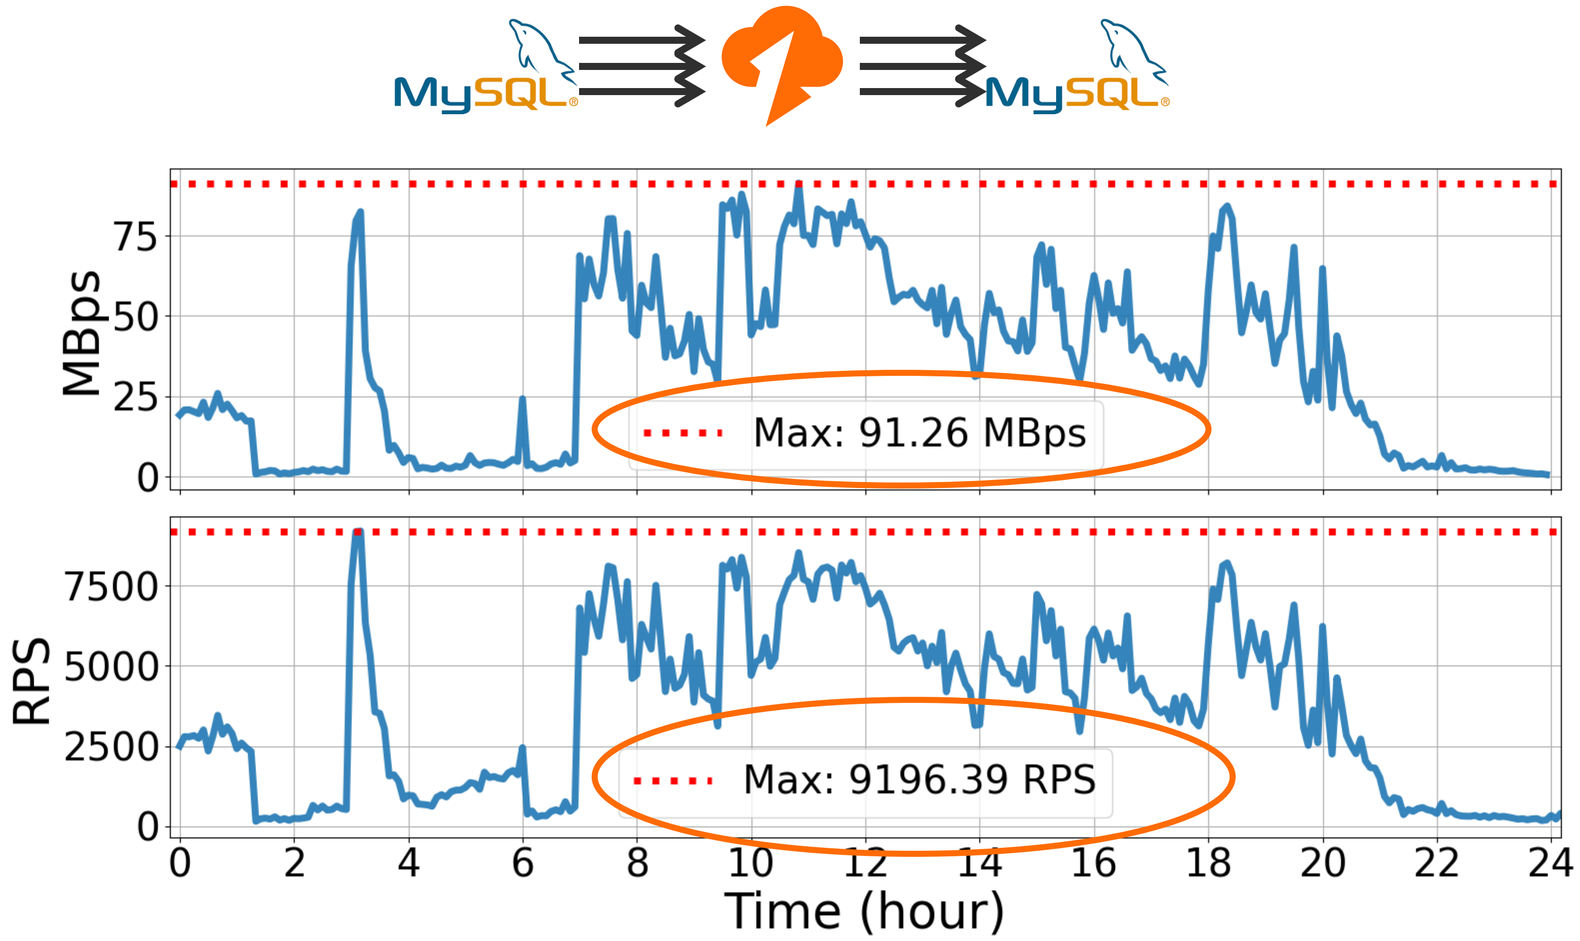

Alibaba Cloud

Contributions

Challenges	Calls for	Contribution
C1: High database type diversity	Reduce development cost	Any-to-Any model
C2: Large transmission task number	Efficient management of cloud resources	DTS serverless Bandwidth scheduler
C3: High transmission velocity	High throughput and low latency	Sink optimization

Any-to-Any (A2A) model

General Internal Representation (GIR): data from all database types are transformed into a standard format.

Record: schema, operation, primary key, before image and after images

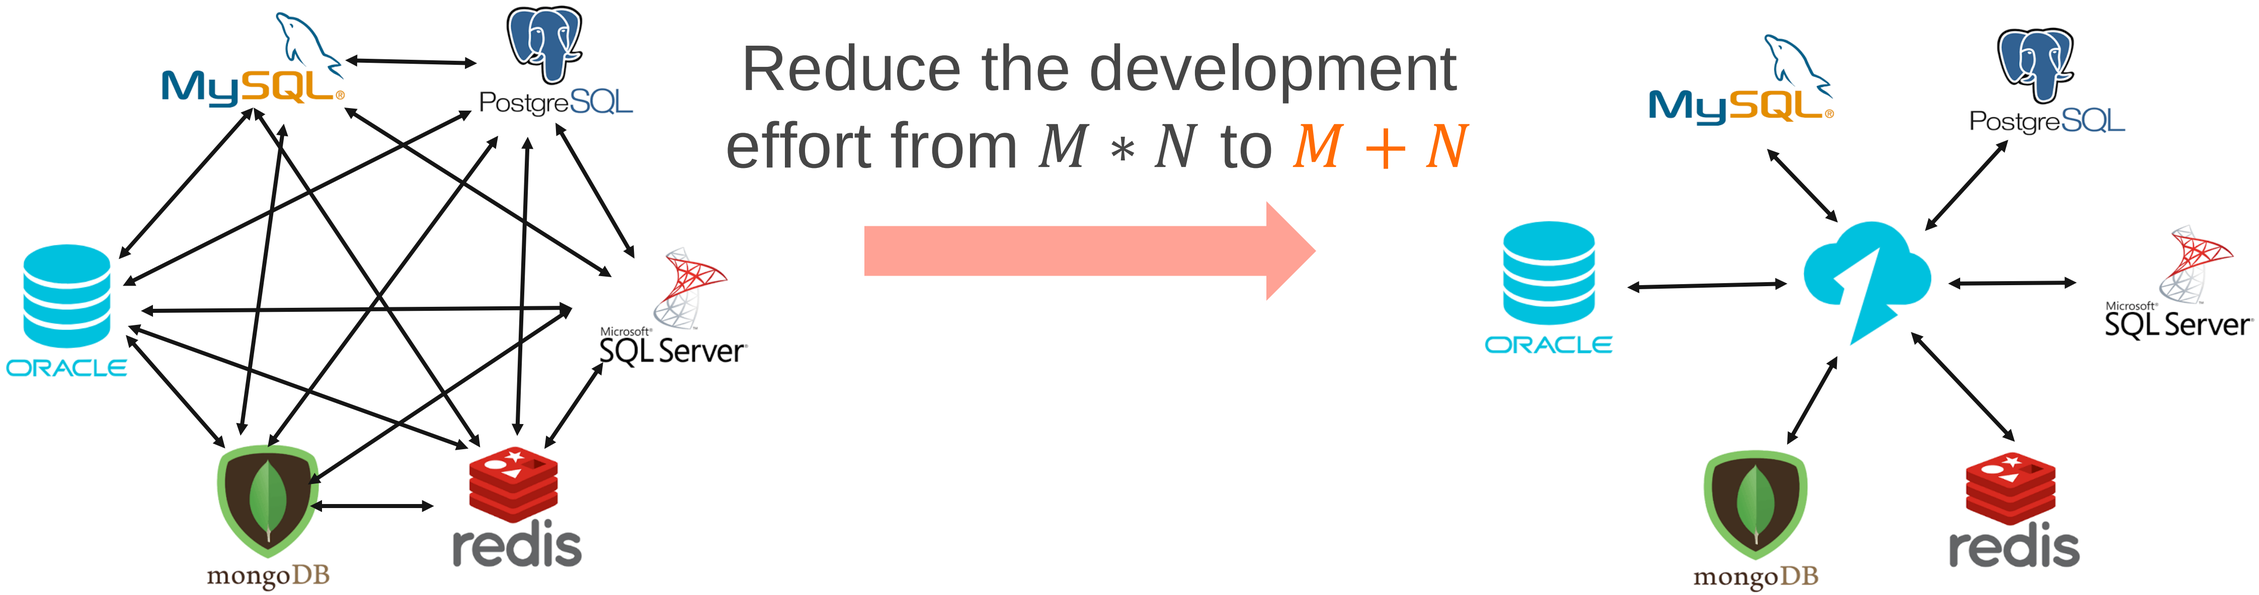

Any-to-Any (A2A) model

Source

Extract data from source DB

- Full Source: extract from tables
- Schema Source: retrieves DDL operations from logs
- CDC Source: retrieve DML operations from logs
- Supports various plugins

Tunnel

Pub-Sub system that bridges source and sink

- Persists data in general internal representation

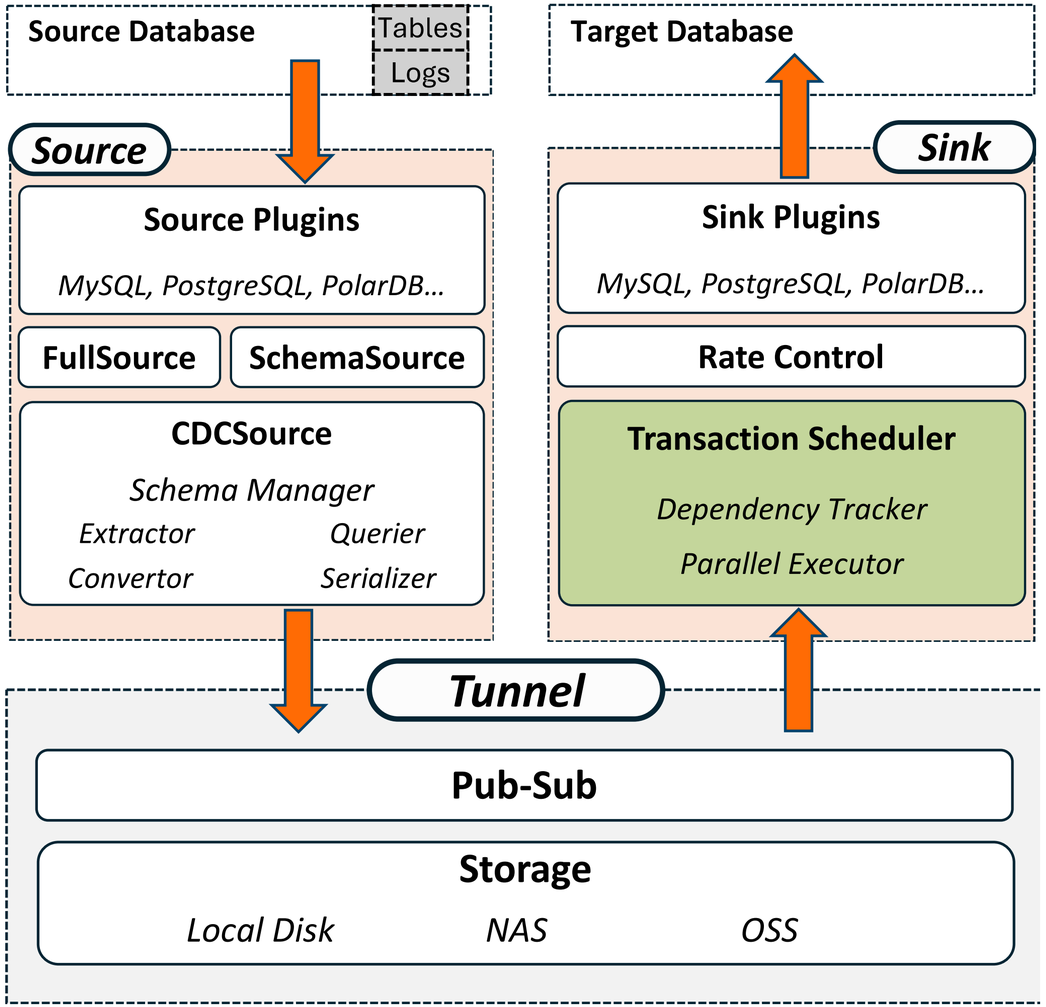

Sink

Applies data changes to the target DB

- Transaction Scheduler
  - Dependency tracking
  - Hotspot consolidation
  - Parallel execution
- Supports various plugins

Sink optimization

Consistency: DTS ensures eventual consistency for real time tasks, which allows non-conflicting transactions to be executed in parallel & out of order.

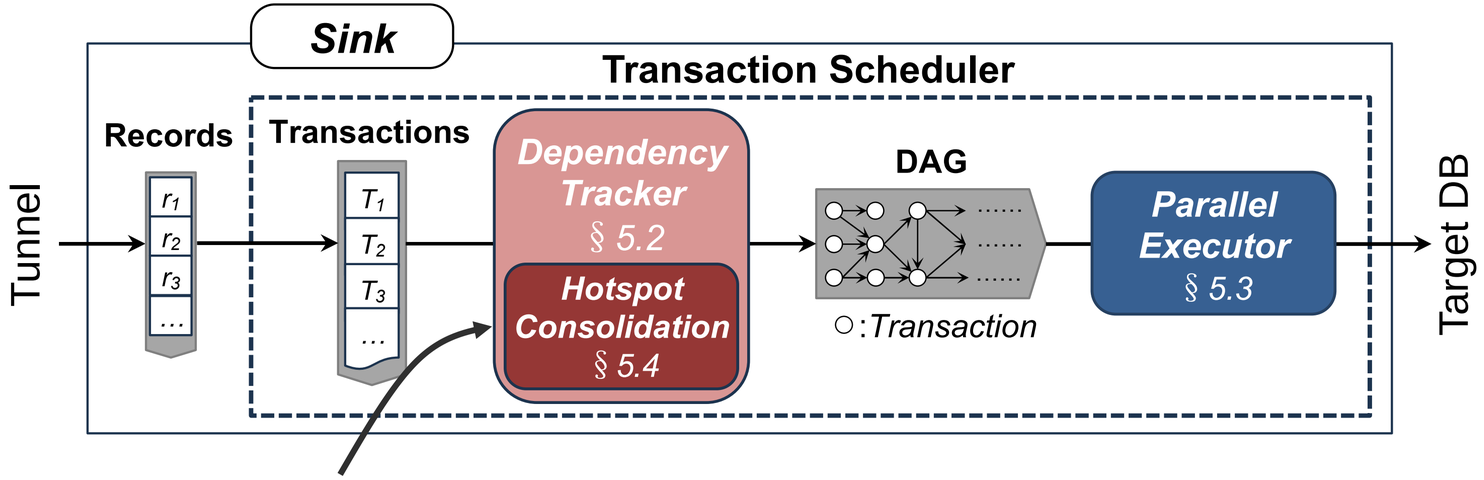

Dependency tracker: constructs a directed acyclic graph (DAG) as a dependency graph

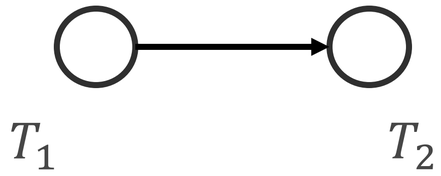

T_1 and T_2 both involve updates with same primary keys, and T_1 is earlier than T_2

Sink optimization

Consistency: DTS ensures eventual consistency for real time tasks, which allows non-conflicting transactions to be executed in parallel & out of order.

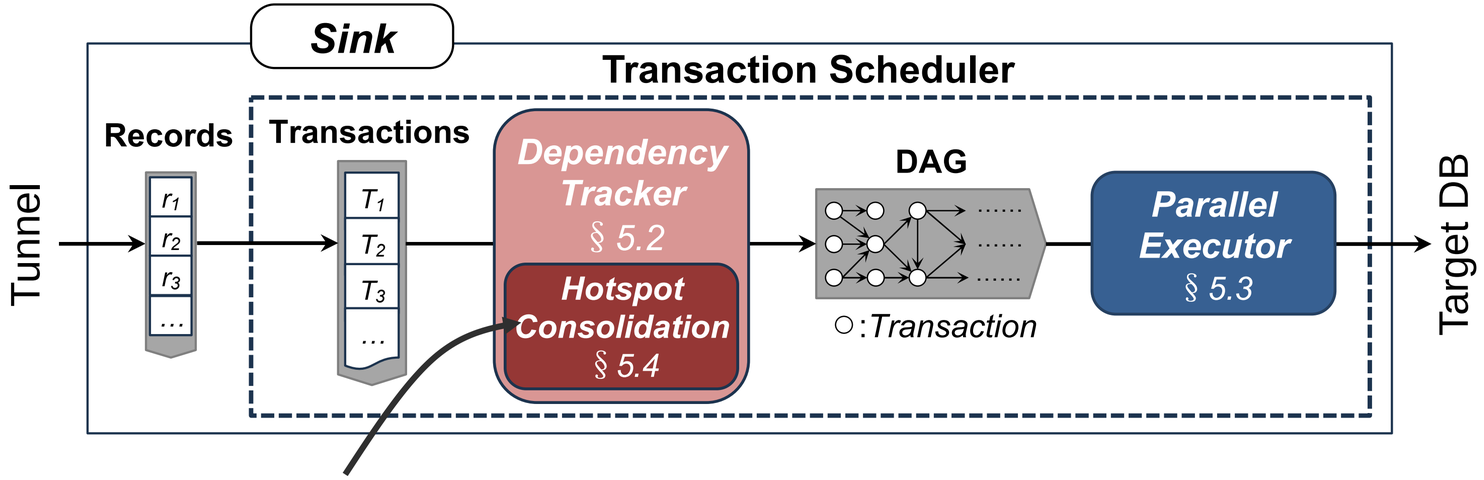

Hotspot consolidation: merges transactions with same frequently modified primary key

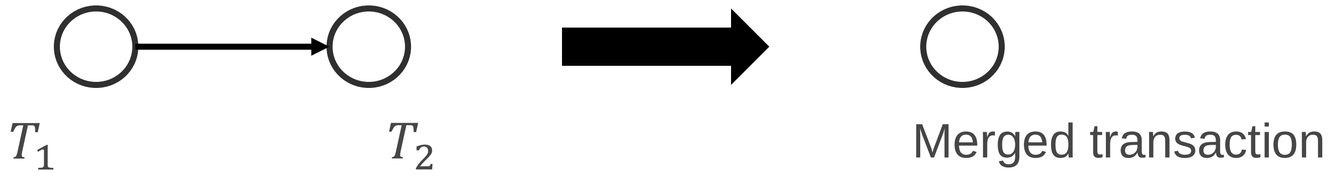

Alibaba Cloud

Sink optimization

Consistency: DTS ensures eventual consistency for real time tasks, which allows non-conflicting transactions to be executed in parallel & out of order.

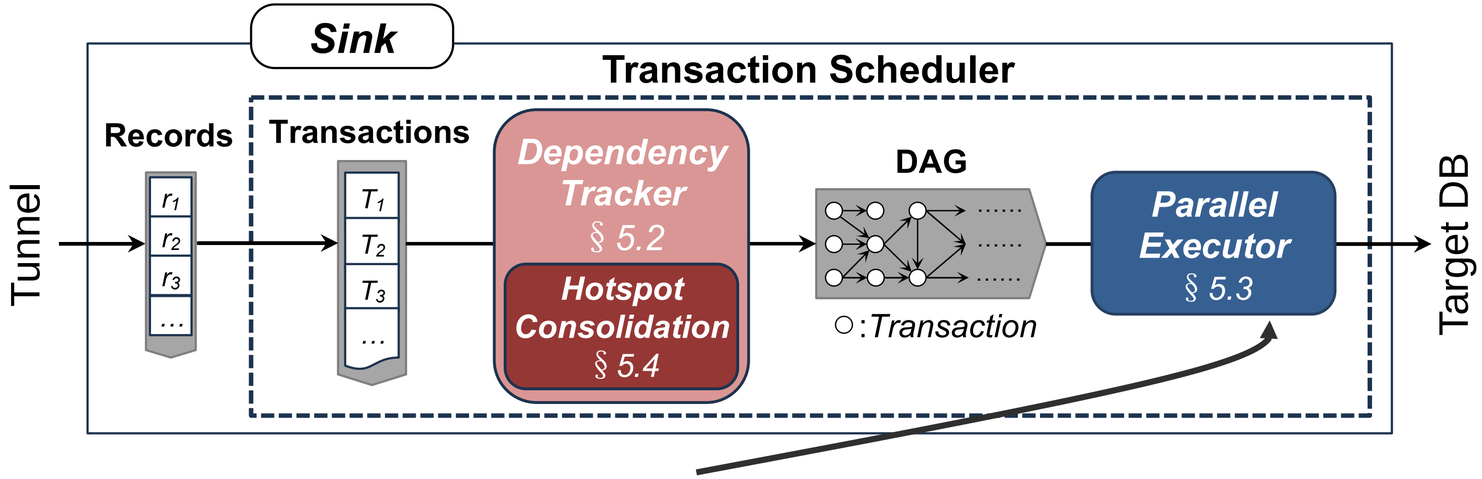

Parallel executor: extract then execute non-dependent transactions in parallel

Indegree = 0

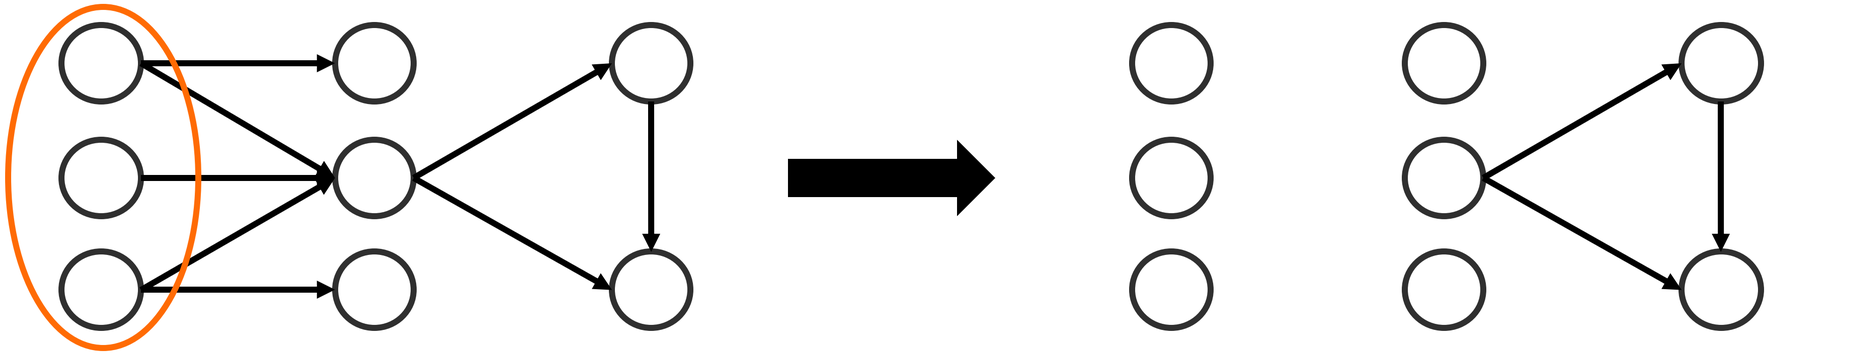

Parallel execution

Remaining graph

Sink optimization

Consistency: DTS ensures eventual consistency for real time tasks, which allows non-conflicting transactions to be executed in parallel & out of order.

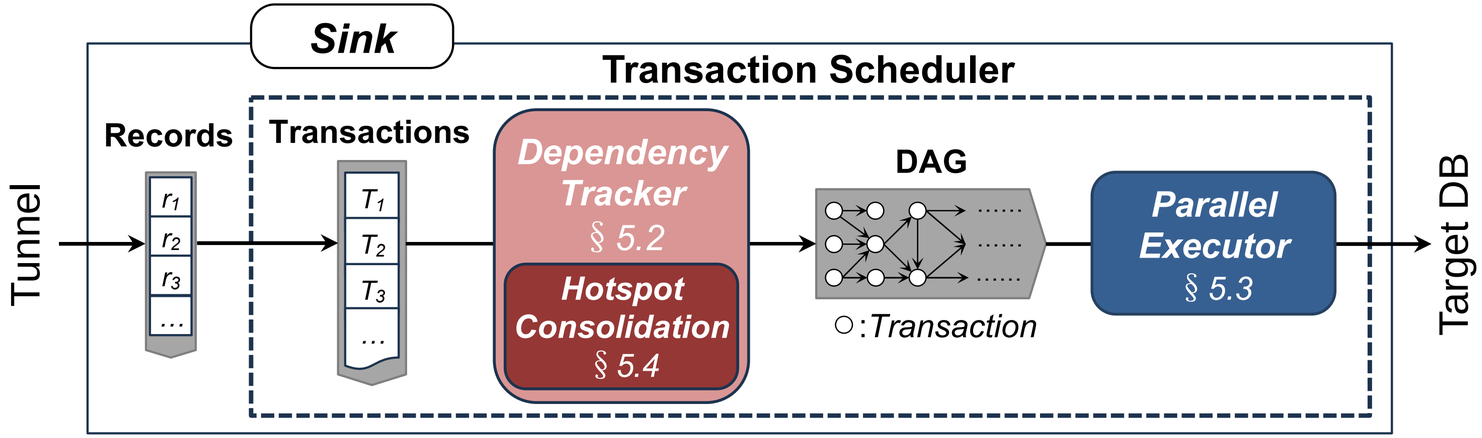

Trade-off: Consistency-performance trade off in sink optimization. We allow users to choose to disable sink optimizations.

Alibaba Cloud

Data transmission performance

DTS achieves modular (source and sink) and end-to-end high throughput and low latency

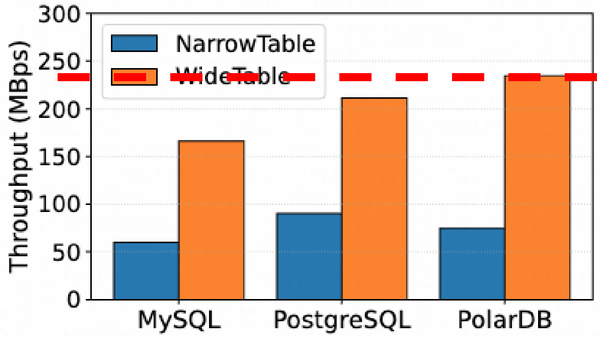

(a) end to end throughputs for Full tasks

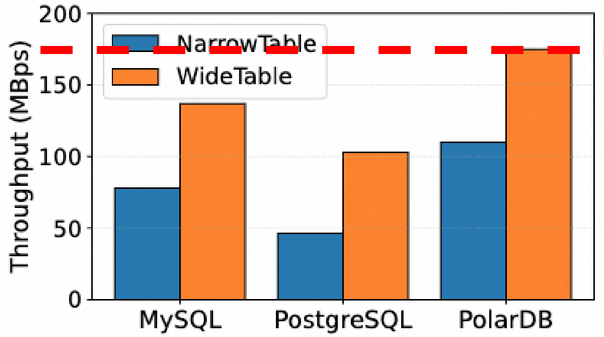

(b) Source throughputs for CDC tasks

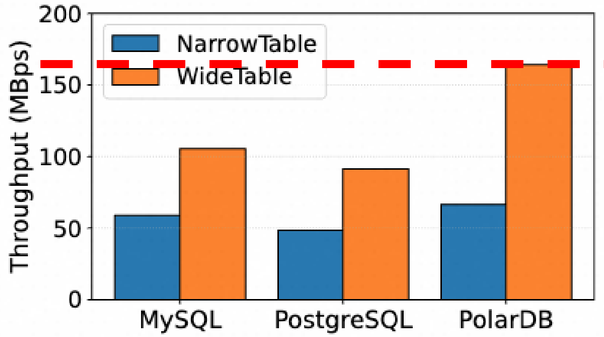

(c) Sink throughputs for CDC tasks

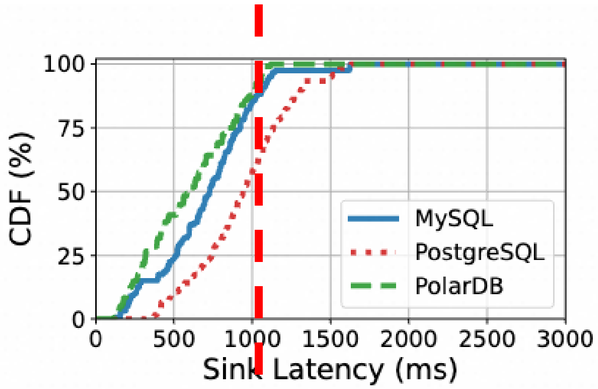

(d) latency distribution under 20MBps

Alibaba Cloud

Sink optimization performance

Parallel execution and hotspots consolidation improve transmission throughput to OLTP/OLAP target databases

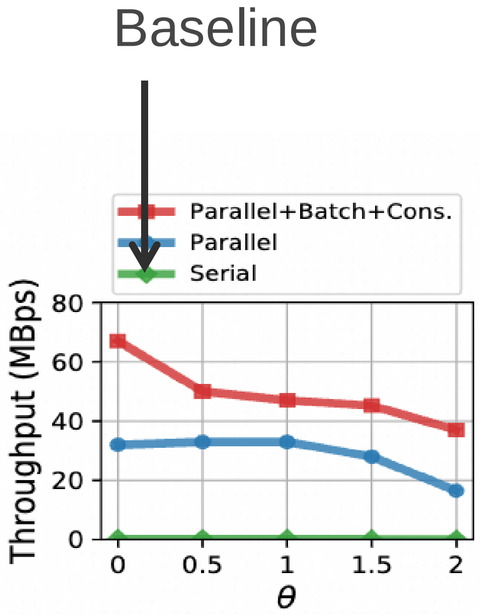

(a) Sink to MySQL varying skewness

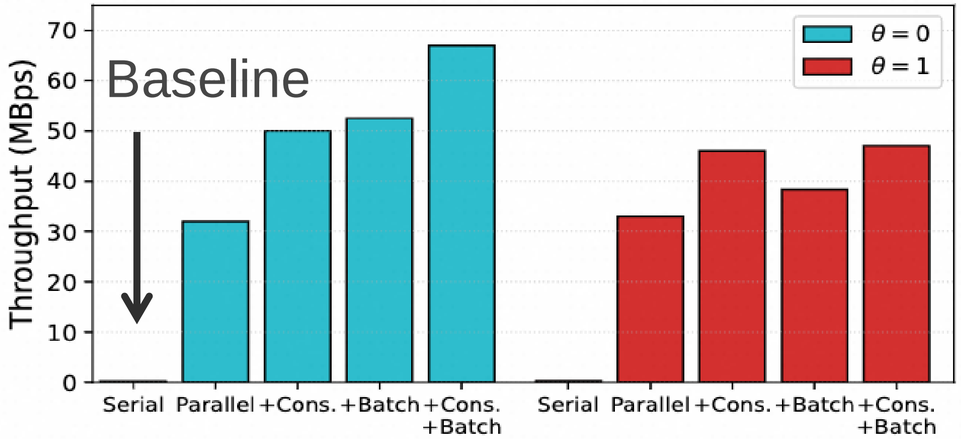

(b) Breakdown of Optimization: Sink to MySQL

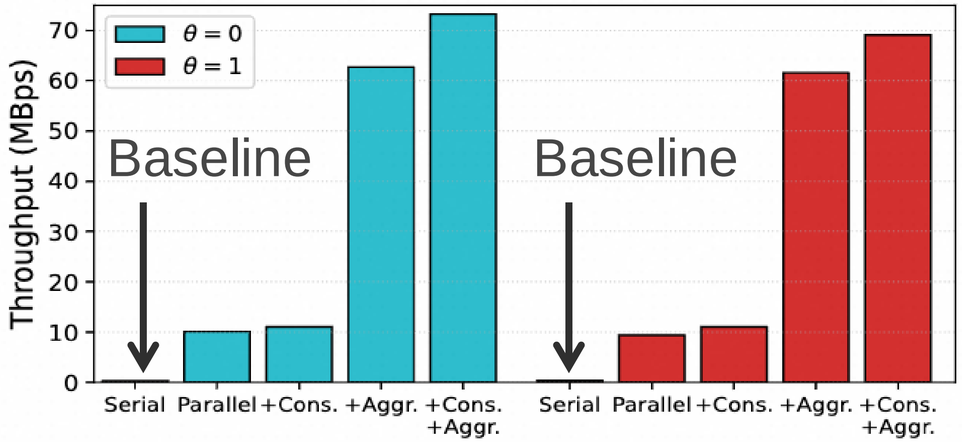

(c) Breakdown of Optimization: Sink to ADB

Real-world performance

DTS serverless achieves low user cost with low latency

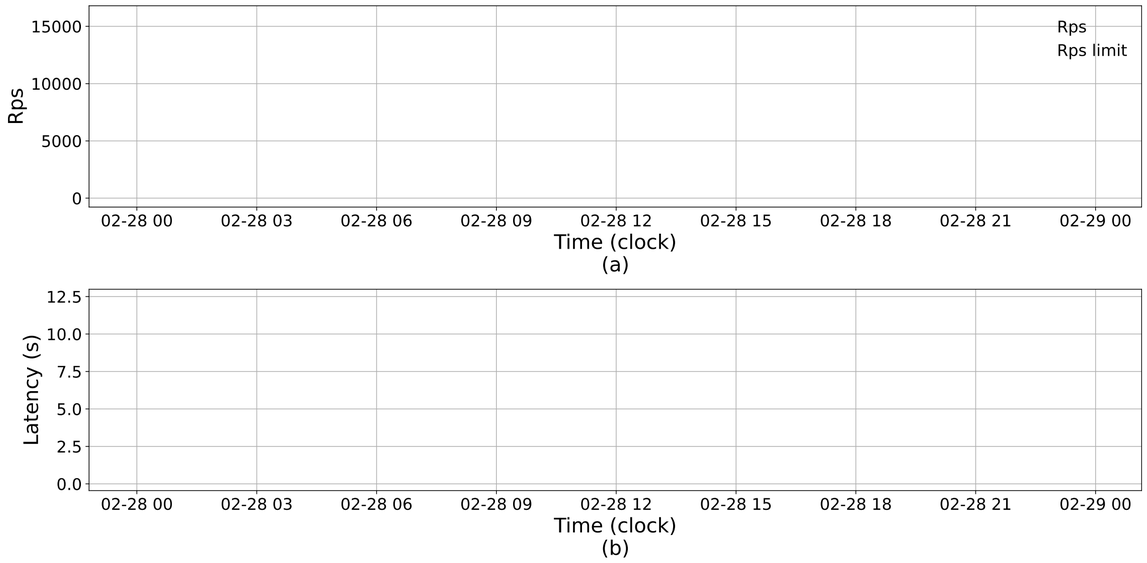

DTS achieves high throughput with low latency

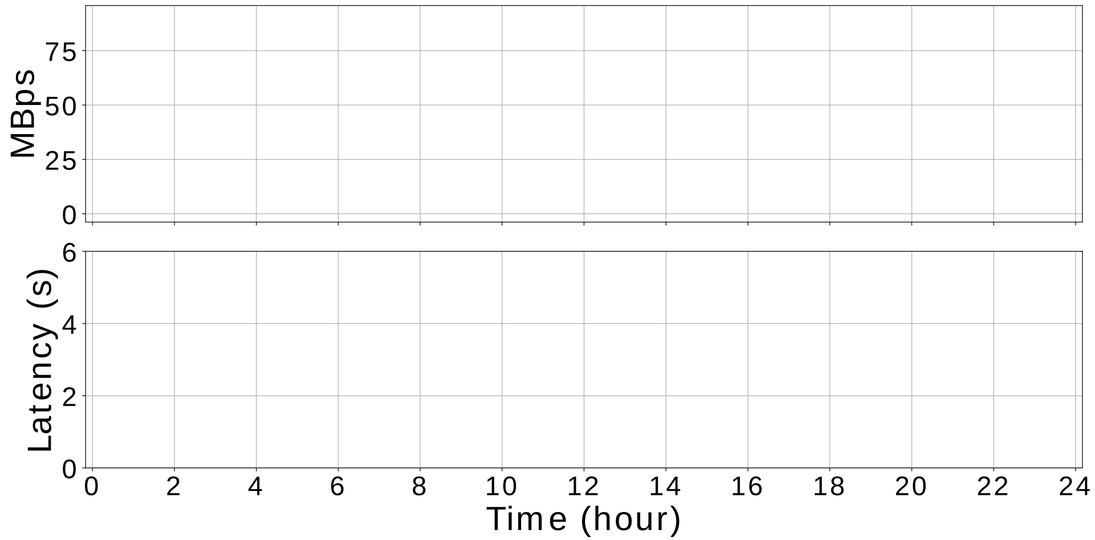

Alibaba Cloud

Summary

Challenges	Solutions	Evaluations
C1: High database type diversity	Any-to-any (A2A) model	Low development effort
C2: Large transmission task number	DTS serverless Bandwidth scheduler	Low user cost & low latency
C3: High transmission velocity	Sink optimizations	High throughput & low latency