DTCC 2015

Database Technology
Conference China 2015

Postgres-XL/XC
Scale-out Approach in PostgreSQL

May 18th, 2015

NTT DATA INTELLILINK Corporation
Koichi Suzuki

NTT DATA

Introduction

About the Speaker

- Fellow at NTT DATA Intellilink Corporation
- Principal Technology Professionals at NTT DATA Group

In Charge Of

- General Database Technology
- Database in huge data warehouse and its design
- PostgreSQL and its cluster technology

In The Past

- Character Set Standard (Extended Unix Code, Unicode, etc)
- Heisei-font development (Technical Committee)
- Oracle Porting
- Object-Relational Database

Motivation

- Growing Database Workload both for OLTP (OnLine Transaction Processing) and OLAP (OnLine Analytical Processing) applications.
- Shared-Nothing Approach
  - Performance with commodity hardware/software
- Extension to existing PostgreSQL
- Transparent API
  - Internal API could be different
  - Transparent libpq Interface
  - No significant restriction on transaction ACID properties and SQL language.

Scale-out approach

- Distribution/Replication of table rows among different database “nodes”
  - Parallelism
  - Local join operation
- SQL planning for row distribution/replication
- Consistent and synchronous transaction management among “nodes”
- Performance with commodity hardware/software

Read Scale-out in PostgreSQL Master/Slave

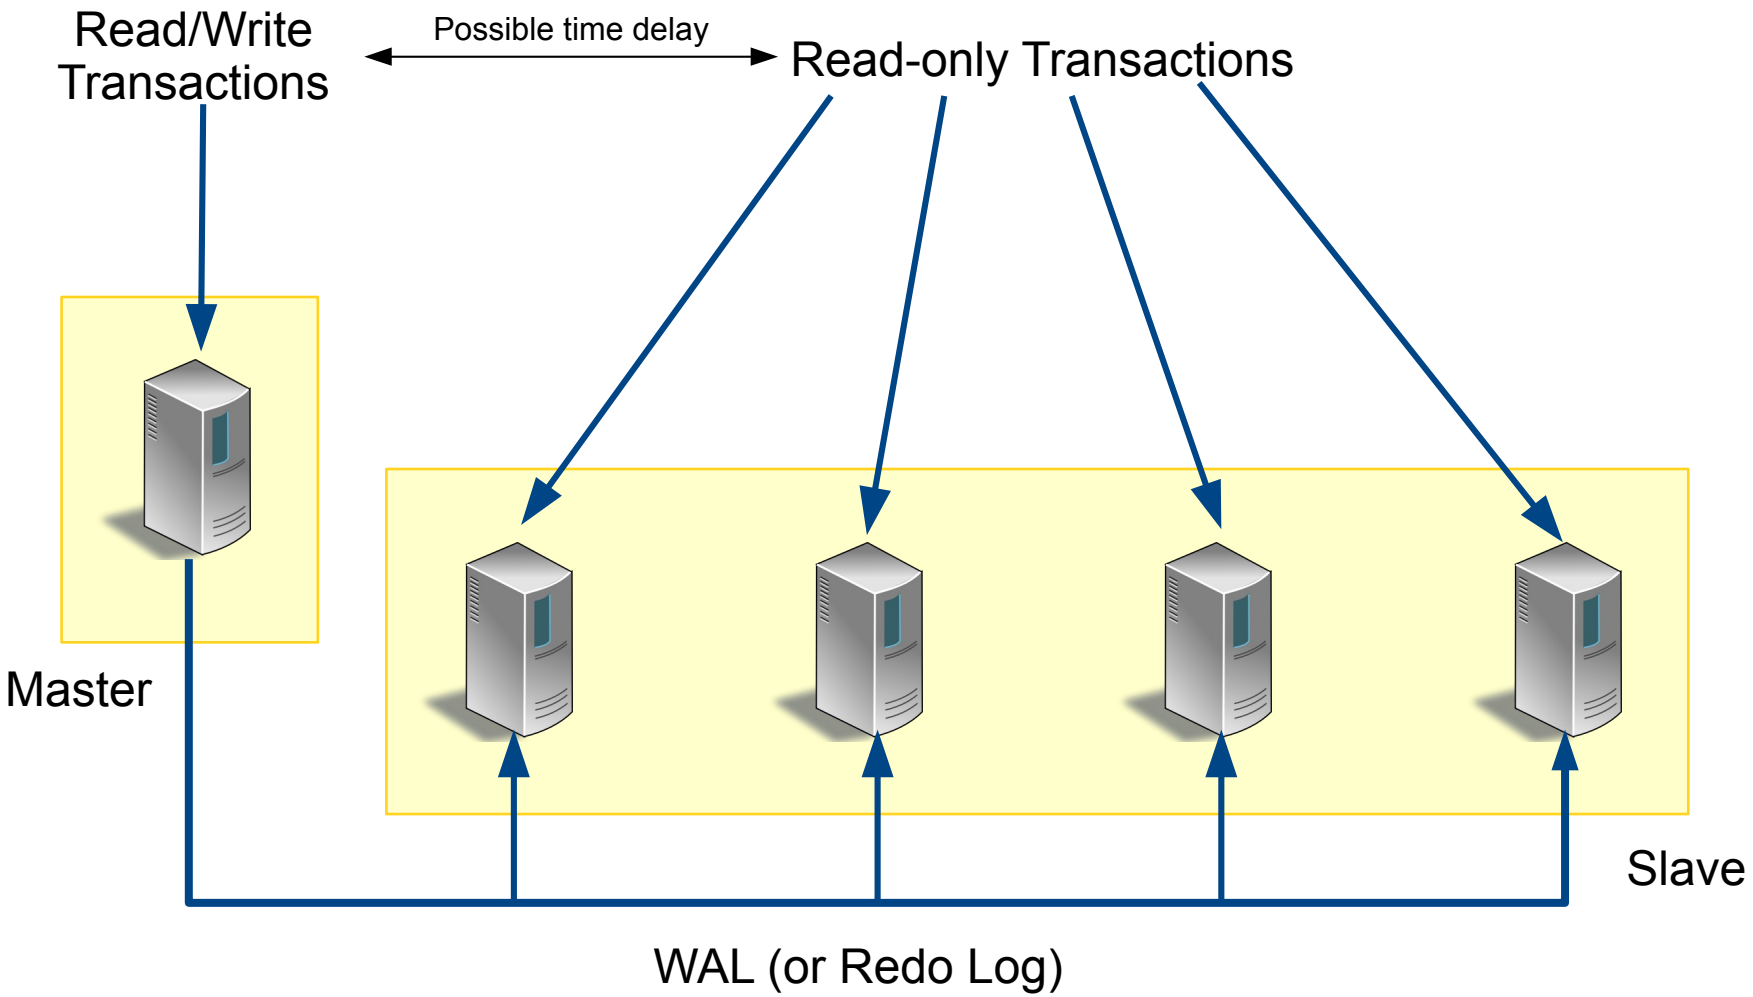

Scaling Out in Postgres XL/XC

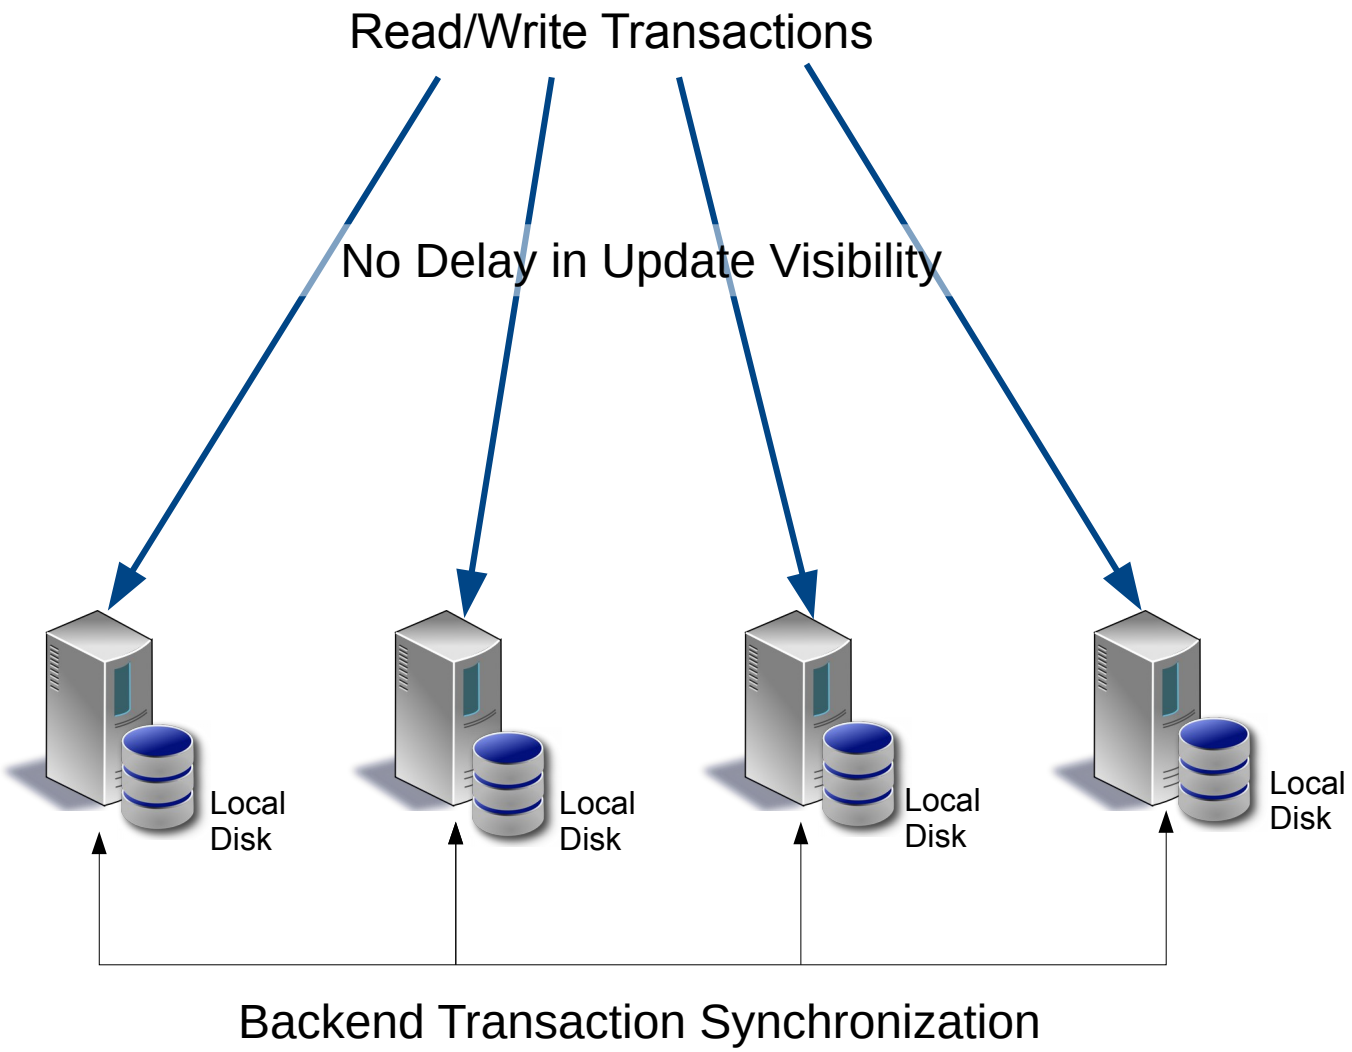

OLTP Workload Scalability and Table Design

DBT-1 Workload Scalability

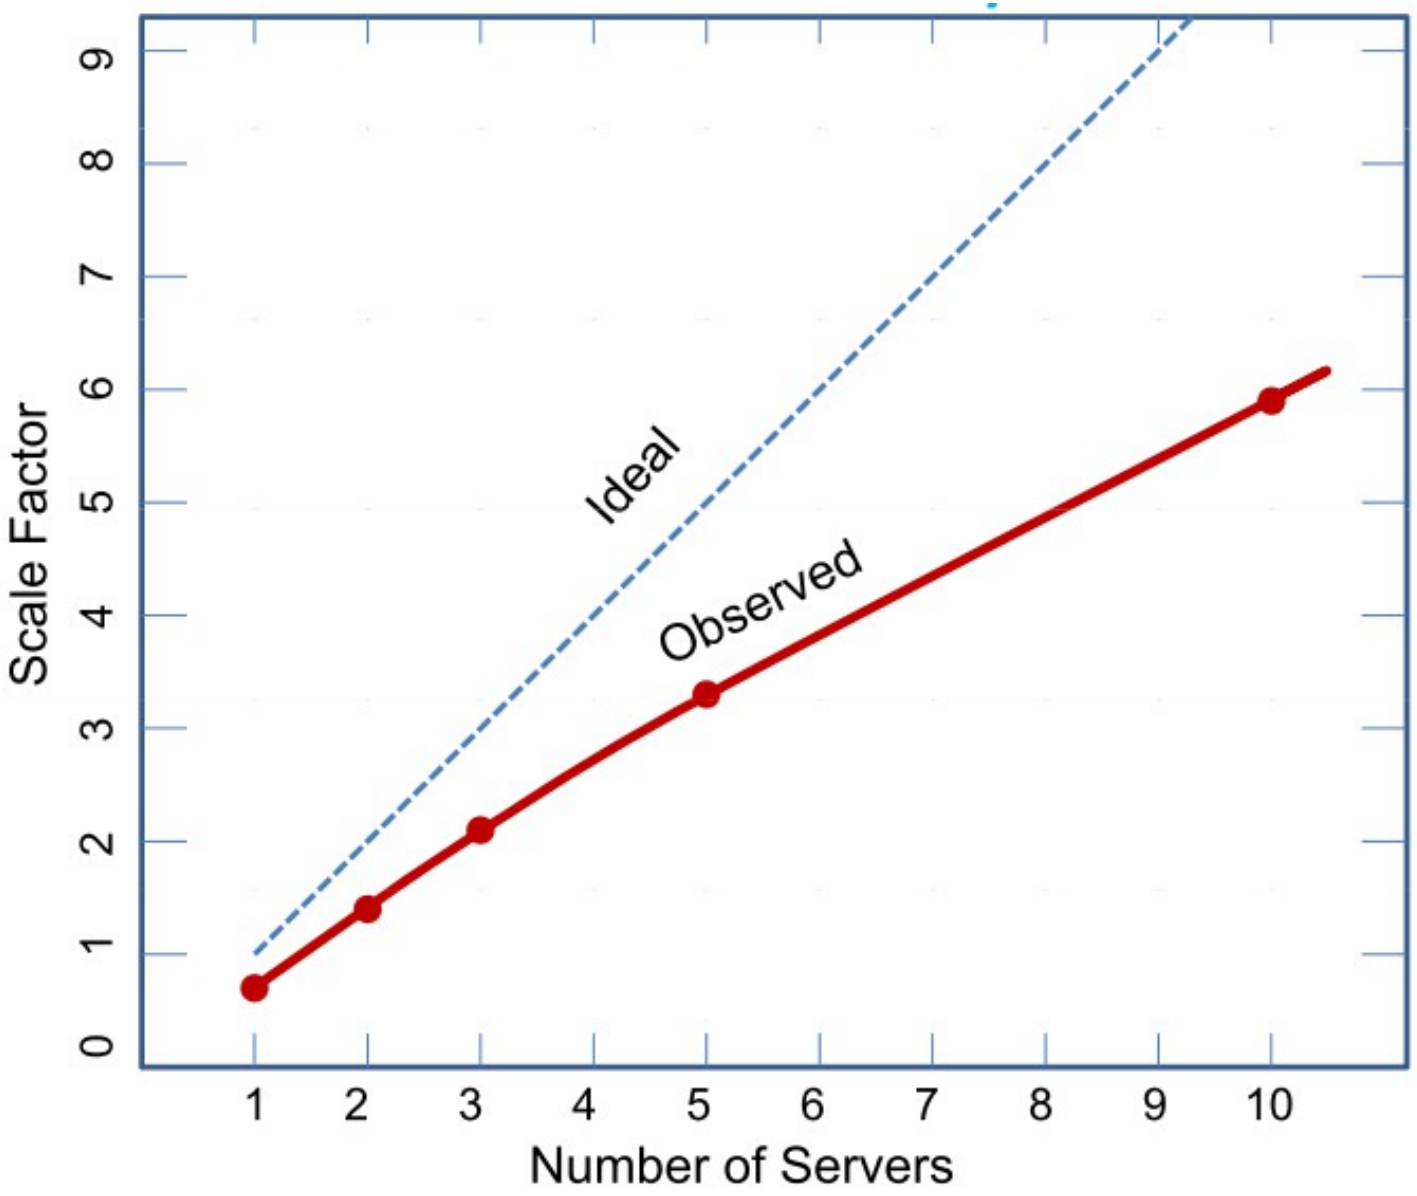

DBT-1 (Rev)

Table Design in DBT-1 Benchmark

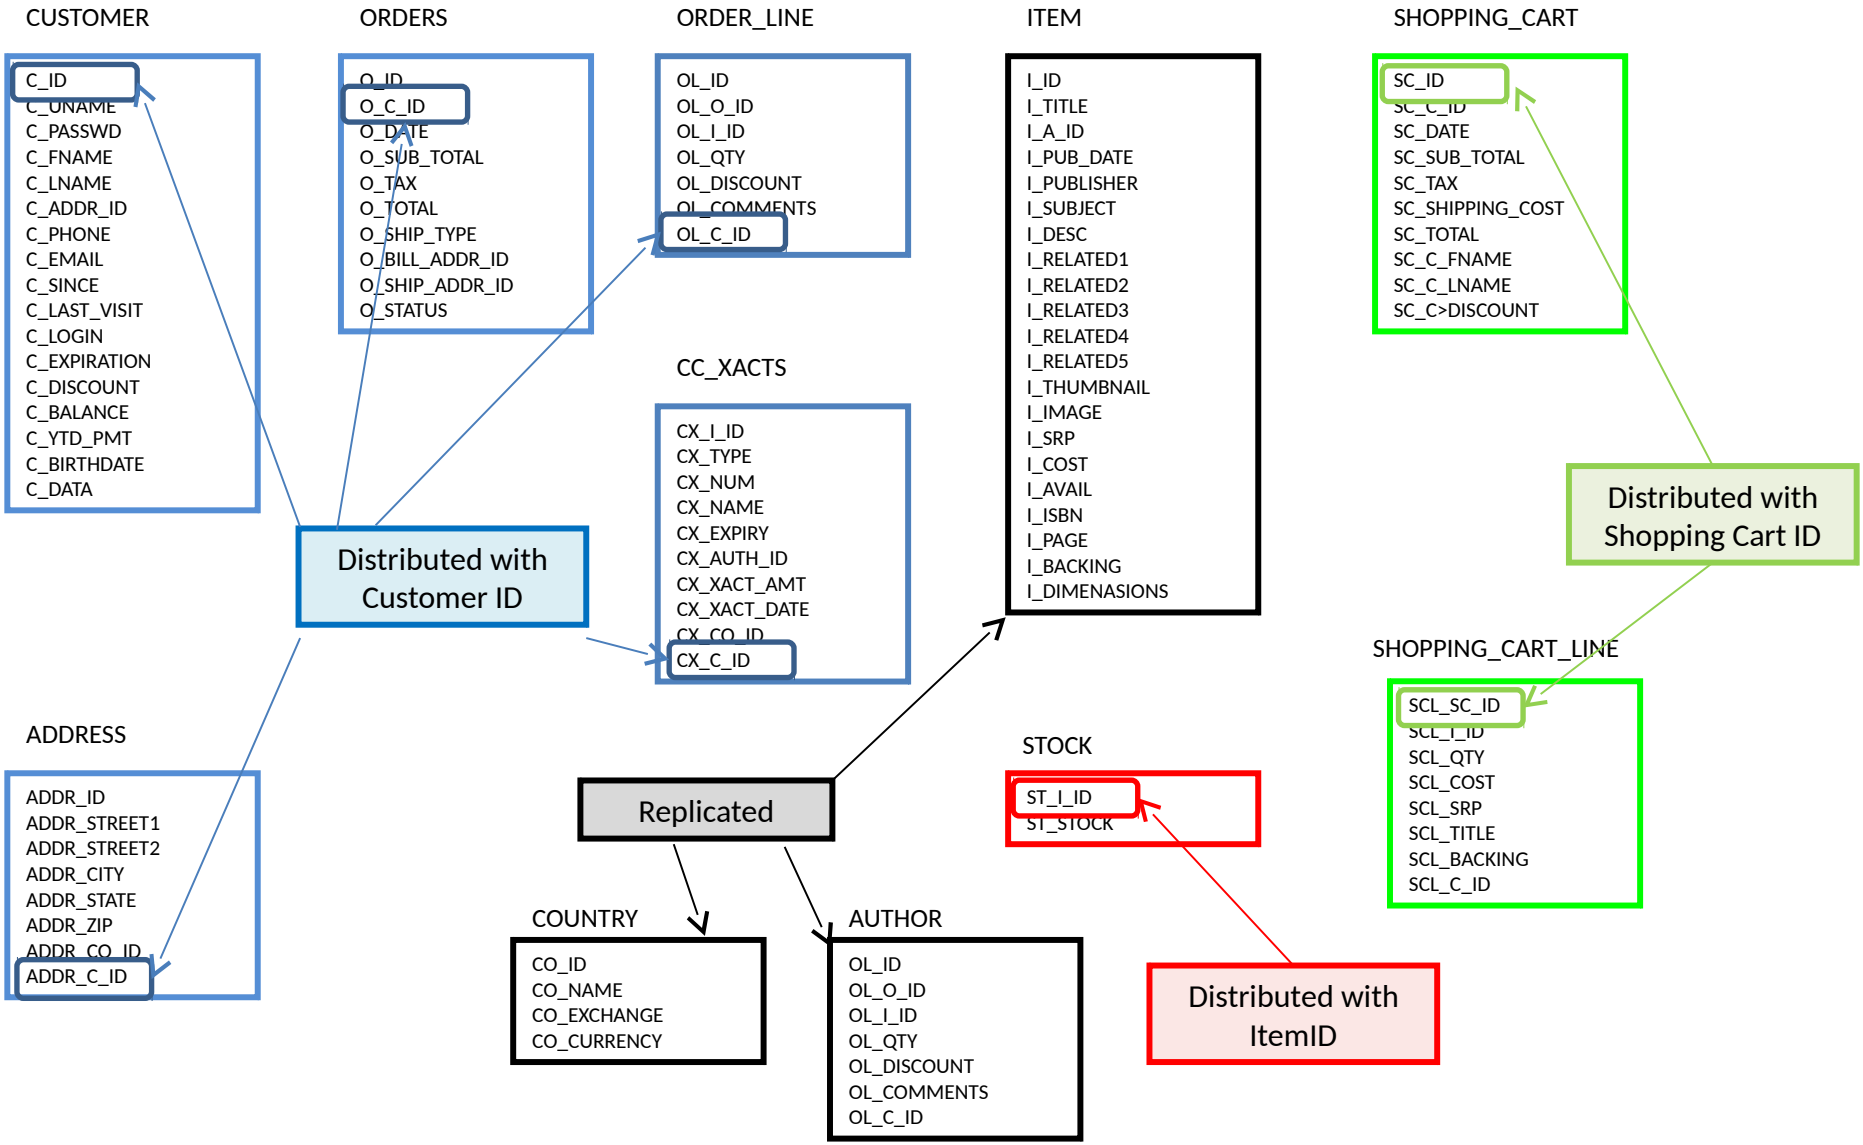

Scale Out Approach (1): Table Distribution/Replication

NTT DATA

Categorize tables into two groups:

Large and frequently-updated tables

→ Distribute rows among nodes (Distributed Tables)
→ Based on a column value (distribution key)
→ Hash, modulo or round-robin

→ Parallelism among transactions (OLTP) or in SQL processing (OLAP)

Smaller and stable tables

→ Replicate among nodes (Replicated Tables)

→ Join Pushdown

Avoid joins between Distributed Tables with join keys different from distribution key as possible.

Scale Out Table Design in DBT-1

Three distribution keys:

- Customer ID
- Shopping Cart ID
- Item ID

Some transactions involve joins across distributed tables with non-distribution join keys.

Some More in XL/XC Node Configuration

Node Configuration: Two-Tier Approach

Coordinator:

- Holds catalog information
- Build global SQL plan
- Tell Datanode what to do in SQL statements

Datanode

- Holds actual data
- Run local SQL statement from Coordiator
  (In XL, datanode may ask other datanodes for their local data)

Coordinator and Datanode

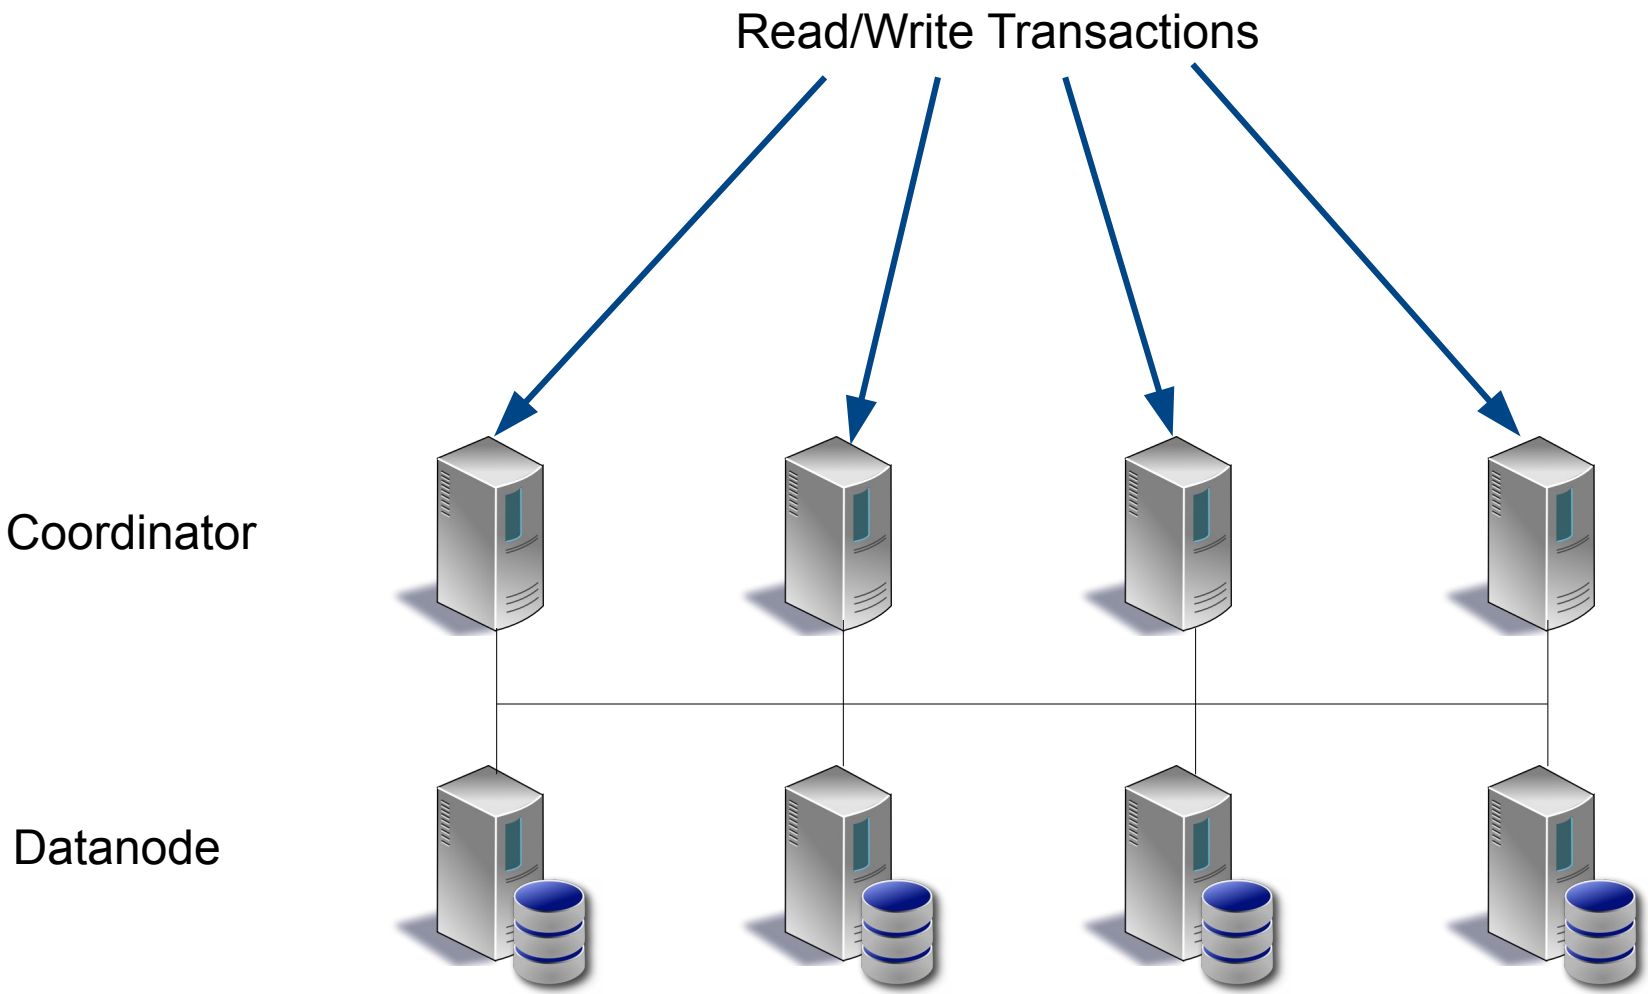

Node Configuration: Yet Another Node: GTM

GTM: Global Transaction Manager

Synchronizes each node's transaction status

Why GTM? Two-Phase Commit Protocol doesn't work?

NTT DATA

Two-Phase Commit Protocol Does:

- Maintain database consistency in transactions updating more than one node.

Two-Phase Commit Protocol Doesn't:

- Maintain Atomic Visibility of Updates to other transactions (next slide)

Atomic Visibility and GTM

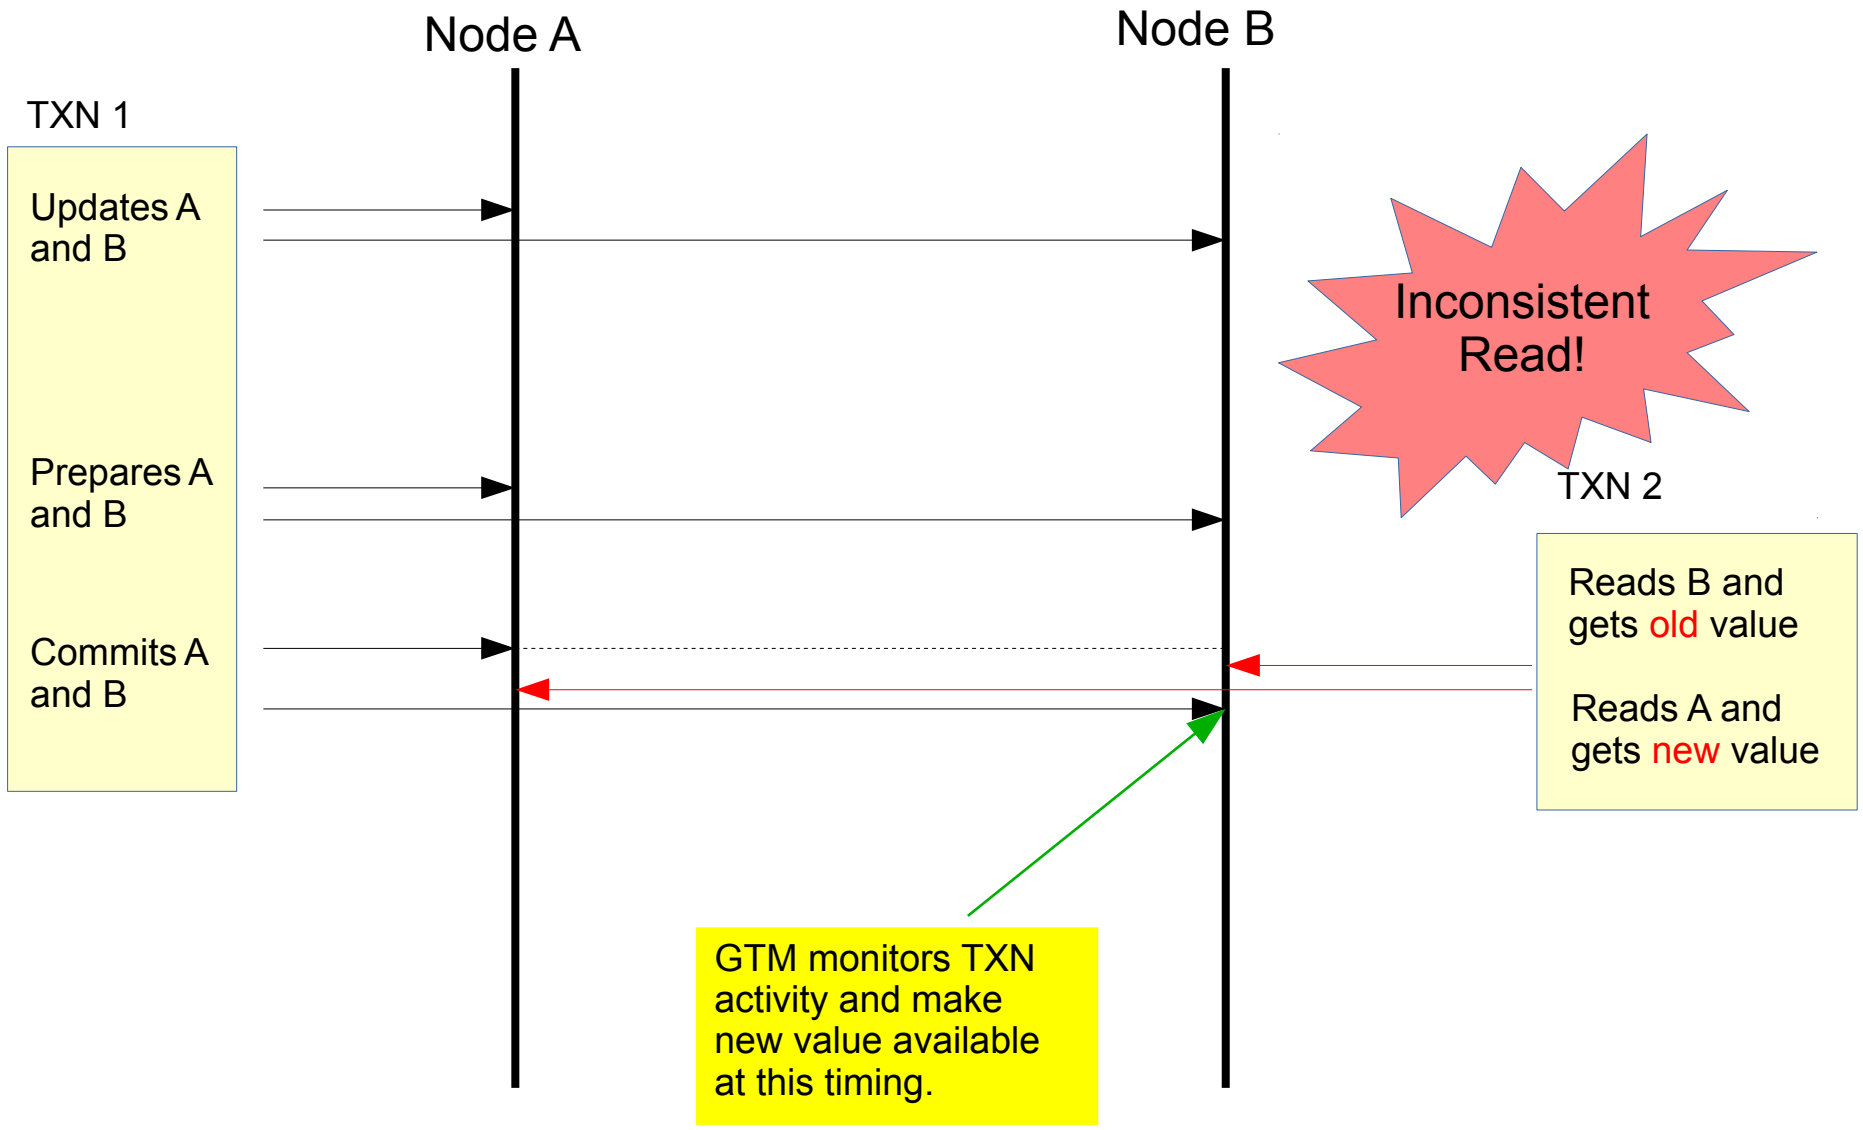

Final Configuration: GTM, Coordinator and Datanode

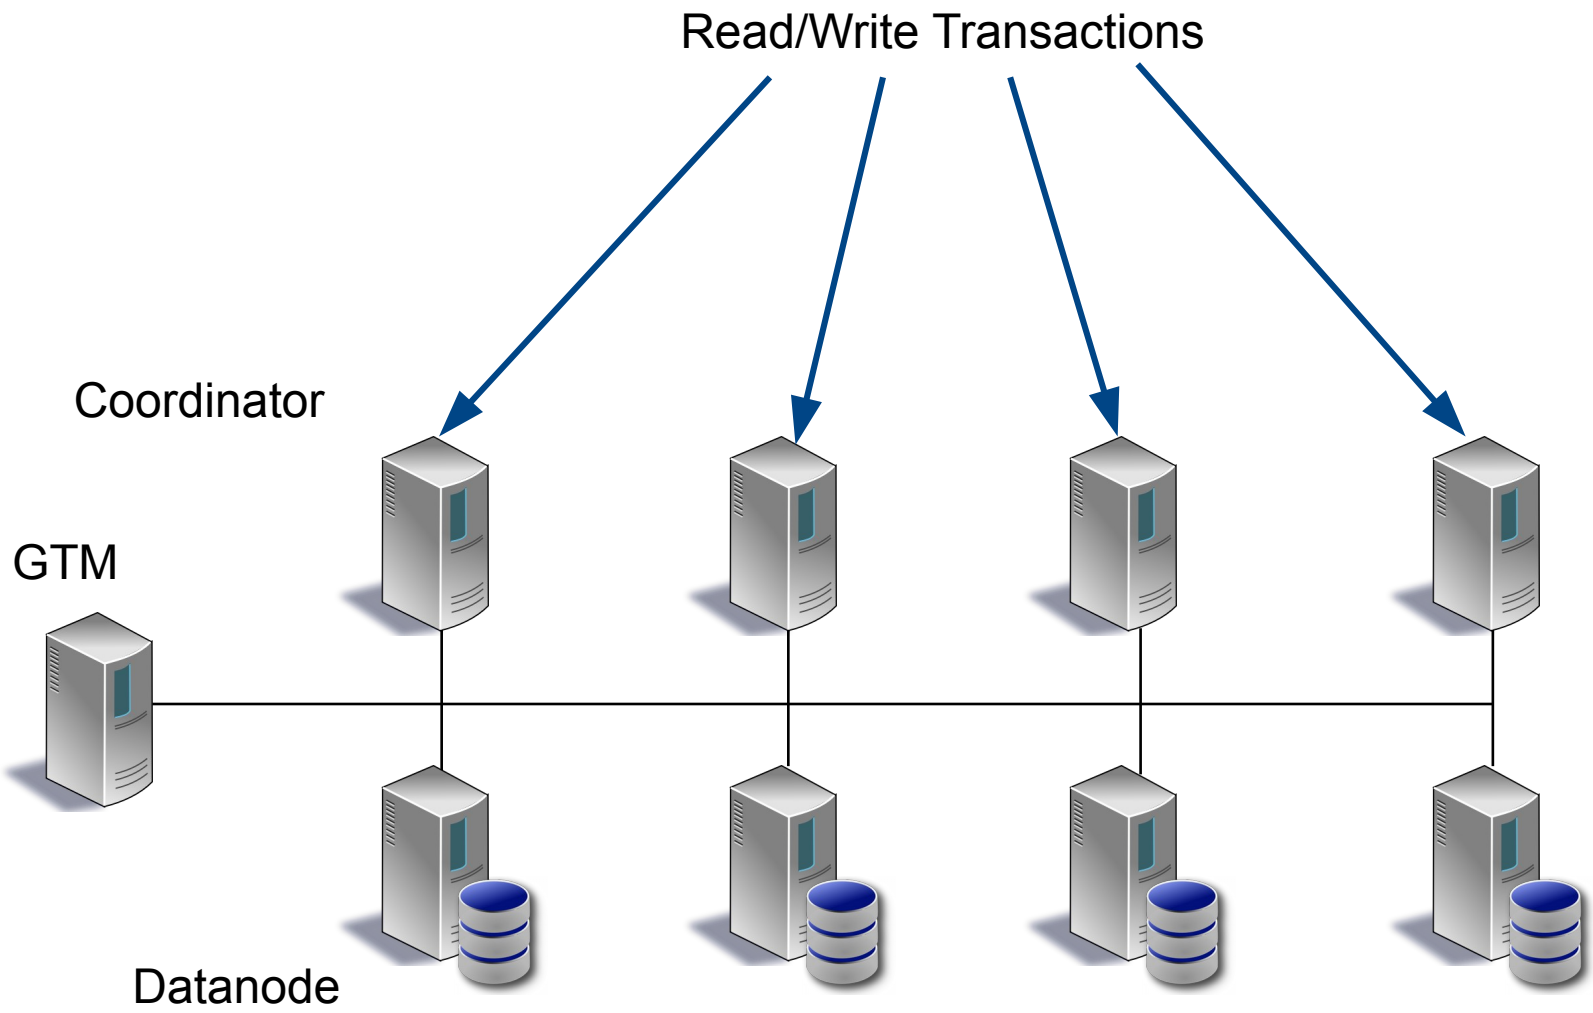

Configuration in Practice

Just like configuring many database servers to talk each other

- Many pitfalls
- Pgxc_ctl provides simpler way to configure the whole cluster
  - Provide only needed parameters
  - Pgxc_ctl will do the rest to issue needed commands and SQL statements.

– Visit
`http://sourceforge.net/p/postgres-xc/xc-wiki/PGOpen2013_Postgres_Open_2013/`

Scalability in OLTP Workloads

OLTP Workload Characteristics

Number of Transactions: Many
Number of Involved Table Rows: Small
Locality of Row Allocation: High
Update Frequency: High

Scaling Out OLTP Workload

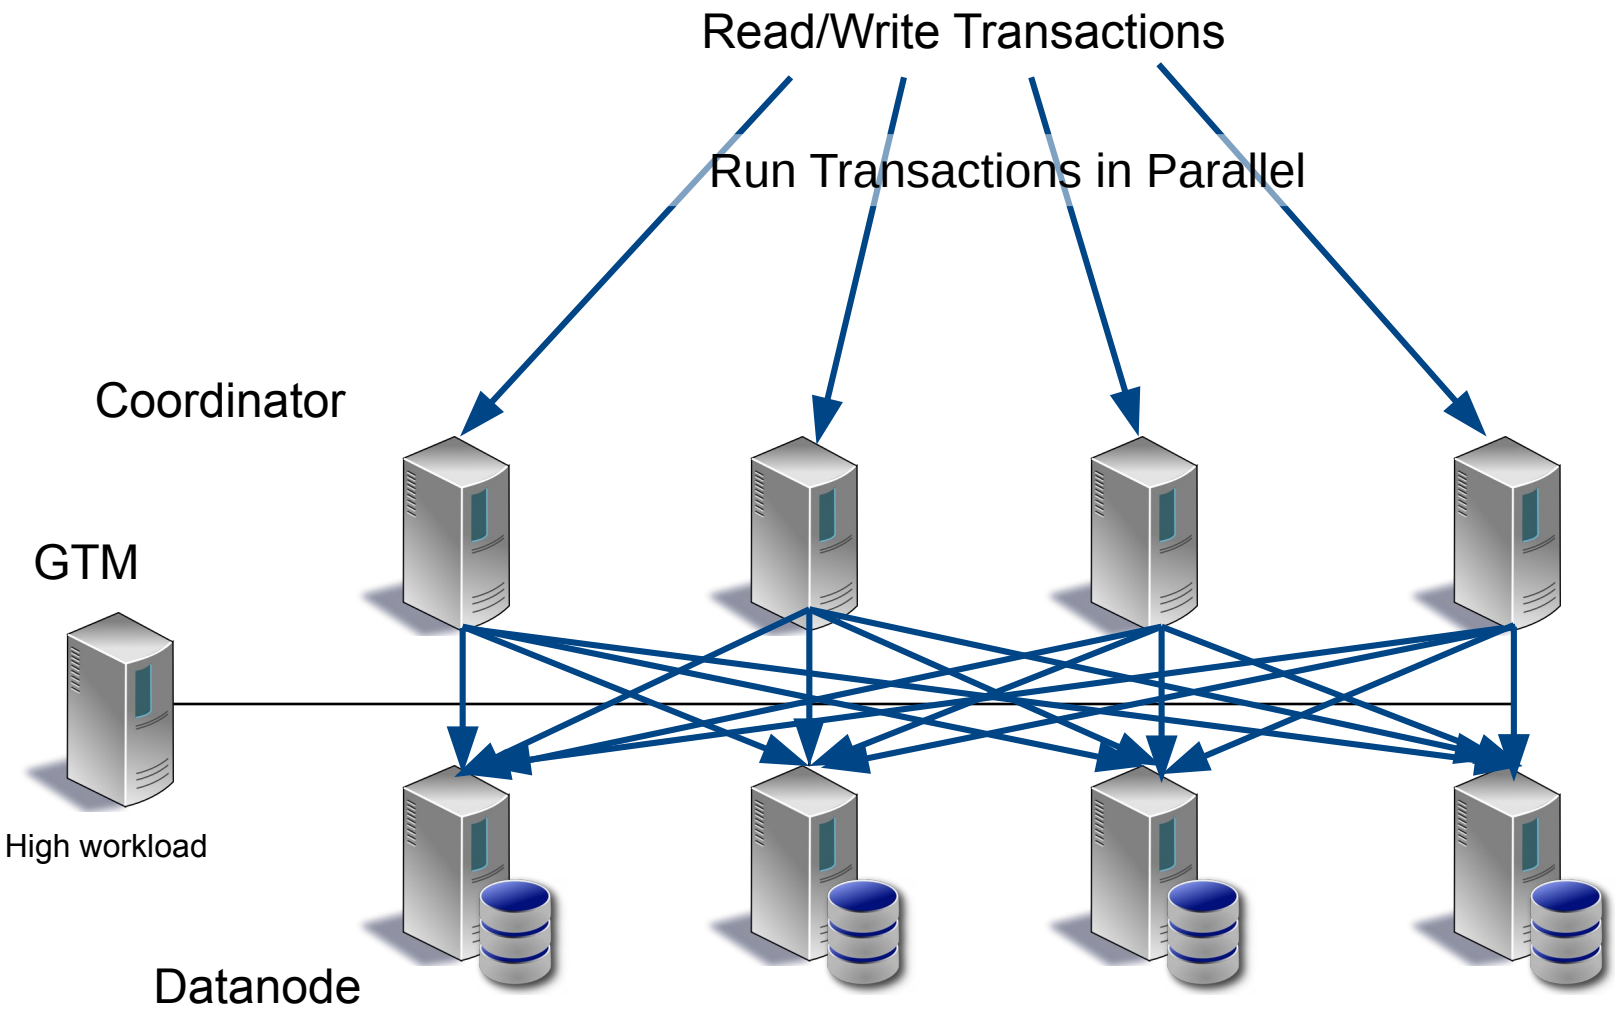

Scalability in OLAP (Analytic) Workloads

OLAP Workload Characteristics

Number of Transactions: Small
Number of Involved Table Rows: Huge
Locality of Row Allocation: Low
Update Frequency: Low

Scaling Out OLAP Workload

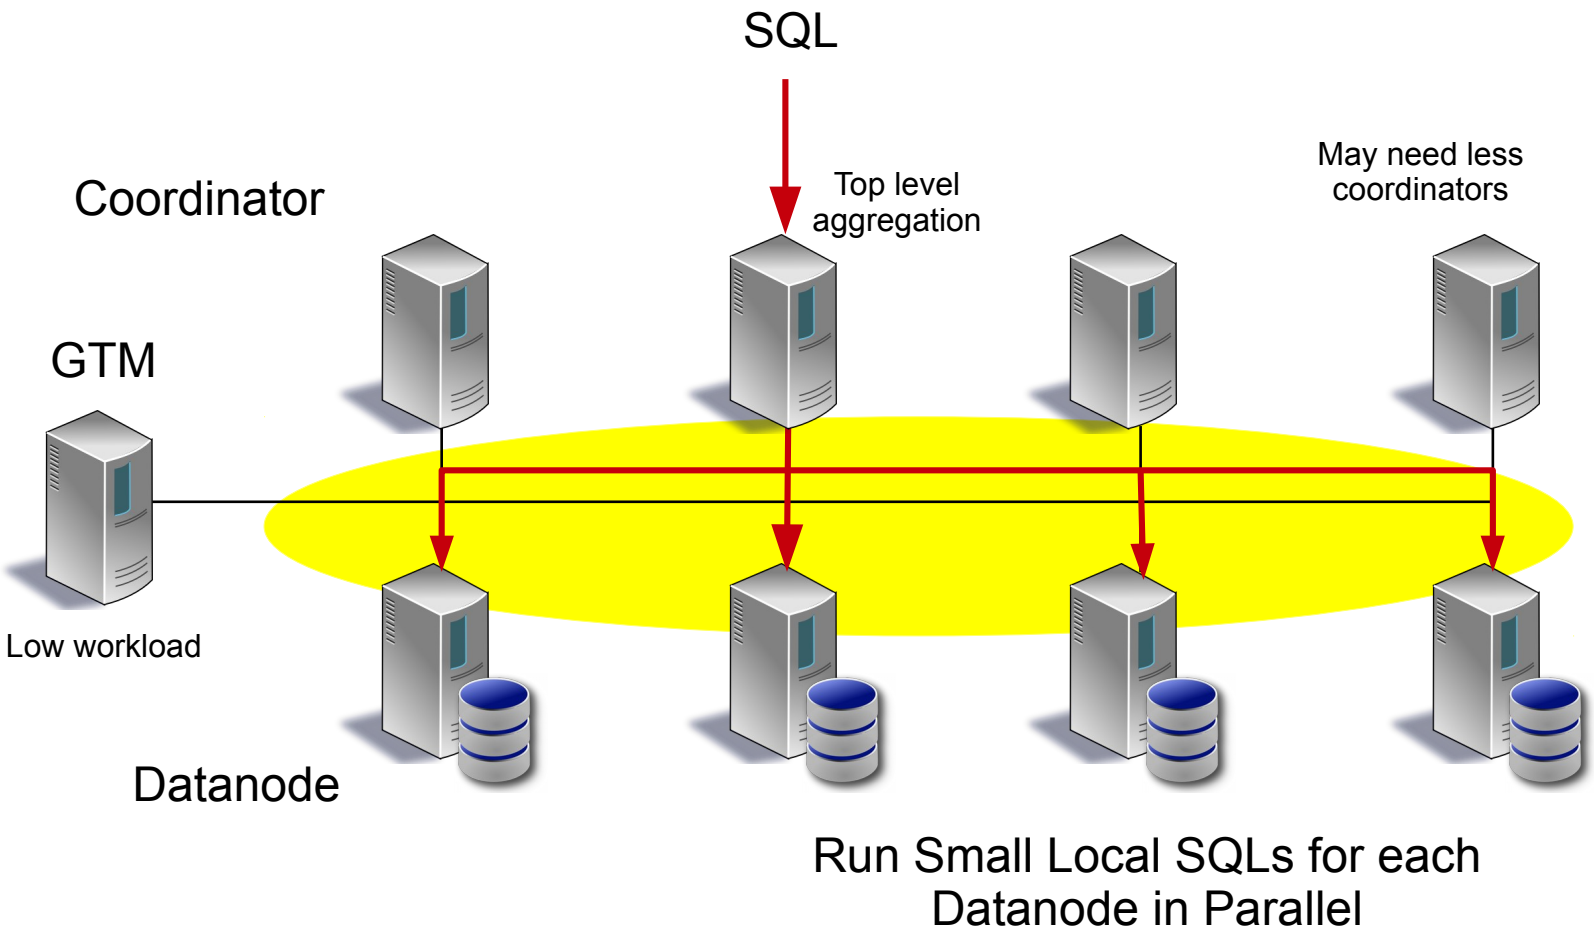

Join Offloading

Join Offloading: When row allocation is available

- Replicated Table and Partitioned Table
  - Can determine which datanode to go from WHERE clause

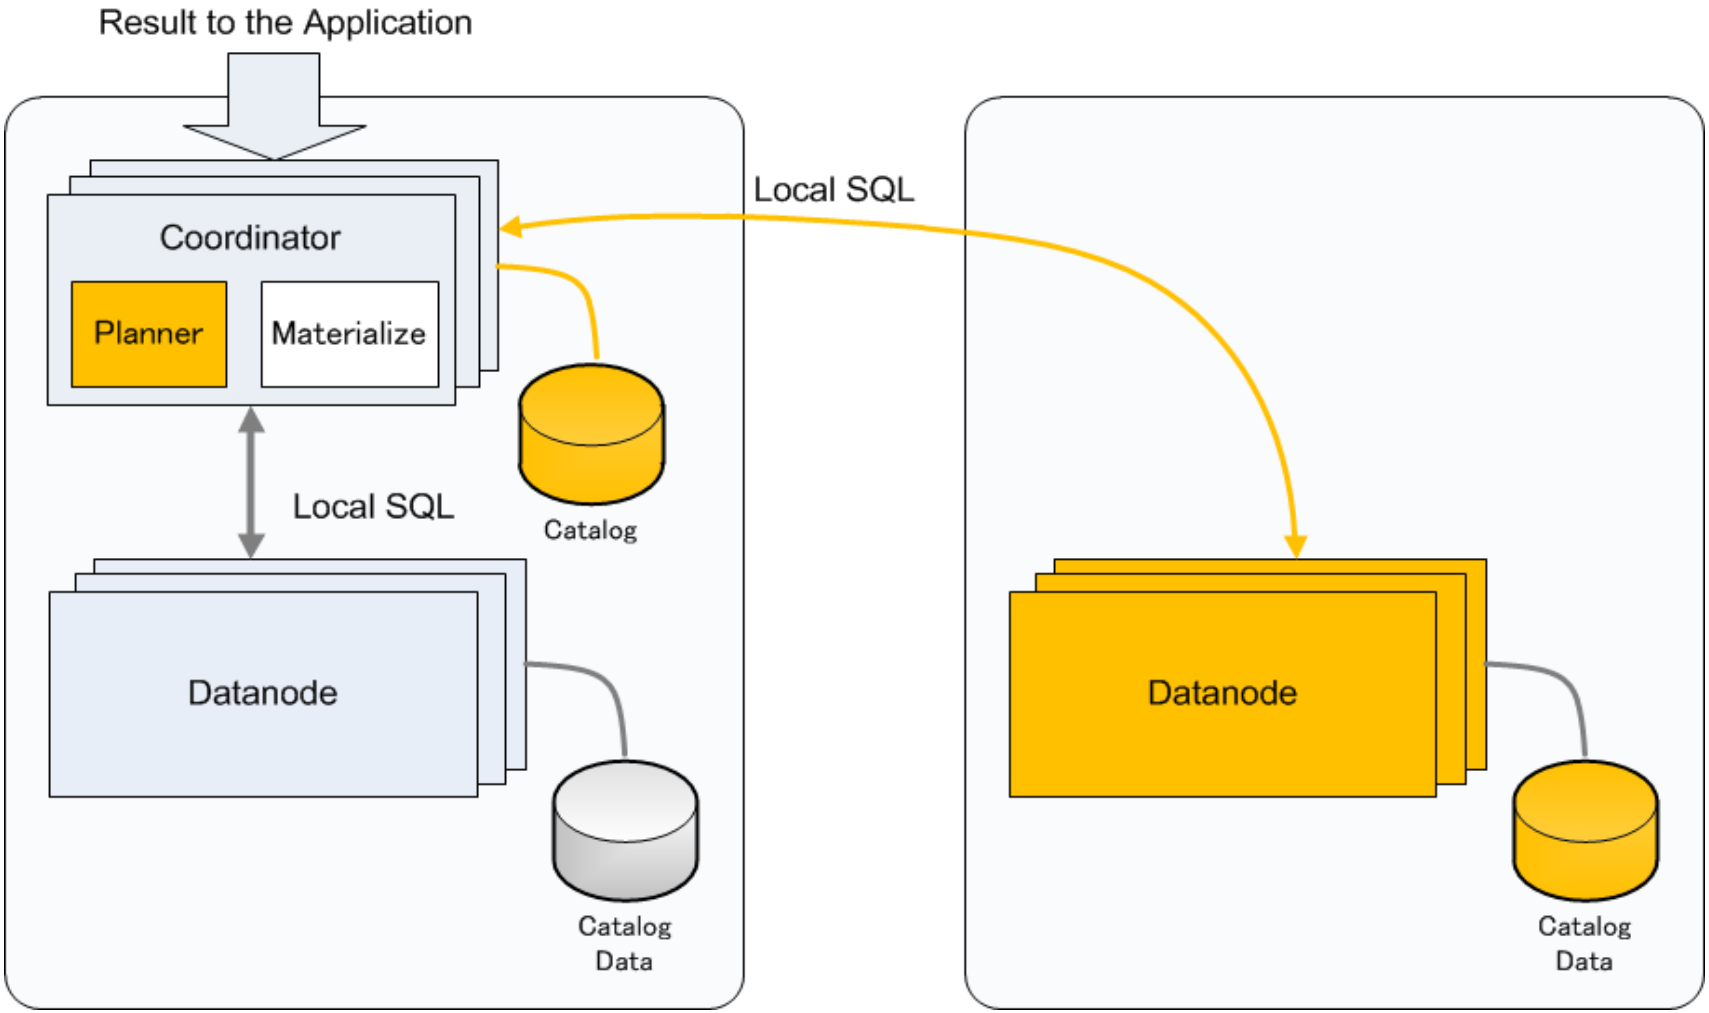

Join Offloading: When row allocation is available

- Replicated Table and Partitioned Table
  - When the coordinator cannot determine which datanode to go from WHERE clause

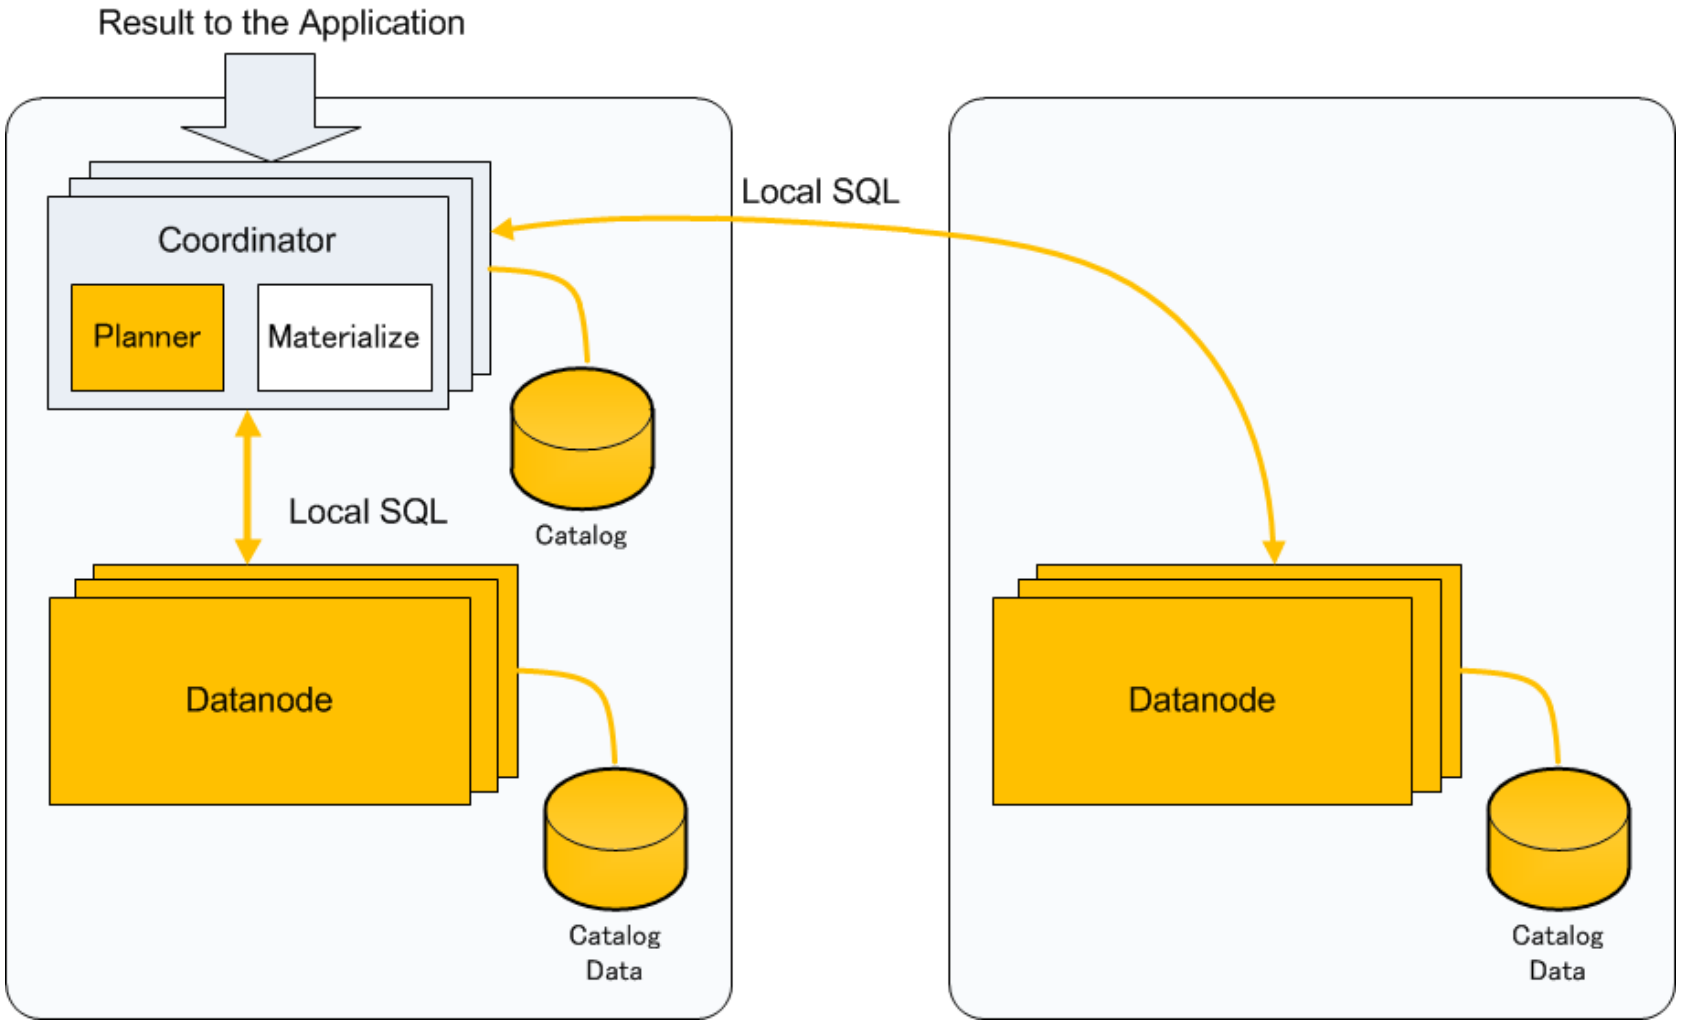

Parallel Aggregation

Aggregate Functions in PostgreSQL

Aggregate Functions in Postgres-XC/XL

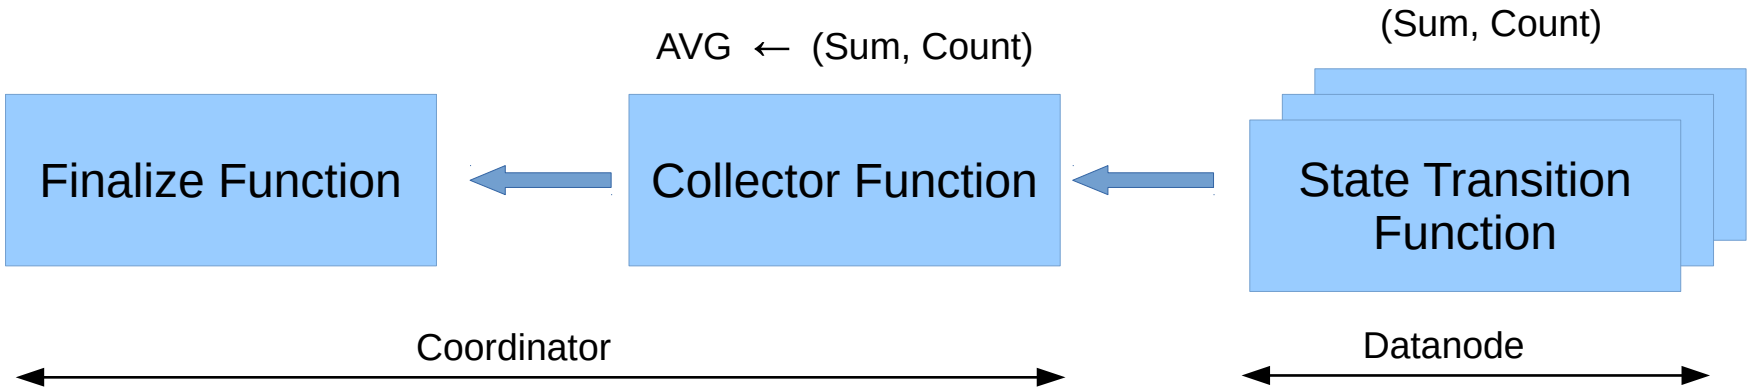

Similar to Map Reduce!

Community status and future

XC and XL community

- Postgres-XC is the original community
  - Based upon PostgreSQL 9.3
  - Tested more for OLPT workload
- Postgres-XL was became separate community for more product-oriented and better stability
  - Based upon PostgreSQL 9.2
  - Shares most of XC code base
  - Tested more for OLAP workload
    - Direct data capture between datanodes
  - Provide many fixes. Most of them apply to XL as well
- Unified again?

Product status

- Source code inherits all the PostgreSQL repository (at some point)
- Fundamental features are all available
  - Global transaction management
  - SQL statements
  - Utilities
- Further challenges
  - Subtransaction (needed for full function support)
  - Catching up PostgreSQL (needed?)
  -

XC and XL community

- Both communities need much more resource to move forward
  - Developer
  - Tester
  - Real workload

XC and XL community sites

Postgres-XC

http://sourceforge.net/projects/postgres-xc/

koichi.dbms@gmail.com

Postgres-XL

http://www.postgres-xl.org/

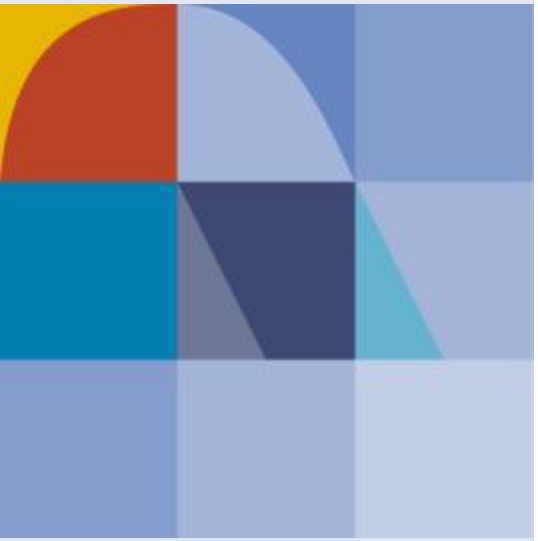

NTT DATA

Global IT Innovator