

Introducing MariaDB Platform X3 and the rise of hybrid everything

Kwangbock LEE (이 광복)
Professional Services Consultant
MariaDB Corporation



Agenda

1. Hybrid workloads
2. MariaDB Platform X3
3. Scalability
4. Real-world use cases
5. Hybrid cloud
6. Database consolidation
7. Getting started with Docker

Hybrid workloads

Database workloads

Transactional

Current data

Range queries


Known queries



Analytical

Historical data

Aggregate queries

Unknown queries







More data

More customers






Database (OLTP)



Data warehouse (OLAP)



Database workloads

Transactional

Current data

Range queries

Known queries

Row-based storage

Indexes

Clustered/Replicated

Analytical

Historical data

Aggregate queries


Unknown queries

Columnar storage


No indexes

Distributed


Application development




Database (OLTP)



BI/reporting + data science




Data warehouse (OLAP)





Transactional

Show me all new products in the science fiction category



Analytical

Show me the top products added to shopping carts or purchased today, and with low inventory.


Actionable insight

I should buy one now because everyone wants one, and they'll be sold out by the end of the day!



Transactional

Show me all pending transactions



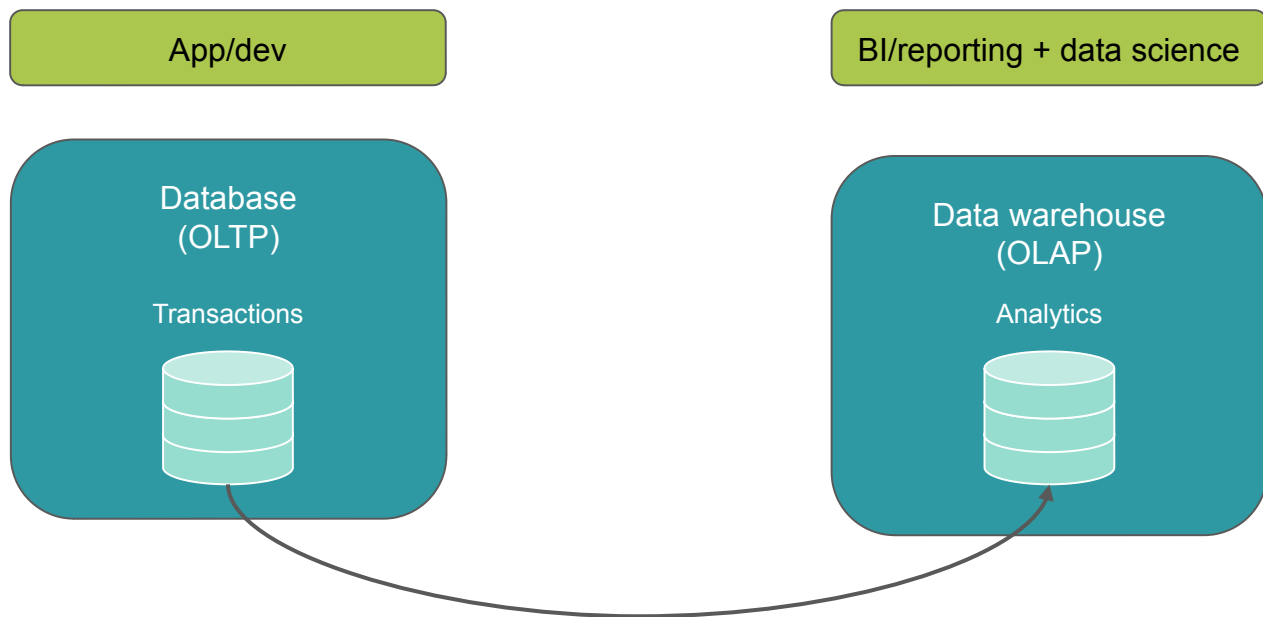
Analytical

Show me when my balance will run low based on historical transactions and my current spending rate

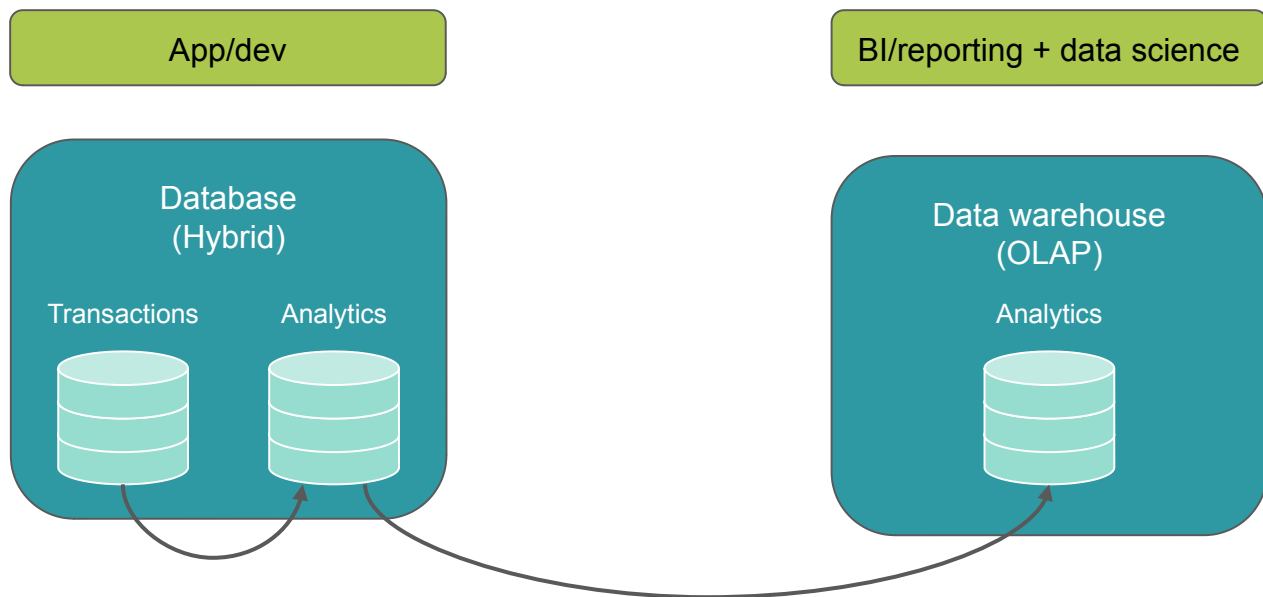
Actionable insight

I should transfer some money from my savings account to my checking account until I get paid again!

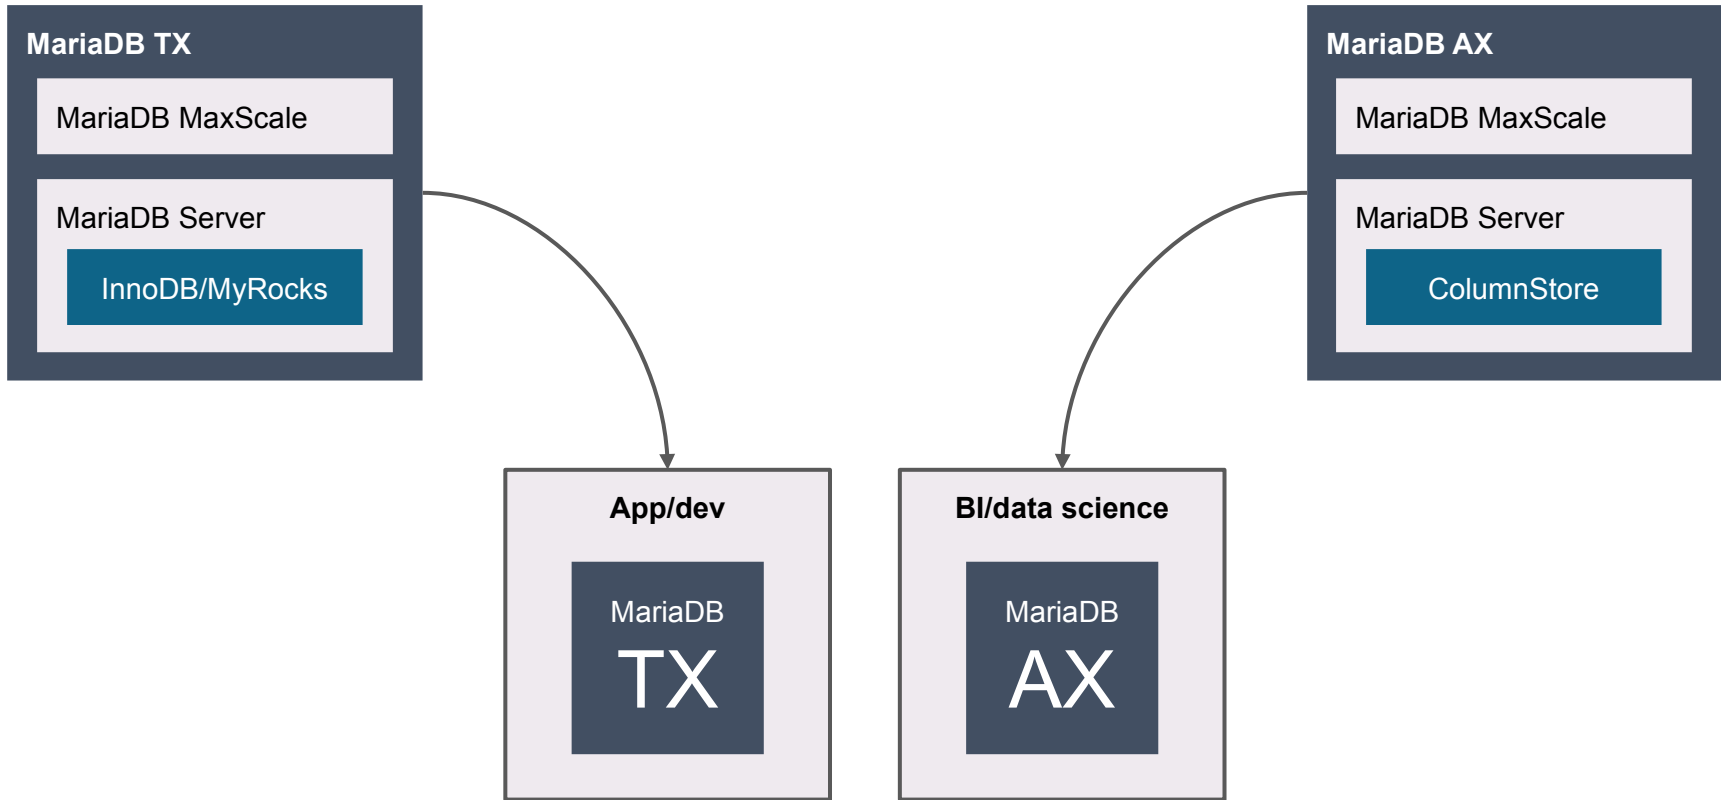
Hybrid workloads: the problem

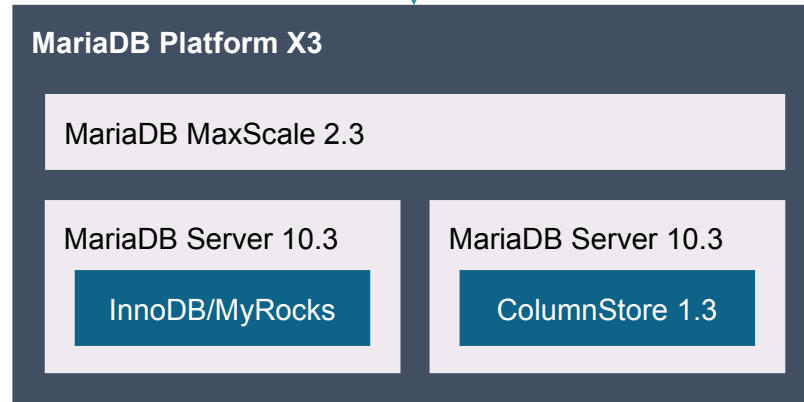
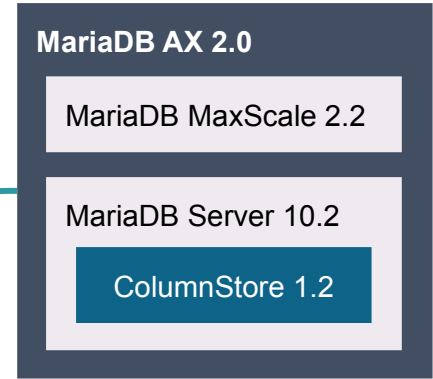
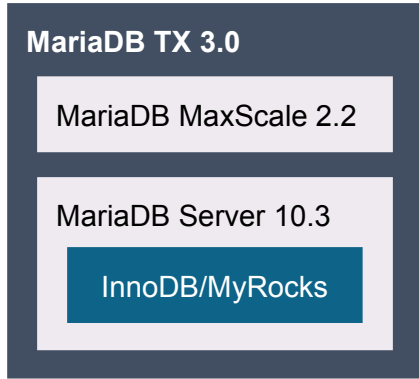


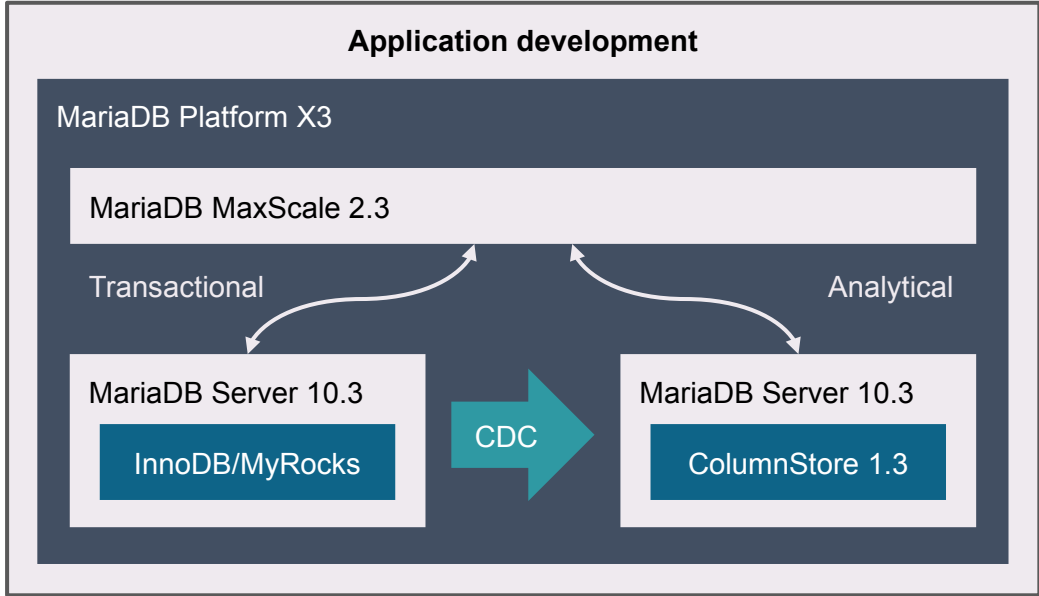
Hybrid workloads: the solution



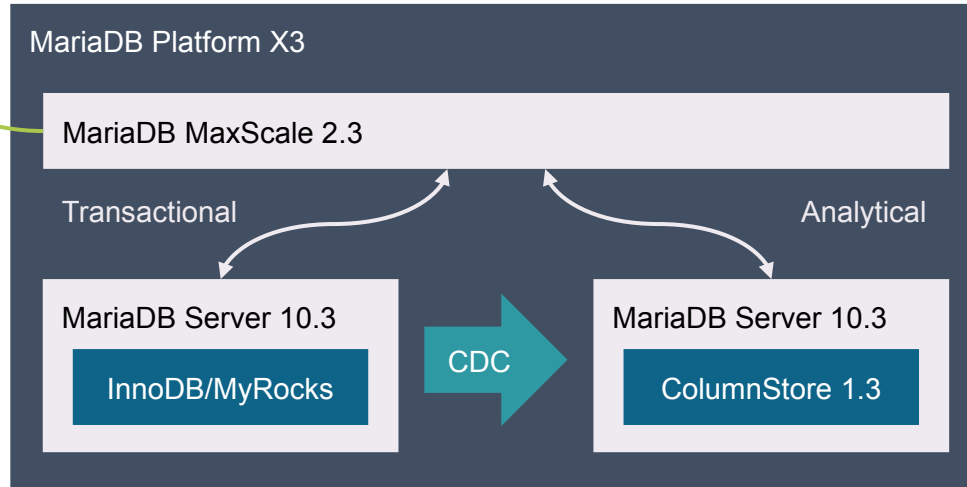
MariaDB Platform X3



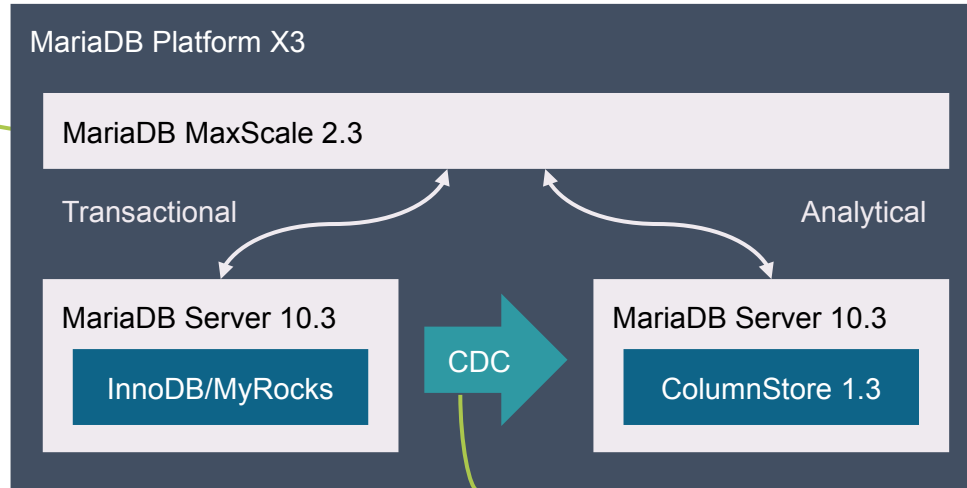




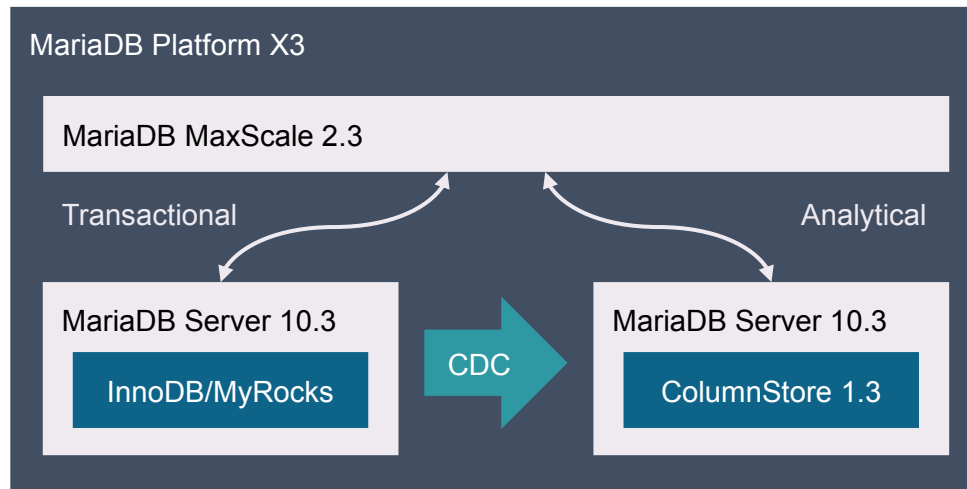
The database proxy inspects queries and routes them to transactional and/or analytical database instances.



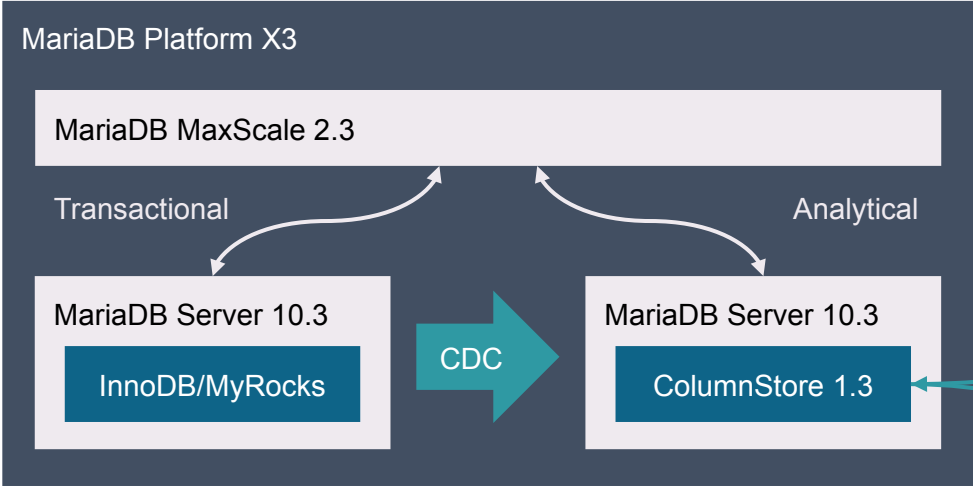




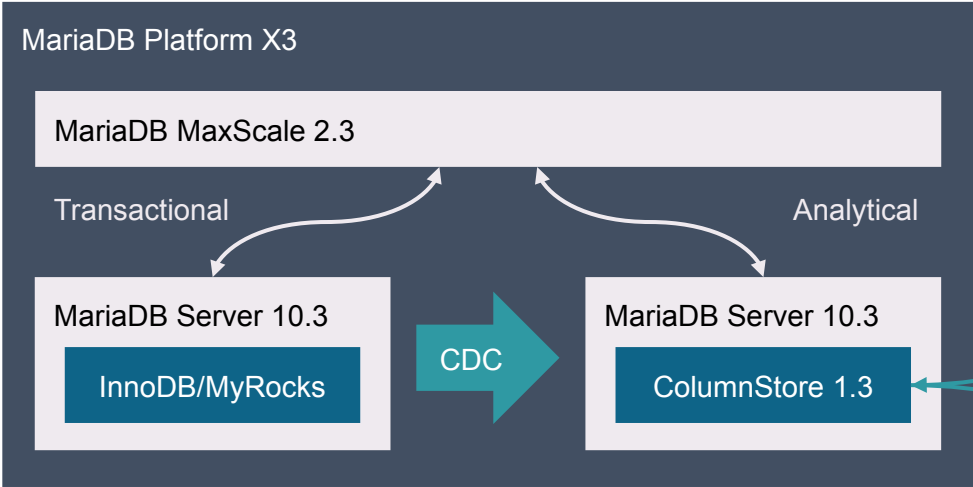

The database proxy inspects queries and routes them to transactional and/or analytical database instances.

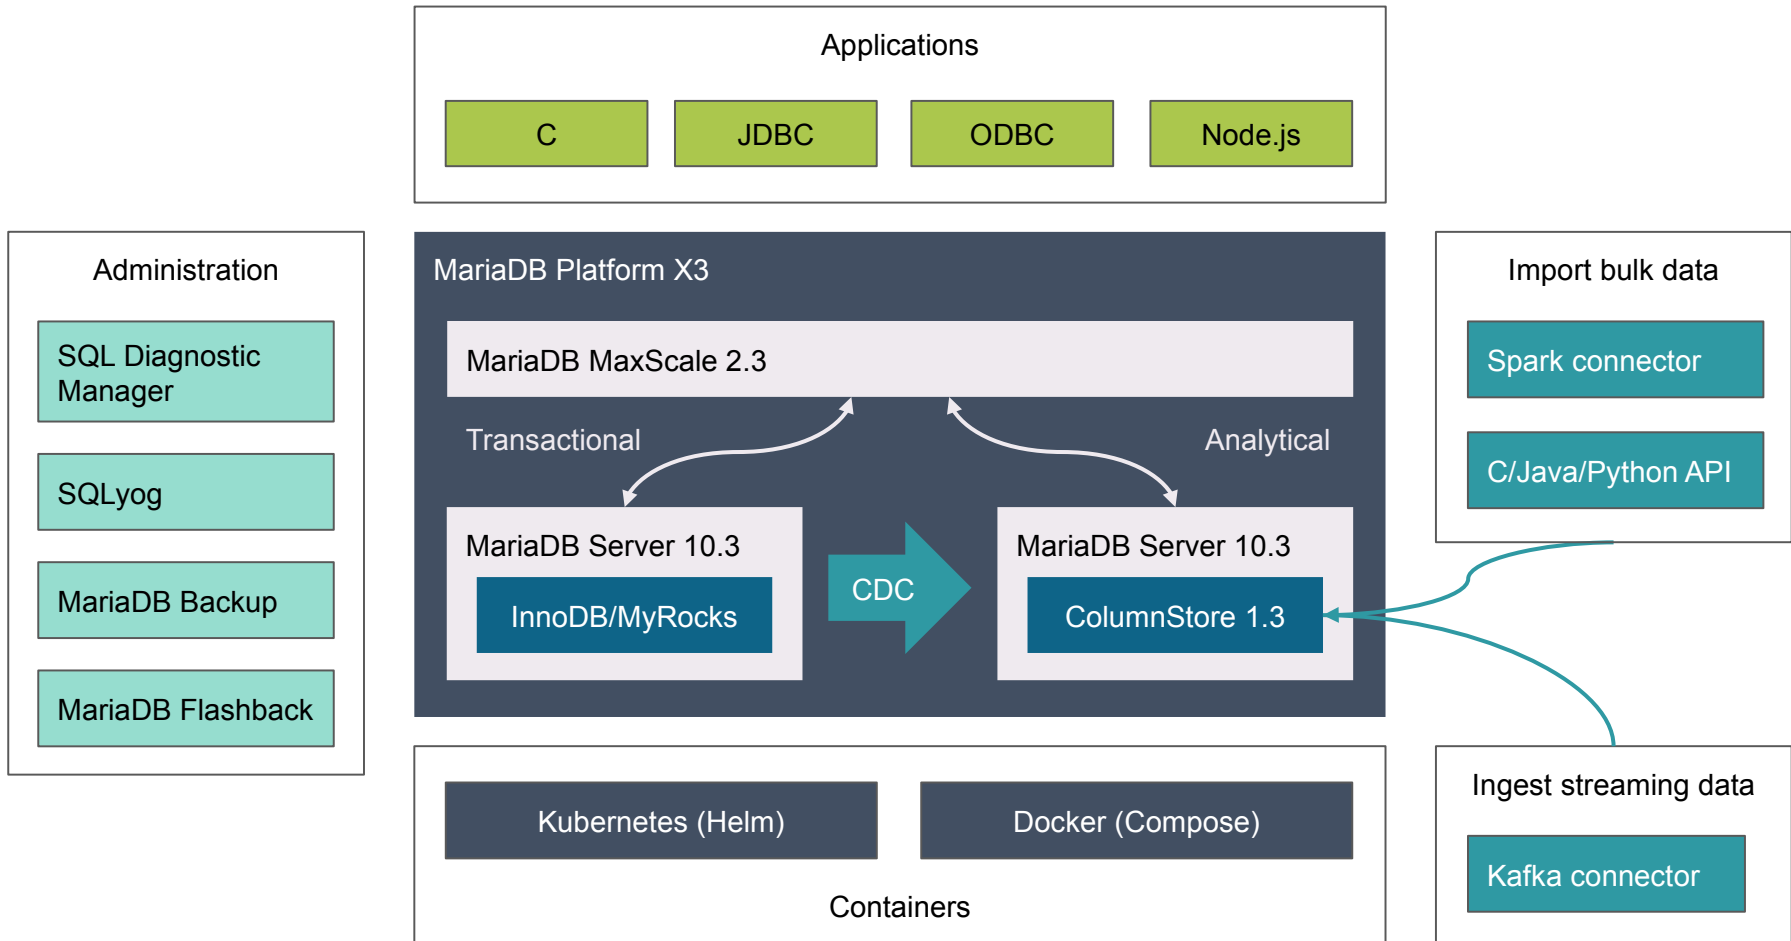


The change-data-capture stream replicates all writes from transactional databases to analytical databases within microbatches.









Scalability

Hybrid workloads: why scalability is needed

Outgrowing OLTP

Applications have transactional and analytical queries

1. Constrained by limited, lightweight analytics
2. Need full analytics to create competitive features

Using historical data

Applications with lots of customers, lots of transactions

1. Limited to current or recent transaction data (months)
2. Need access to all historical data (years)

Exposing analytics

SaaS customers are becoming data-driven organizations

1. They don't have access to their own data
2. They need to analyze it in unknown/unexpected ways

Hybrid workloads: options

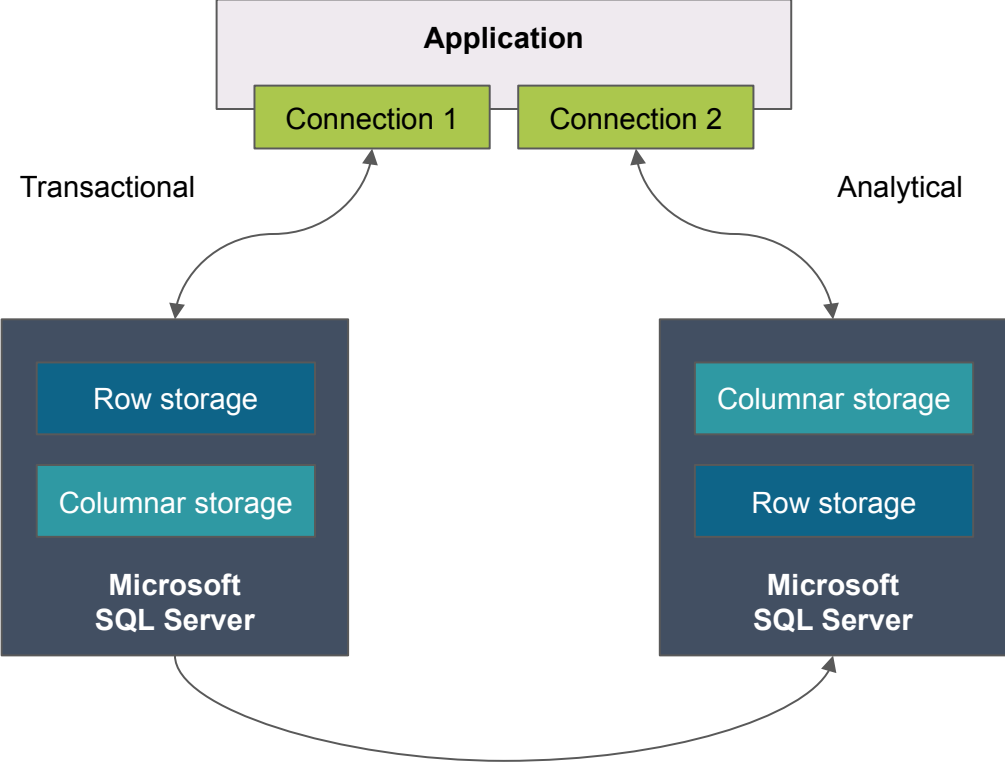
- Oracle
 - In-Memory Column Store
- Microsoft SQL Server
 - ColumnStore Indexes (Clustered or Nonclustered)
- IBM Db2
 - Shadow Tables
- MySQL Enterprise
 - None
- EnterpriseDB Postgres Platform
 - None

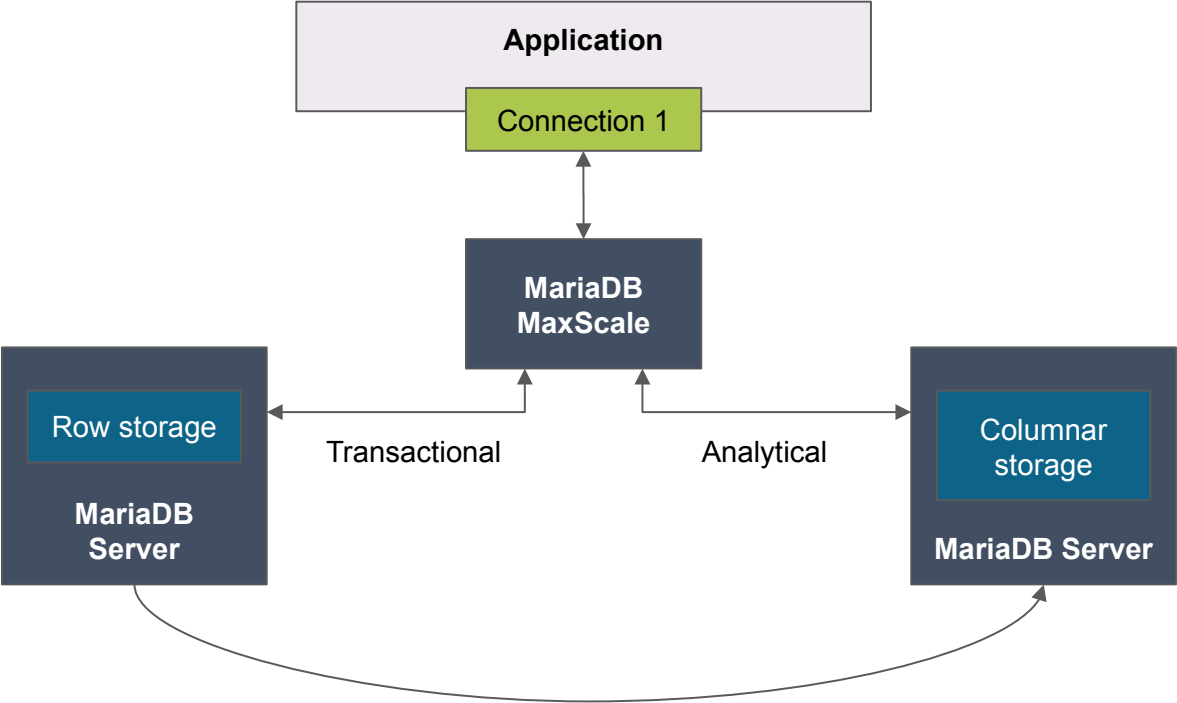
Hybrid workloads: full comparison

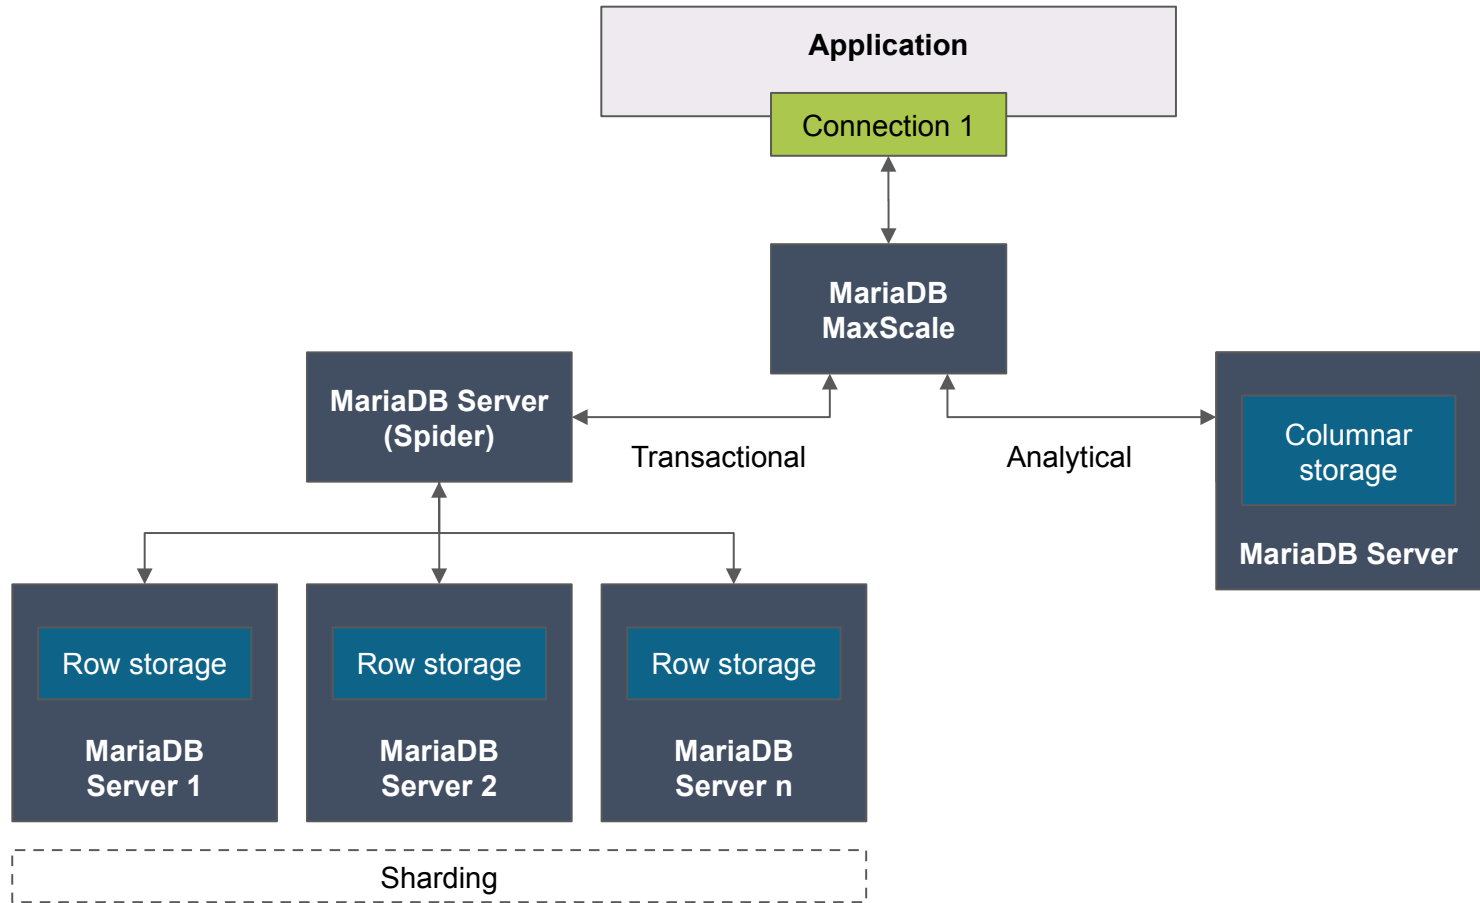
	Oracle	Microsoft	IBM *	MariaDB	MySQL	Postgres
Row storage (OLTP)	Yes	Yes	Yes	Yes	Yes	Yes
Sharded	Yes	No	Yes	Yes	No	No
Columnar storage (OLAP)	Yes	Yes	Yes	Yes	No	No
Disk-based	No **	Yes	Yes	Yes	-	-
Distributed	Yes (RAC)	No	No	Yes	-	-

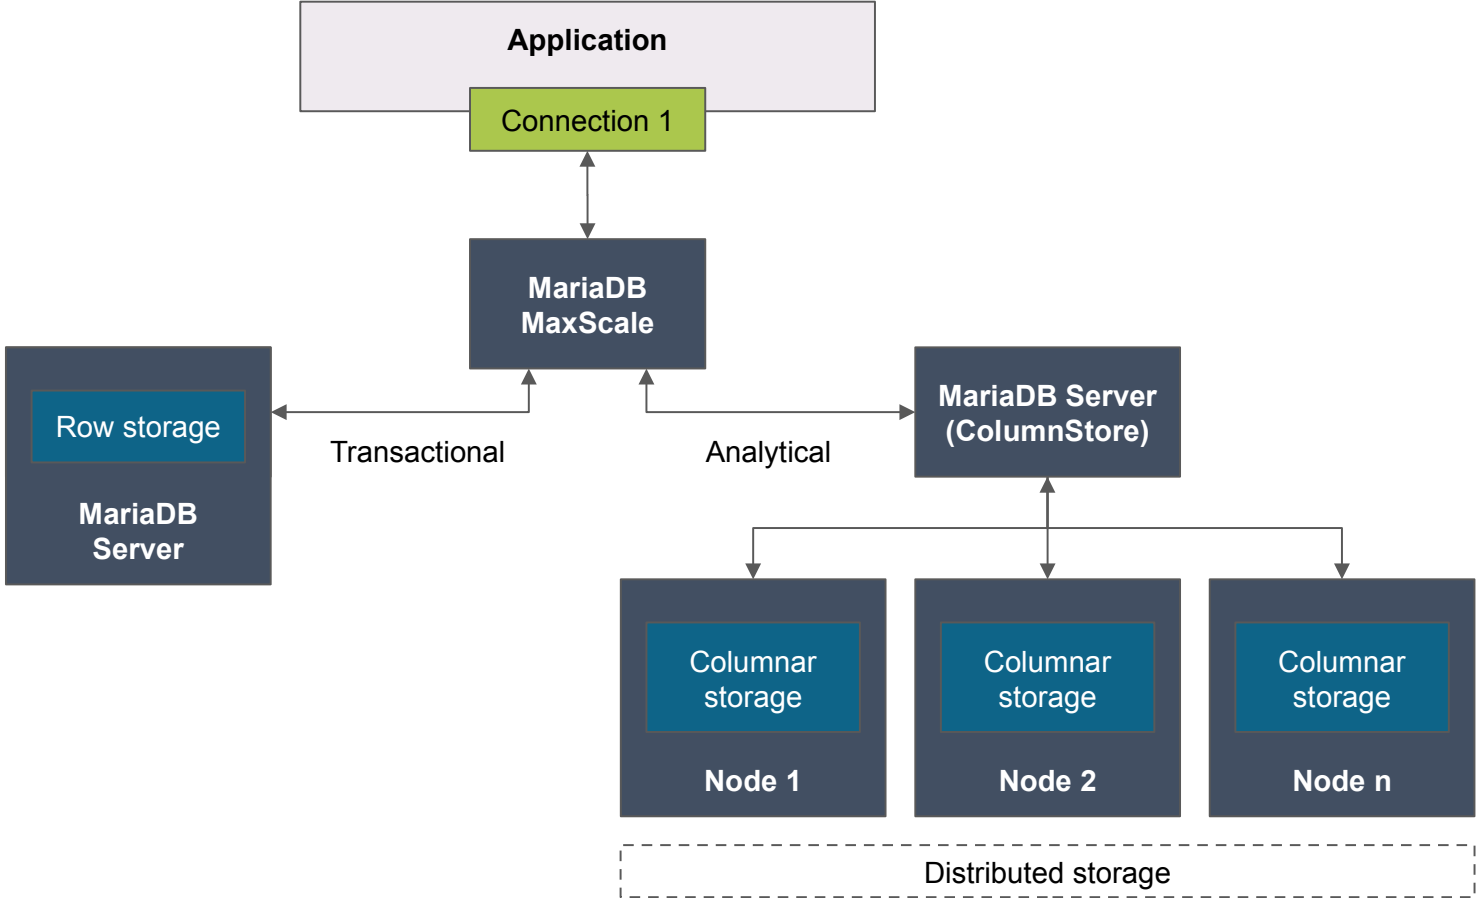
* IBM Db2 Shadow Tables, ** Oracle IM column store can read from on-disk row storage if needed

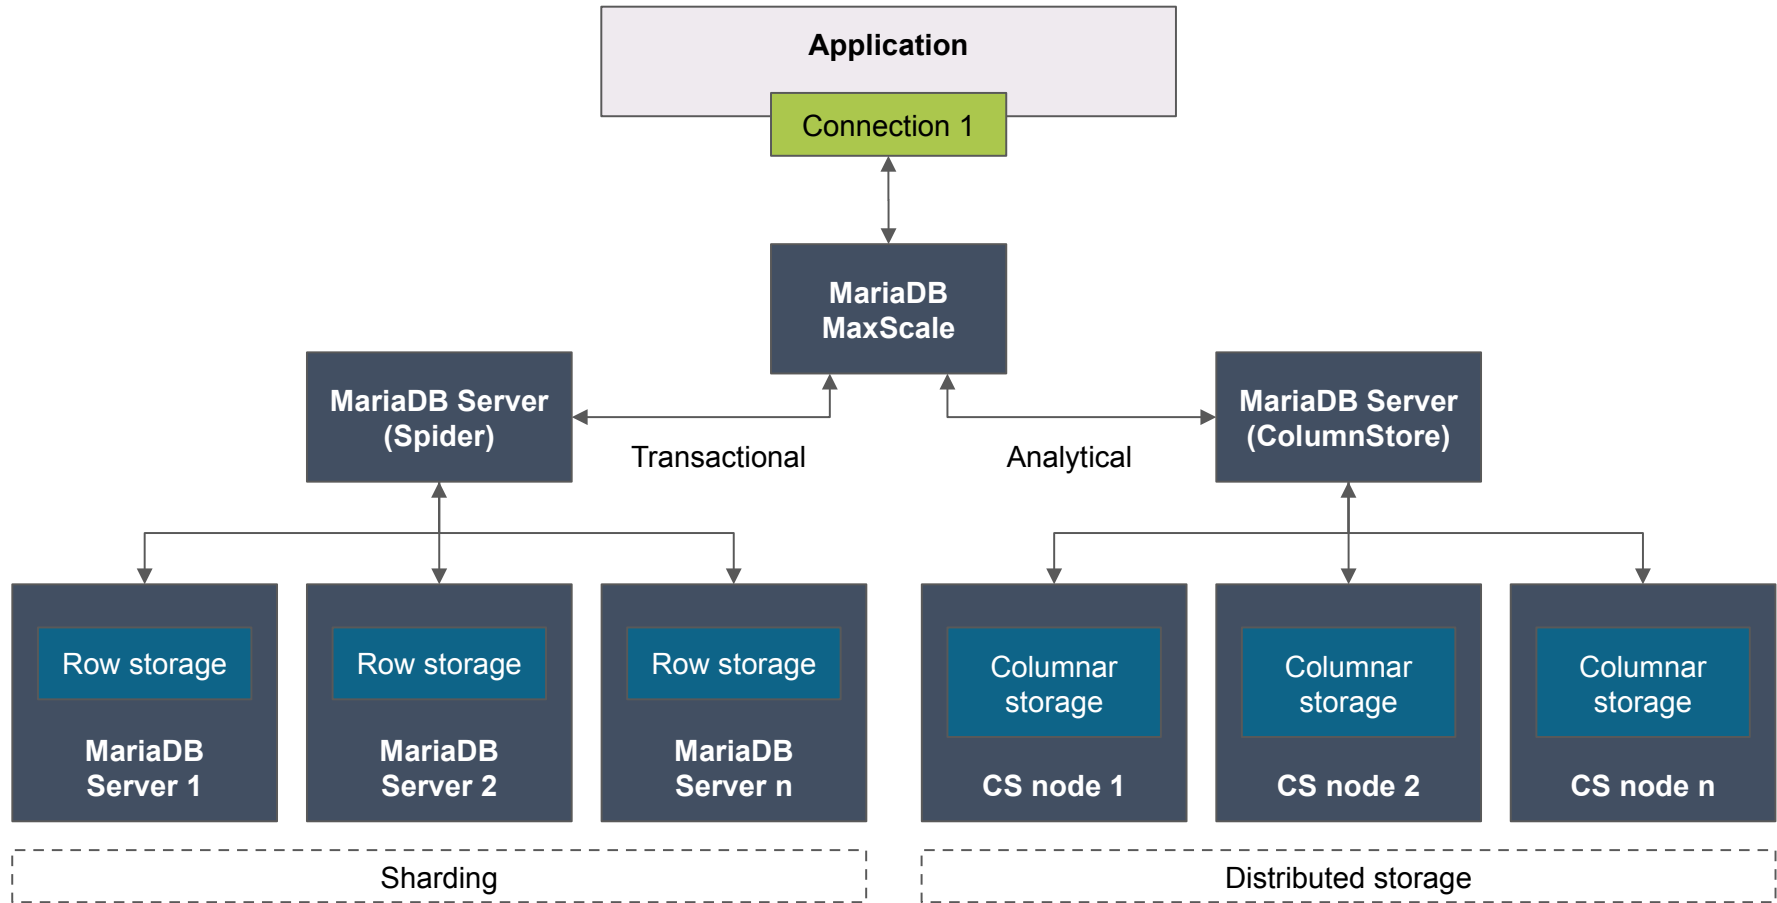










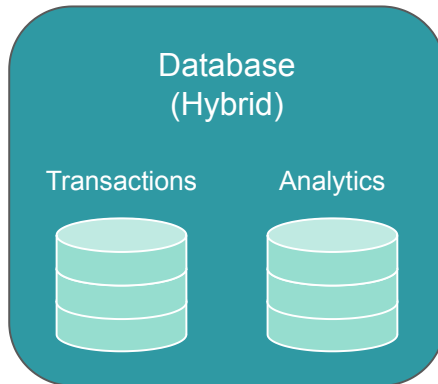


Real-world use cases

Retail (market research)

Transactional

Capture daily product prices
Update product information
Show current prices
Store current pricing data




Analytical

Show prices over time
Self-service analytics
Store historical pricing data

Telecommunications (IP telephony – SaaS)

Transactional

Capture call detail records
Charge by call/message
Generate bills



Analytical

Monitor usage
Identify peak periods
Estimate costs
Self-service analytics

Hybrid cloud

Hybrid workloads: perfect for hybrid cloud

- MariaDB Platform separates and isolates different workloads
 - Run different workloads on different infrastructure
 - Place different workloads closer to different users
 - Scale different workloads on different hardware

Public cloud
for analytics

Incremental path to
the cloud, places
analytics close to
customers

MariaDB Platform

MariaDB Server

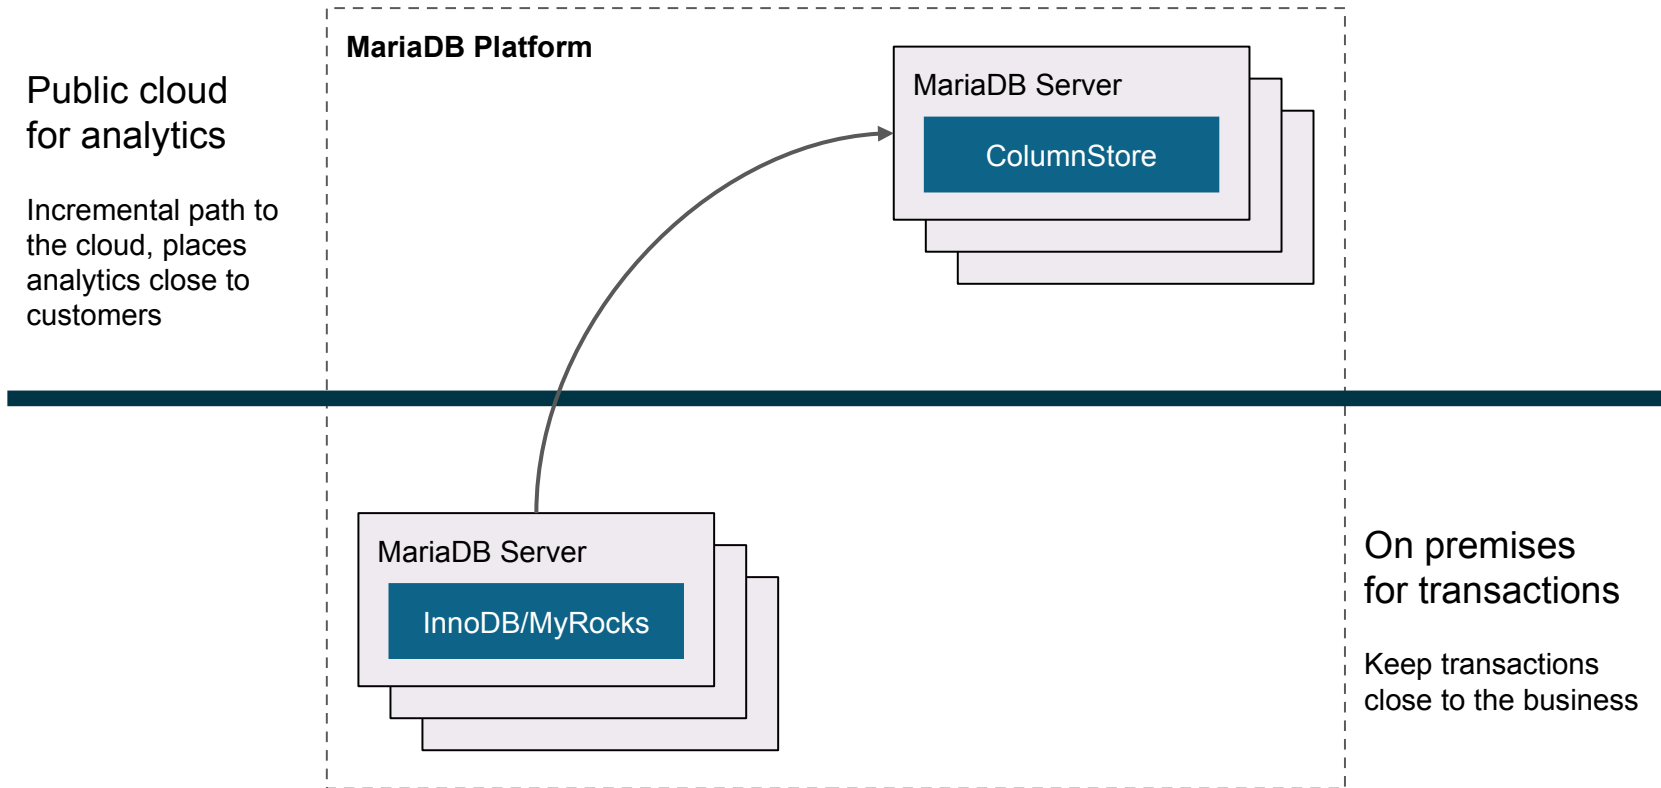
ColumnStore

MariaDB Server

InnoDB/MyRocks

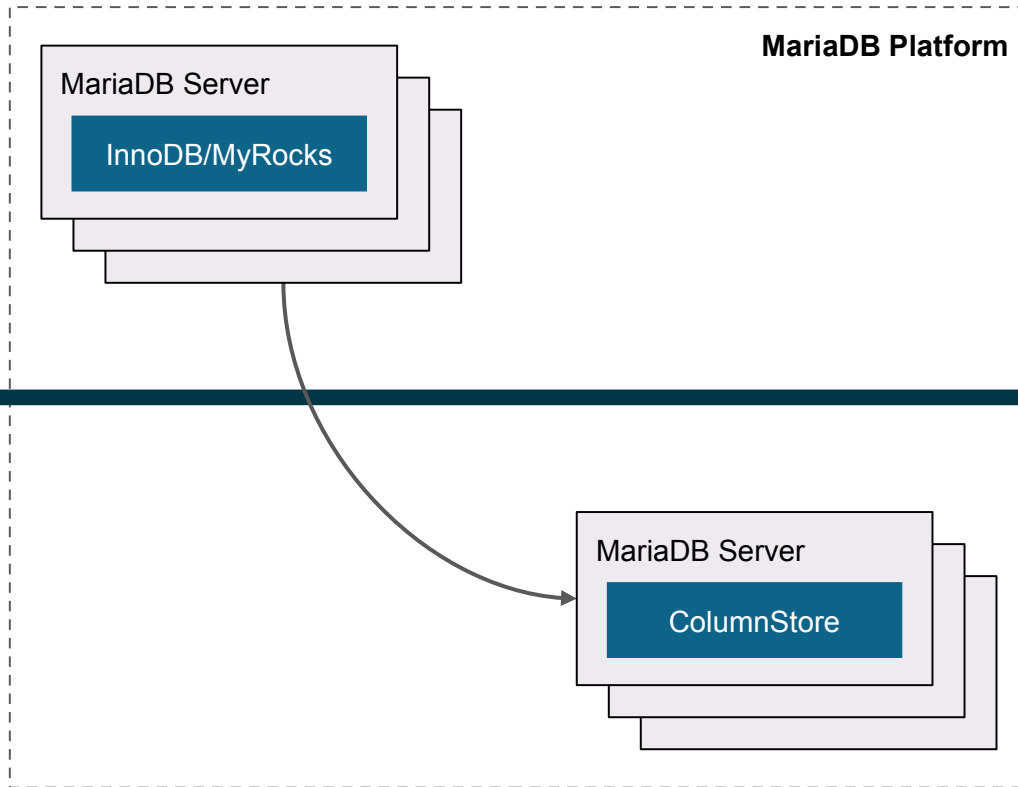
On premises
for transactions

Keep transactions
close to the business



Public cloud
for transactions

Place transactions
close to customers,
geographically
distributed



MariaDB Platform

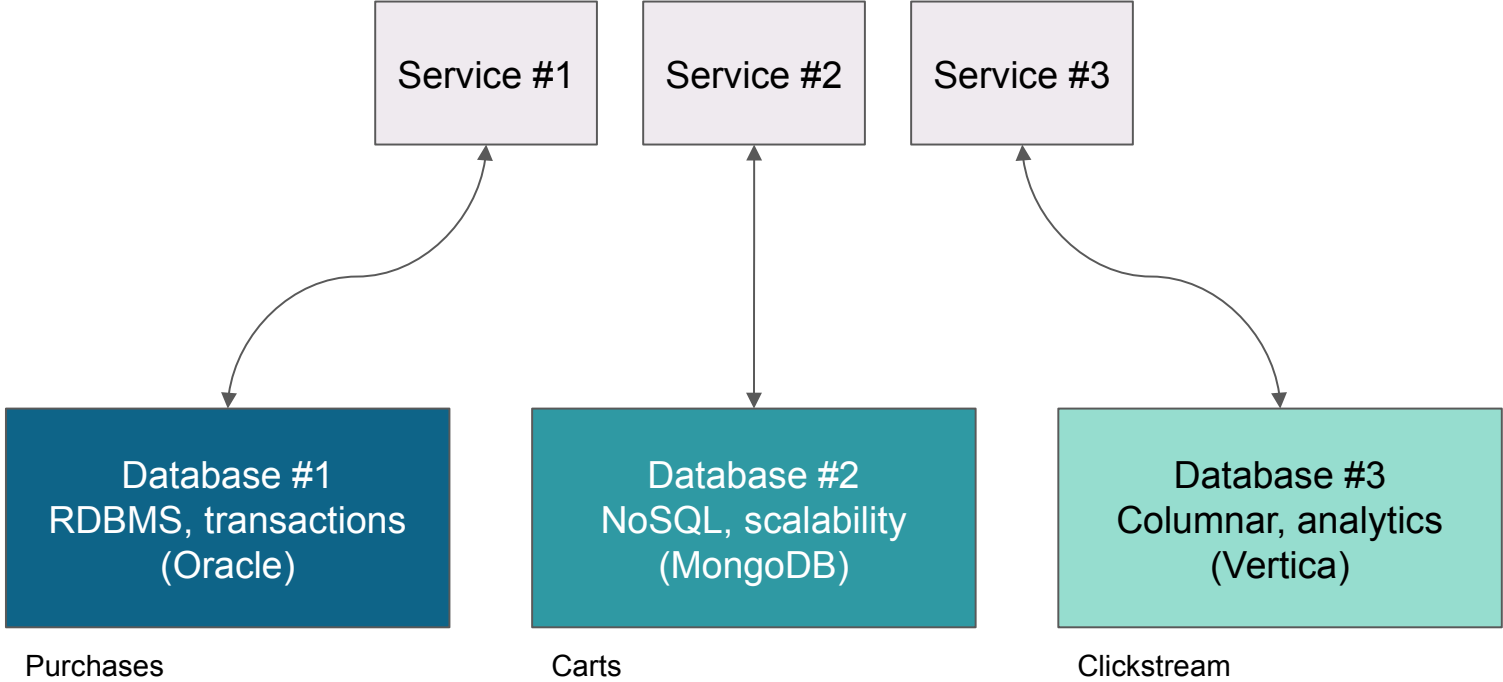
MariaDB Server

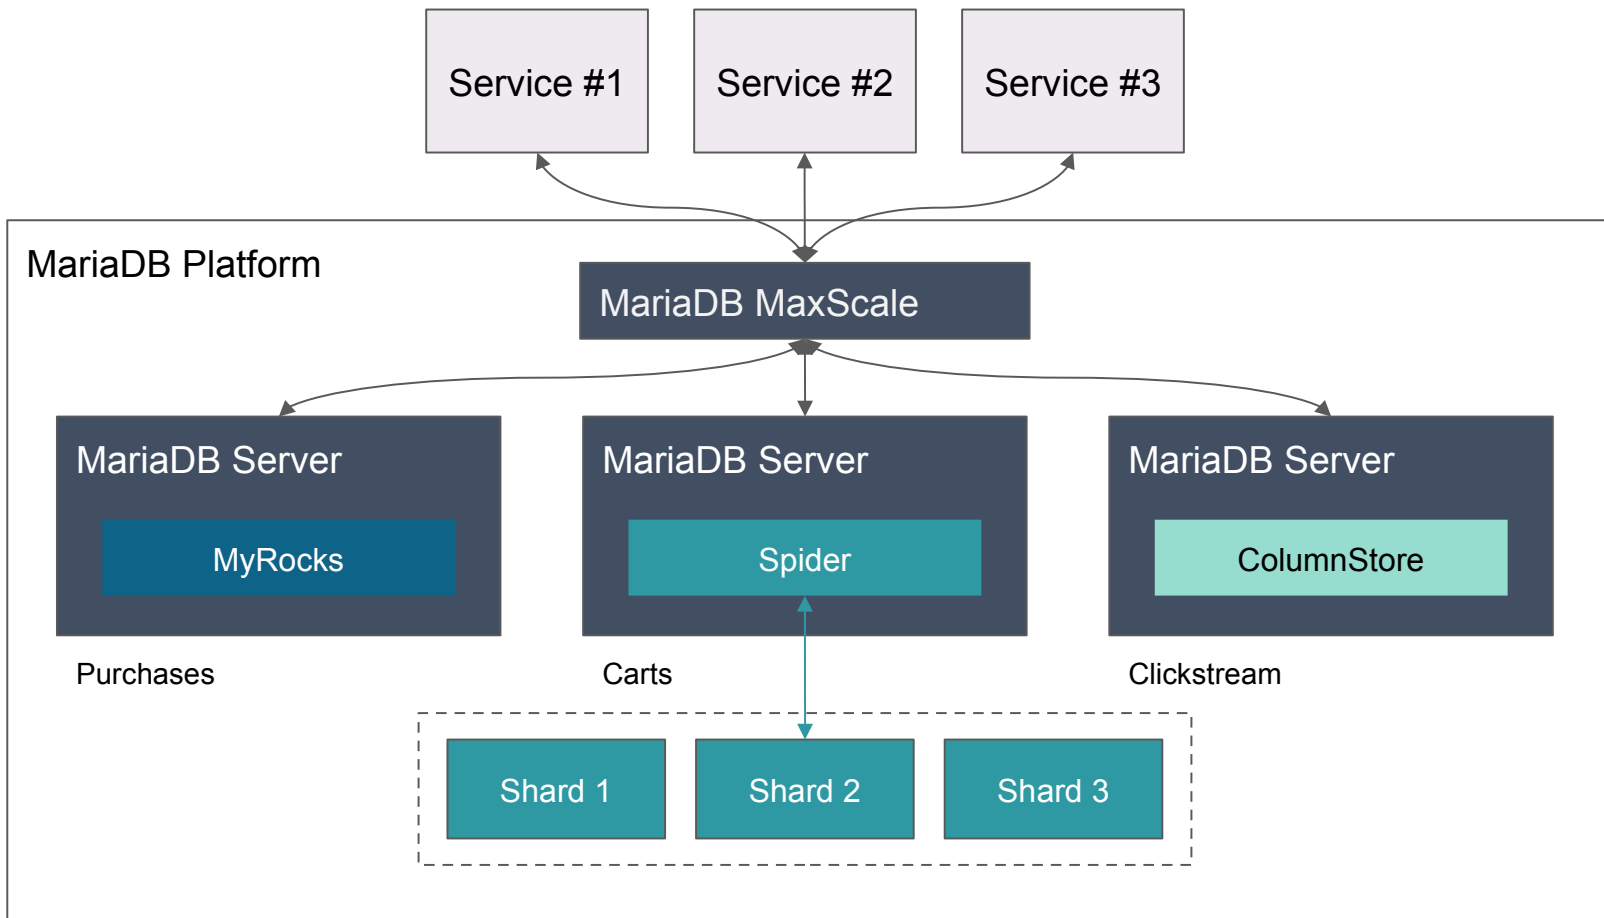
ColumnStore

On premises
for analytics

Place analytics
closer to employees,
aggregate all
transactions

Database consolidation





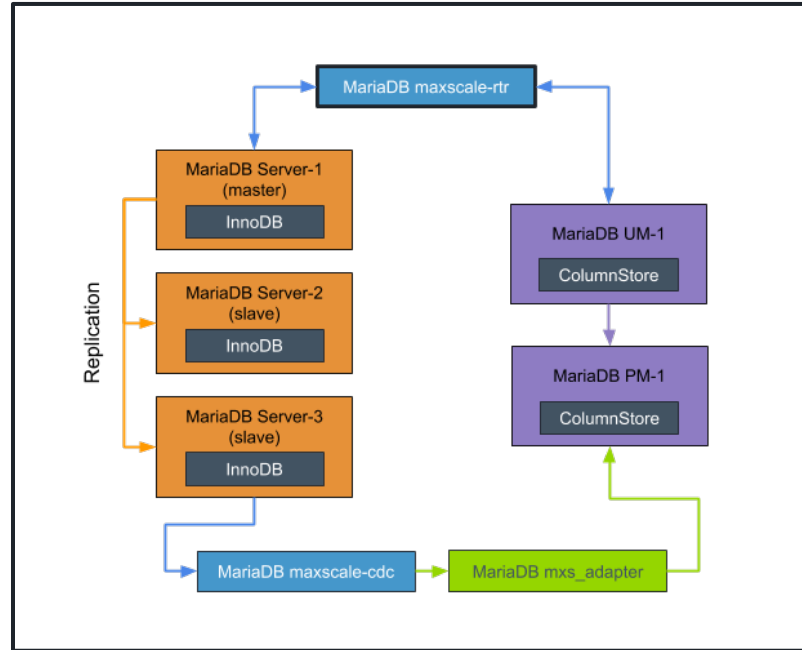
Getting started with Docker

MariaDB Platform X3 in a container

- MariaDB Platform “in-a-box”
- Fully-configured and ready to use out of the box
- Intended as a quickstart for development
- Launched with a Docker command
- Runs on a laptop

<https://github.com/mariadb-corporation/mariadb-platform-docker/tree/master/single-container>

MariaDB Platform X3 in a container



THANK YOU!