

Presto@Alibaba Group

AnalyticDB - High-Performance Real-Time OLAP Cloud Database Service

Jessica Zhang
Presto Meetup
10/17/2019

About Us – Alibaba Group



Digital Media & Entertainment



Core E-Commerce



Local Services



870 Million

**Active Users of Alipay
Around the Globe**

Serving 15 million SMEs worldwide

100 Million

**Packages Processed by Cai Niao
Logistic Network Each Day**

Covering 15 hundred cities and
counties in mainland China

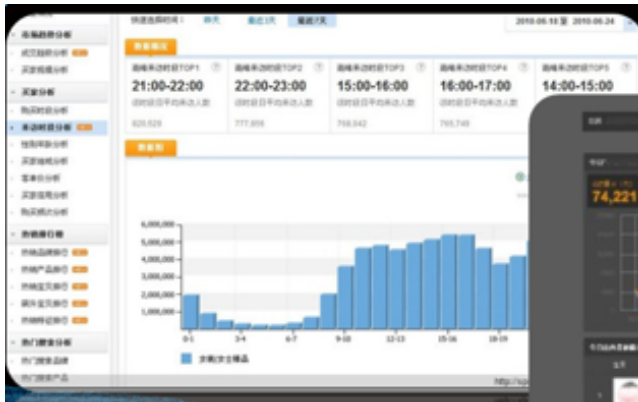
\$768 Billion

GMV of FY2018

552 million monthly active users

online data analytics to explore business value of data for Alibaba cloud customers and Alibaba Group

Report



Delivers high performance through multi-dimensional query to support BI analysis

Interactive



Real-time interactive query from millions labeled customers to find high value targets (E-commerce, CRM...)

Intelligence



Internet company to build data analytics platform to manage & explore data value (Marketing, AD, O2O...)

Our Customers

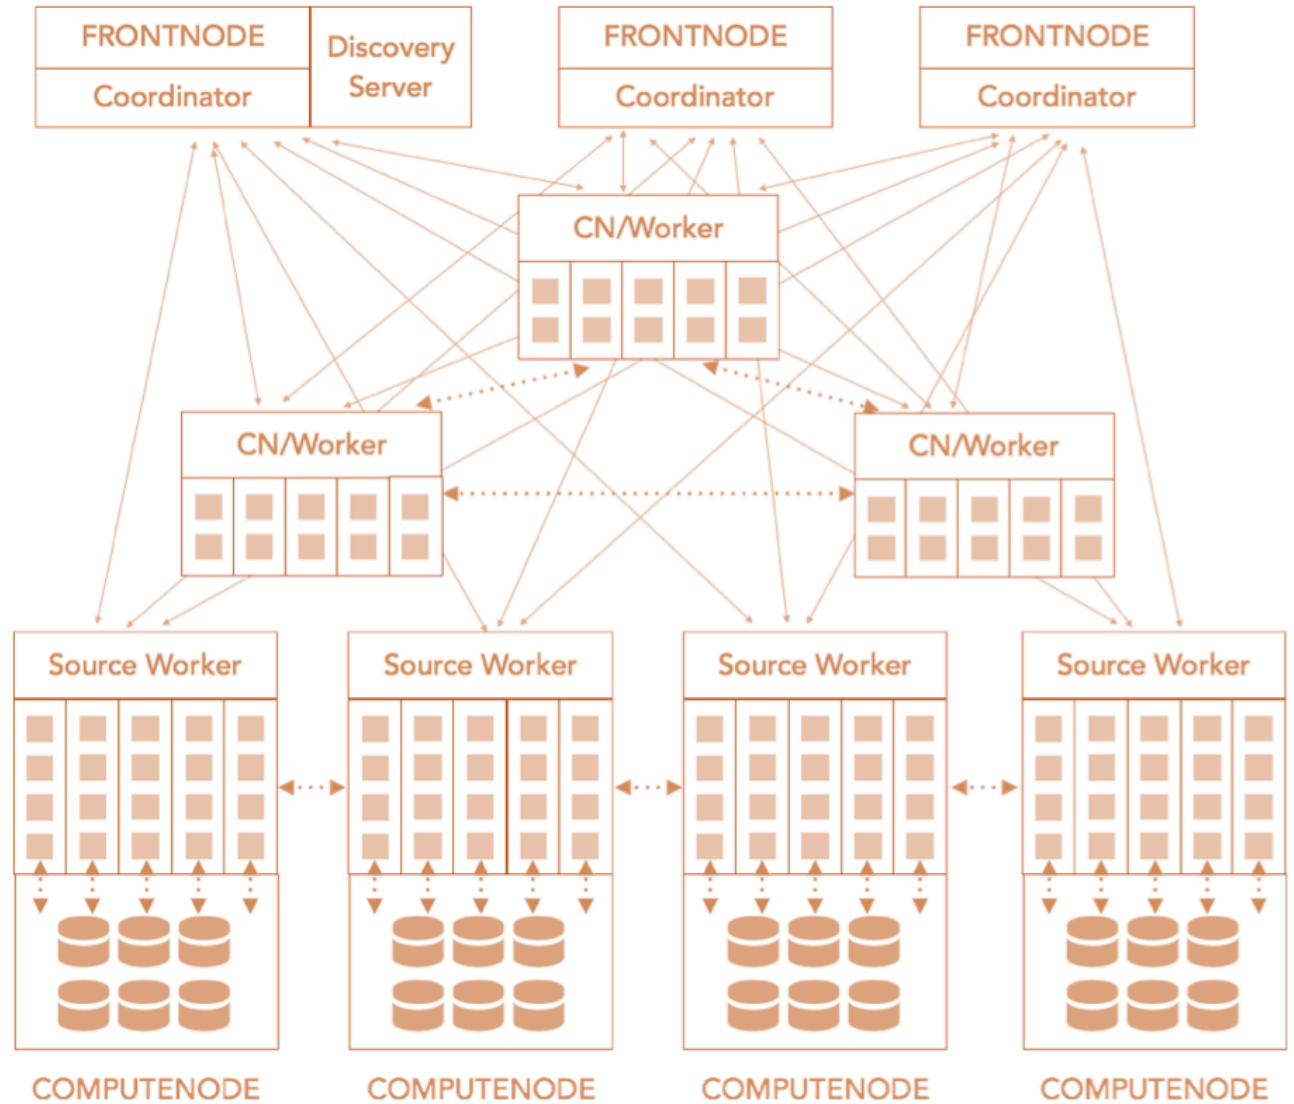


- Public Cloud & Private Cloud & Alibaba Group
 - Across 10+ industries

- Single DB
 - thousands nodes
 - thousands QPS
 - hundreds Concurrent Queries

AnalyticDB Architecture

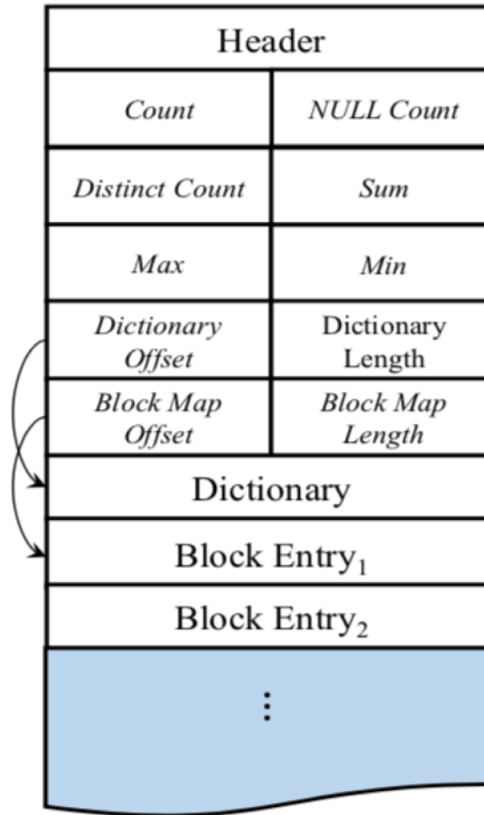
- Distributed processing system
- High availability
- Capable of scaling to thousands of nodes
- Extremely low latency, ~ms level



AnalyticDB Storage Engine

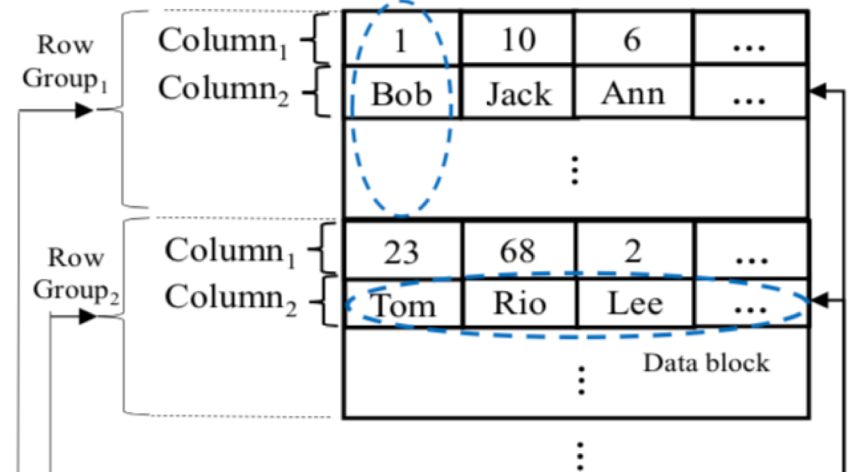
- File format
 - Hybrid row-column store
 - Fixed-size storage model for complex data type (JSON, Full Text, Vector data)
- Index engine
 - All-column index
 - Space-efficient async index building
 - Support Insert / Delete / Update baseline data + incremental data
- Distributed storage engine
 - Hash partition
 - High availability

Detail Meta File for One Column




Data of One Row

Detail File



Index File for One Column



Presto in AnalyticDB

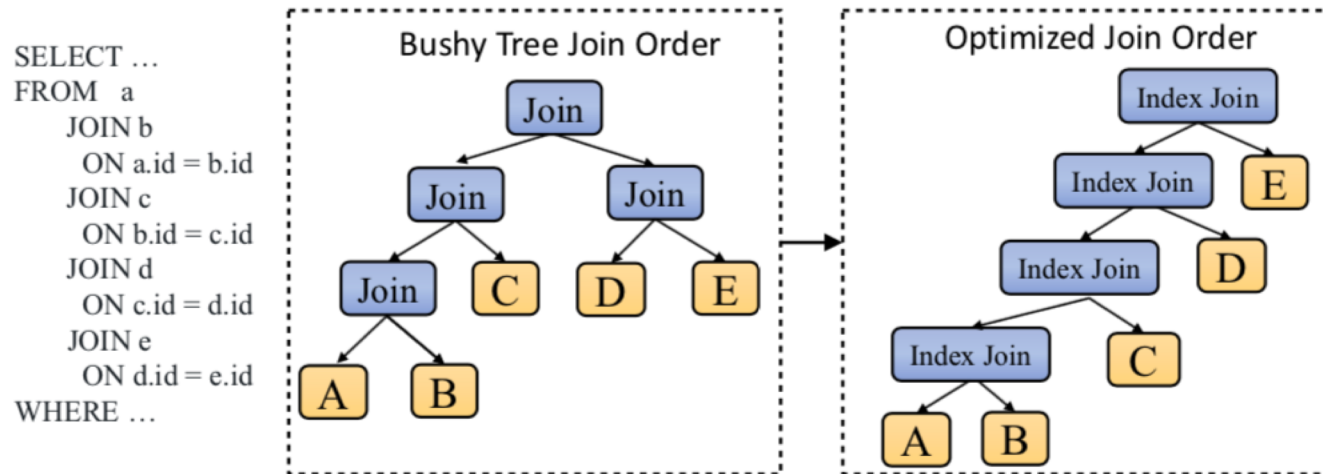
- Connector to storage engine
- execution engine improvement

Connector – Storage-aware optimization

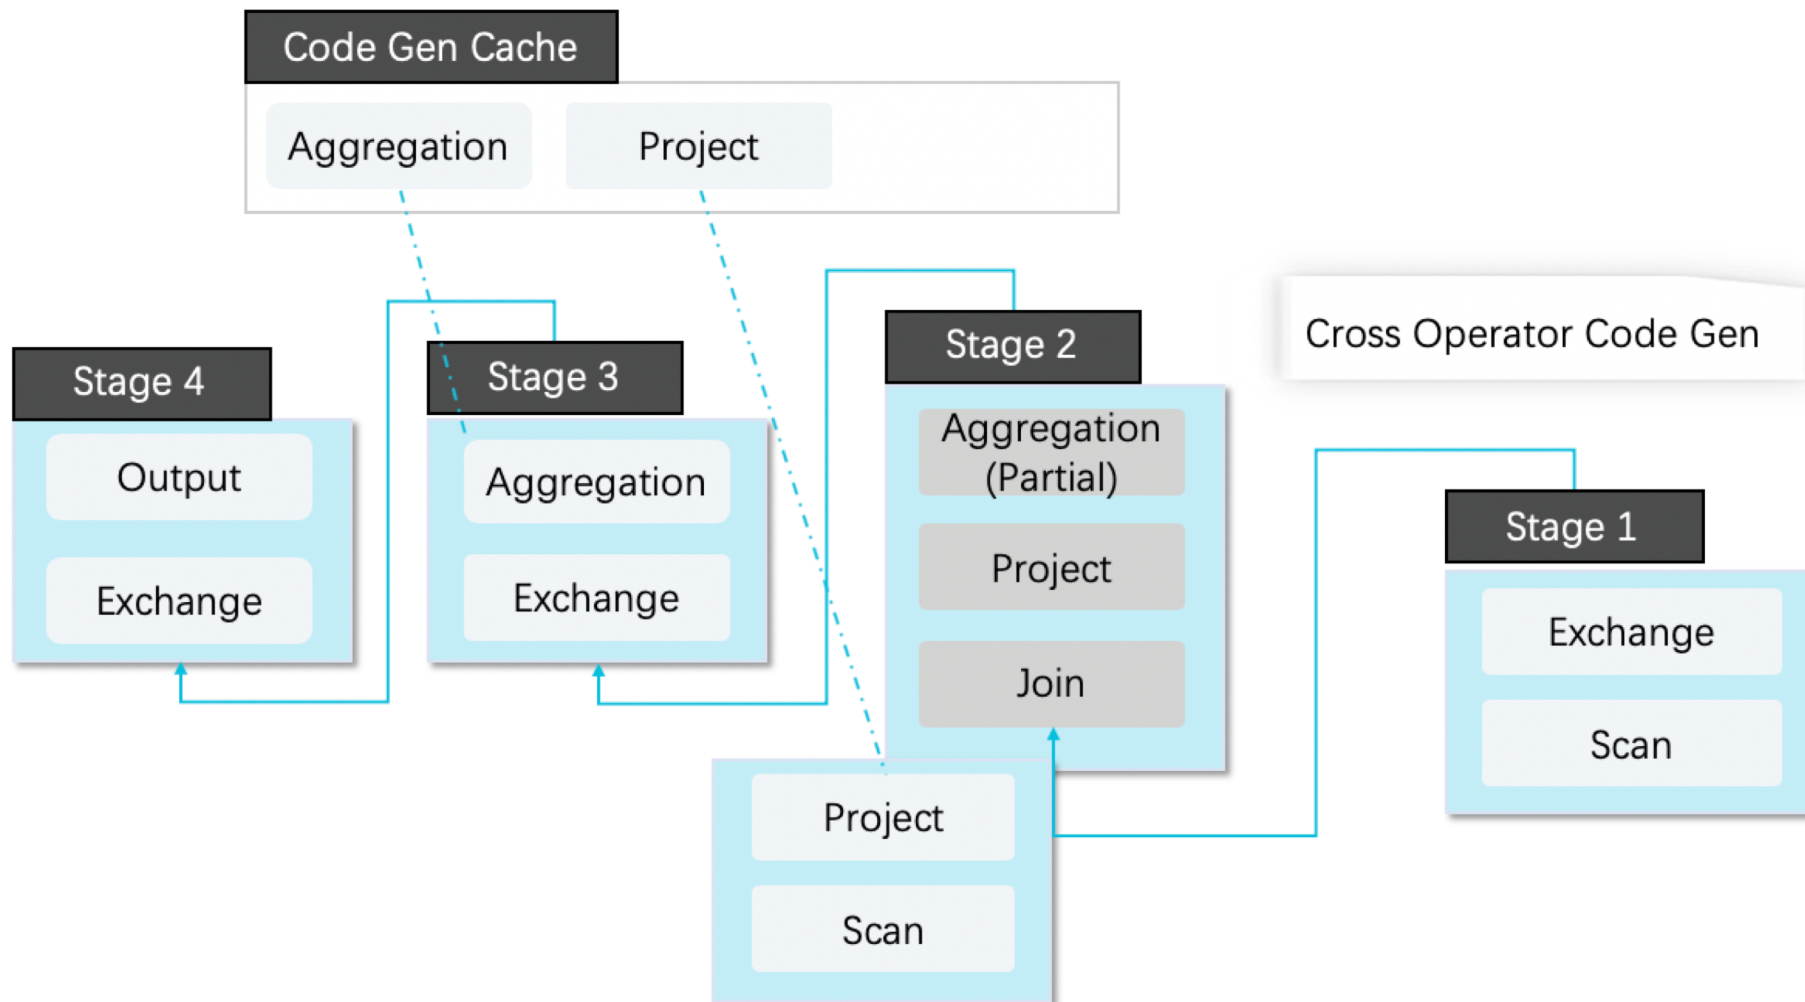
- Predicate pushdown
 - Filter condition on columns (comparison, LIKE, IN, etc.)
 - UDF (time function, vector function, etc.)
- Projection pushdown
- Colocated join / aggregation
 - When join keys or aggregation keys match hash partitioning keys
 - Reduce shuffle cost
- Bucket Pruning
 - When Equal or IN predicate on hash partitioning column
 - Avoid scheduling tasks on unnecessary nodes

Connector – Nested loop index join

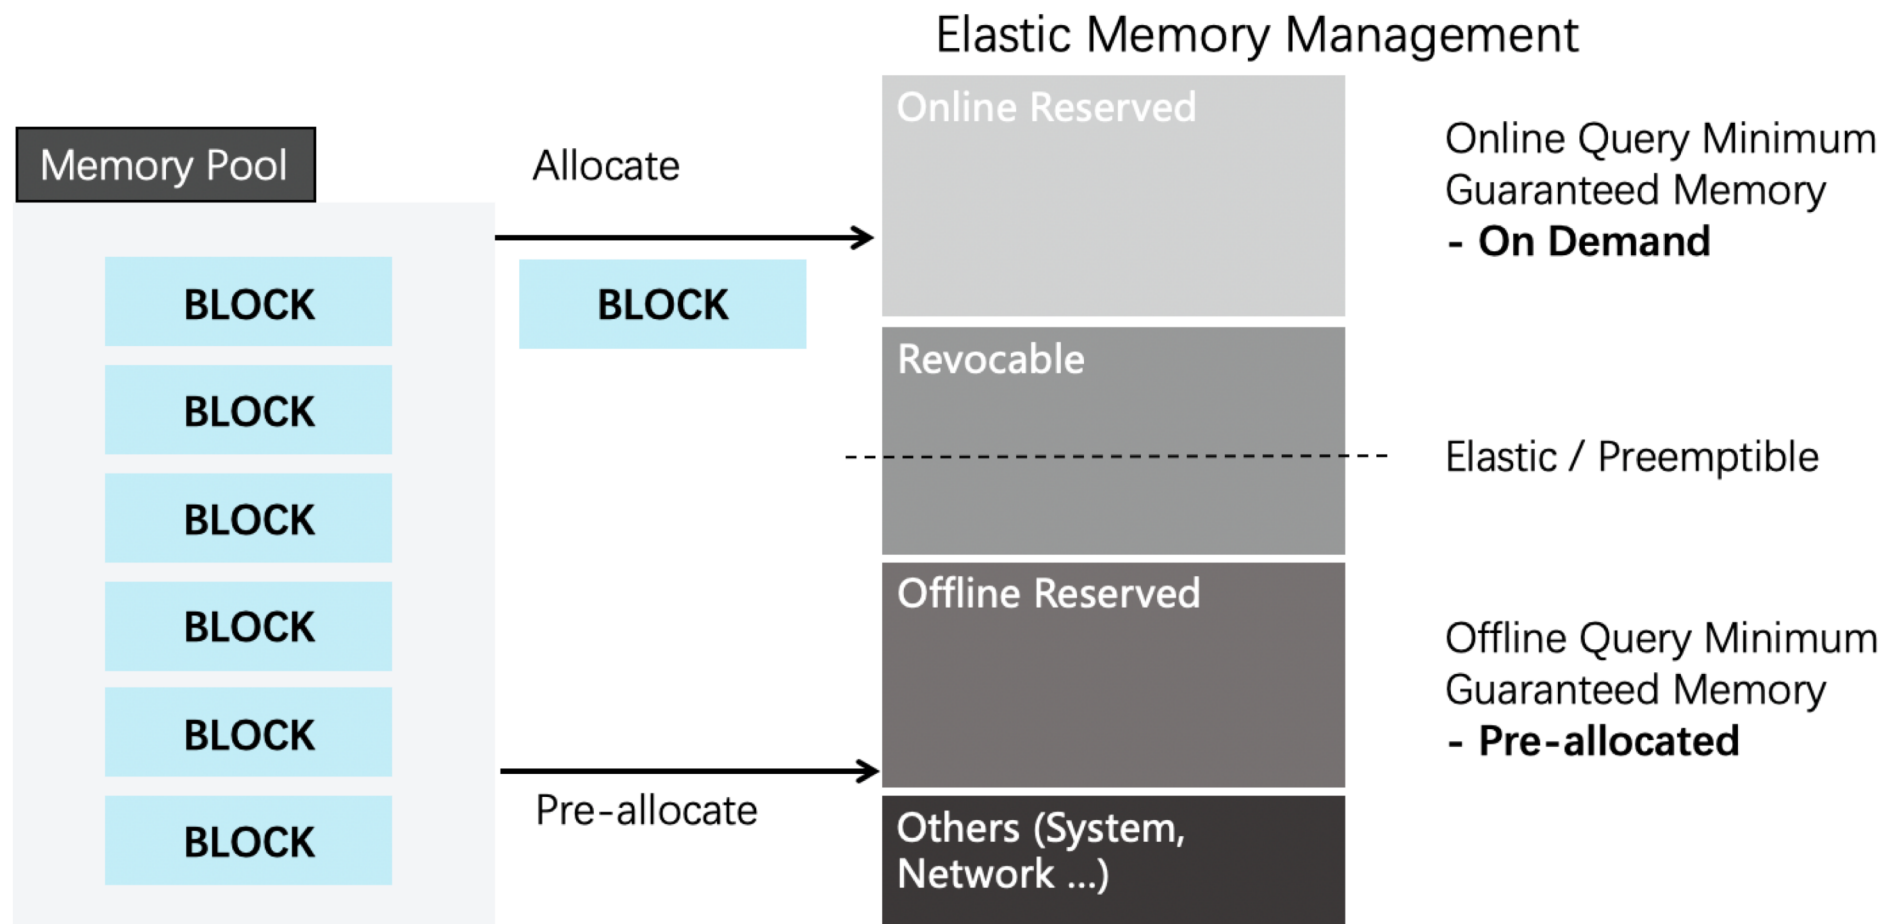
- Fit high-filtering lookup join case
- Preserve left side data distribution



Compute Engine – Code Gen



Compute Engine - Memory



- Pros

- Eliminate Serde overhead
- Reduce data shuffle cost
- Reduce memory fragmentation
- Cooperate with memory pool

- Cons

- Degrade read performance

- Analyze
 - Table cardinality
 - Column cardinality
 - Column histogram
- Real-time sampling statistics
 - Efficient light weight
 - Utilize storage engine index

Optimizer -Tools

- Advanced plan explain
- Plan hint
- Planning time analysis
- Query pattern analysis

- Cost model
 - Filter factor
 - Derived column cardinality
 - Join cardinality
- Cost-based transformation
 - Join reorder
 - Join distribution type
 - Property enforcement
 - Partial aggregation pushdown
 - Filter join pushdown across aggregation

Thanks



Alibaba Group