

Festival Online de la Data



Synaltic

PROCHAINES DATES

Jeudi 02 juillet de 15h à 15h30

MariaDB Columstore

Sébastien Giraud, MariaDB

Laurent Madec, MariaDB

Jeudi 9 juillet de 15h à 15h30

Sensibilisation à la DataVisualisation

Franck Nguyen, Synaltic

SYNALTIC EN QUELQUES MOTS

Acteur innovant, dénicheur de solutions et une équipe de collaborateurs engagés



30

30 collaborateurs formés et certifiés contribuant aux communautés d'utilisateurs.

15

Spécialiste en Data Management depuis plus de 15 ans et plus de **250 projets** réalisés.

160

Plus de **160 clients** dont certains s'appuient sur nous avec succès depuis plus de **8 ans**.

Experts en Data Management, passionnés d'Open Source et d'Innovation !

MARIADB & SYNALTIC

Plus de 10 ans de collaboration



La base de données à la croissance la plus rapide du marché.

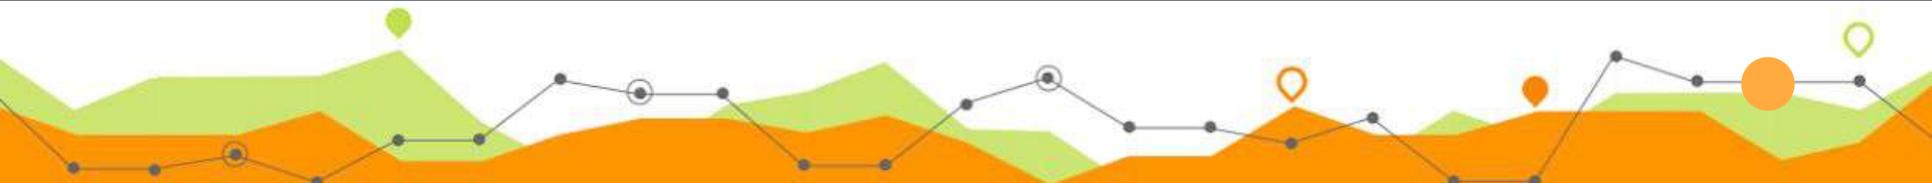


MariaDB a une approche fondamentalement différente des bases de données pour s'adapter au monde d'aujourd'hui.



Allégez votre budget en remplaçant vos bases de données propriétaires. Cela permet un développement rapide d'applications clients innovantes.

**Synaltic soutient MariaDB depuis plus de 10 ans.
Optez pour MariaDB avec Synaltic !**



Synaltic présente MariaDB

Sebastien Giraud - Solution Engineer
Laurent Madec - Senior Sales Manager
MariaDB Corporation



Agenda

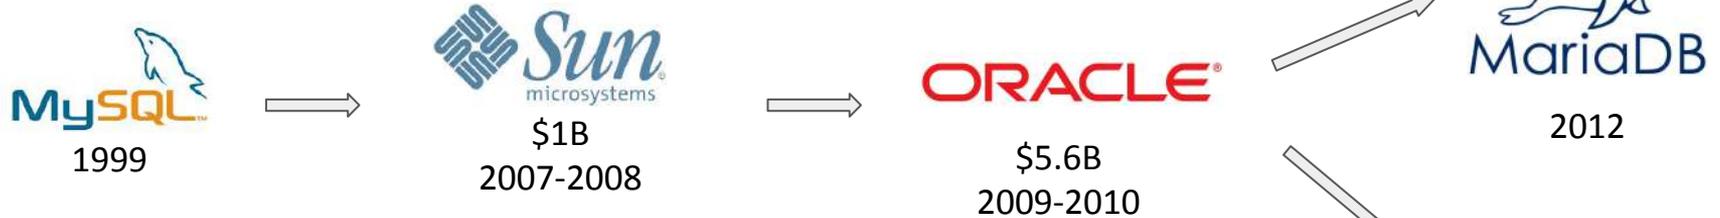
- Présentation de MariaDB
- La vision MariaDB
- Les moteurs de stockage proposés par MariaDB
- L'analytique à haute valeur ajoutée grâce à MariaDB
- Les smart transactions (HTAP)
- Annexe

Présentation de MariaDB



The Soul of Open Source

“ MariaDB was created to preserve openness and community, so that we can push ahead faster with the capabilities for tomorrow’s applications. ”



MariaDB Corporation -v- MariaDB Foundation

- 10.3.22 VS 10.4.12-6-MariaDB-enterprise-log (2 years old release)
- Community mainly used by Dev people
- Enterprise edition increase security level for production needs
- Home made Enterprise grade features - “*Fait Maison*”
 - Automatic Failover
 - Read/Write splitting
 - Database firewall
 - Consultative and corrective support powered by MariaDB Engineers
 - MariaDB toolset for production
 - Several storage engines (transactional, analytics, sharding, columnstore, highly distributed)

THE RESULTS OF UBIQUITY

- MariaDB is One of the most loved and most used databases (top 5)

(Source: Stackoverflow developer survey with more than 100k participants)

- And high in demand:
 - More than 100M Docker Hub downloads
 - In the top 5 databases on Docker Hub
 - 2M+ direct downloads from MariaDB sites
 - 60M+ circulation via Linux distros

La vision MariaDB



MariaDB architecture

Purpose-built database

Relational
Database
(mixed)

Ex. User account

Wide-column
Database
(write-intensive)

Ex. clickstream

Document
Database
(scalable)

Ex. product datasheets

Columnar
Database
(analytical)

Ex. Orders archives

Huge Relational
Database
(distributed)

Ex. cloud storage solution

APPLICATION CONNECTORS

(C, JDBC/ODBC and Node.js)

BI/REPORTING CONNECTORS

(Tableau)

Queries  Queries

MARIADB MAXSCALE

Platform Services - Query Routing, Security, High Availability

EXTERNAL DATA

TRANSACTIONAL (OLTP)

MariaDB Server

MariaDB Server

MariaDB Server

+ Add Database

Clustered / Replicated / Sharded / Distributed

Change-Data Capture

ANALYTICAL (OLAP)

MariaDB Server

+ Add Database

MariaDB ColumnStore

+ Add Storage

Distributed

CONNECTORS

Spark and Kafka

Informatica & Pentaho

C, Java and Python

ON PREMISES OR PUBLIC/PRIVATE/HYBRID CLOUD

Servers

Virtual Machines

Kubernetes Helm Charts

Docker Images

Les moteurs de stockage proposés par MariaDB



Storages engines included

MariaDB Enterprise Server 10.5

Storage engines	Description
InnoDB	A performant, general purpose storage engine supporting ACID-compliant transactions
MyRocks	A storage engine based on RocksDB, optimized for high writes and low disk usage
Aria	A crash-safe replacement for the older MyISAM storage engine, used internally for system tables
Spider	A specialized storage engine used for sharding table data across multiples Servers
Columnstore	A specialized storage engine used for distributed and scalable analytical processing.
S3 storage engine	A specialized storage engine used for archive with S3 API in any third-party public or private cloud
Xpand	A specialized storage engine used for relational database with extreme scalability requirements

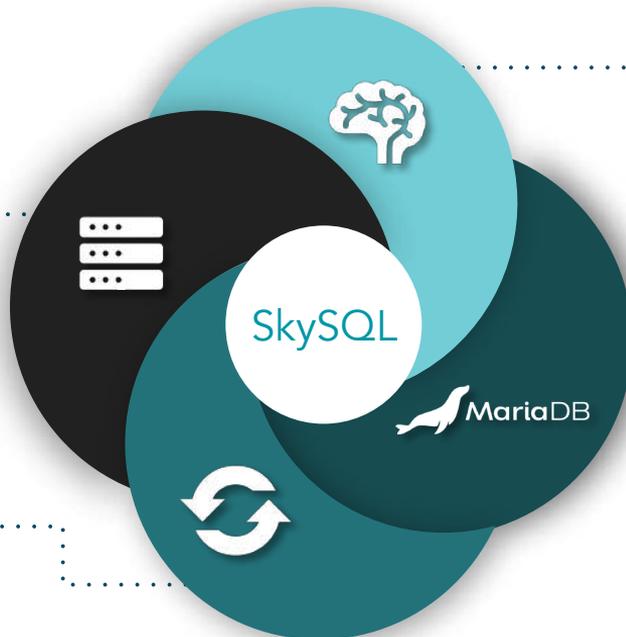
SkySQL FIRSTS

Community to Mission Critical

Expansive Topologies for all Applications Tiers.

Always On, Always Ready

The cloud offering the latest bug fixes and security updates, always.



Smart Transactions

The first DBaaS to support transactional (OLTP), analytical (OLAP) and hybrid (HTAP) workloads for Smart Transactions.

From the Source

A DBaaS architected, engineered, and supported by database architects, engineers, and extraordinary support.

L'analytique à haute valeur ajoutée grâce à MariaDB



MariaDB ColumnStore

What is ColumnStore

New needs, new storage engine, ColumnStore is here

Transactional

Many predefined queries
Point and range
Small subset of data
Many indexes

Ex: recent purchases

BI/reporting

Few predefined queries (mostly)
Range aggregate (mostly)
Medium subset of data
Few indexes

Ex: avg # of purchases (by month)

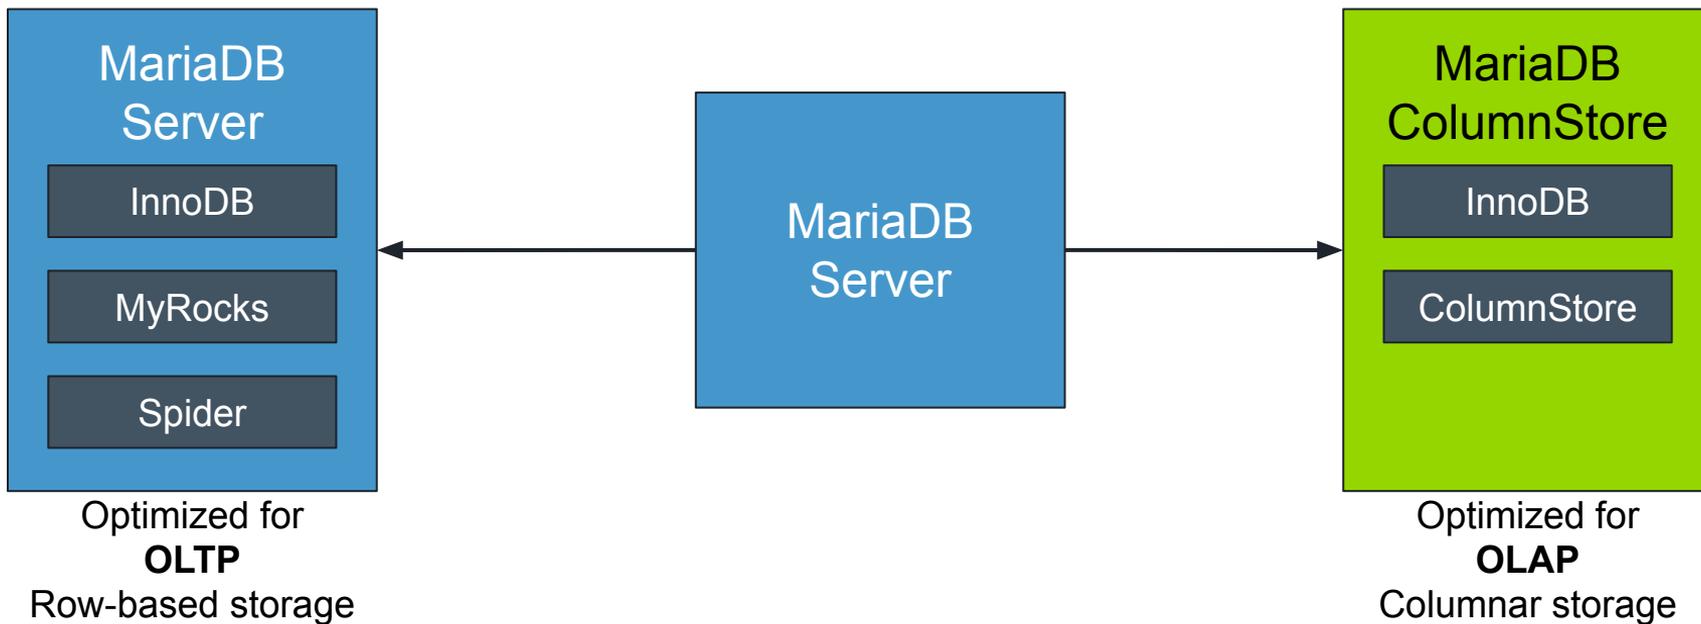
Analytical

Unforeseen queries (mostly)
Aggregate (mostly)
Large subset of data
No indexes

Ex: avg shopping cart life (all time)

MariaDB ColumnStore

What is ColumnStore



ColumnStore

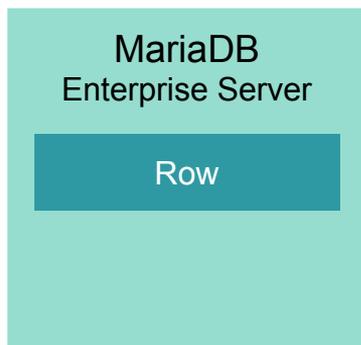


Les smart transactions



MariaDB Platform – use cases

OLTP



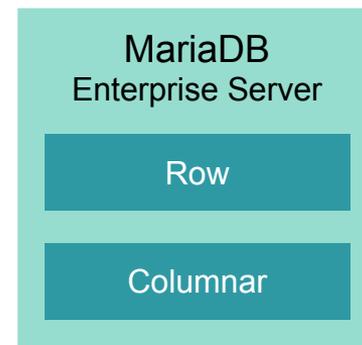
Oracle Database
Microsoft SQL Server
IBM Db2

OLAP



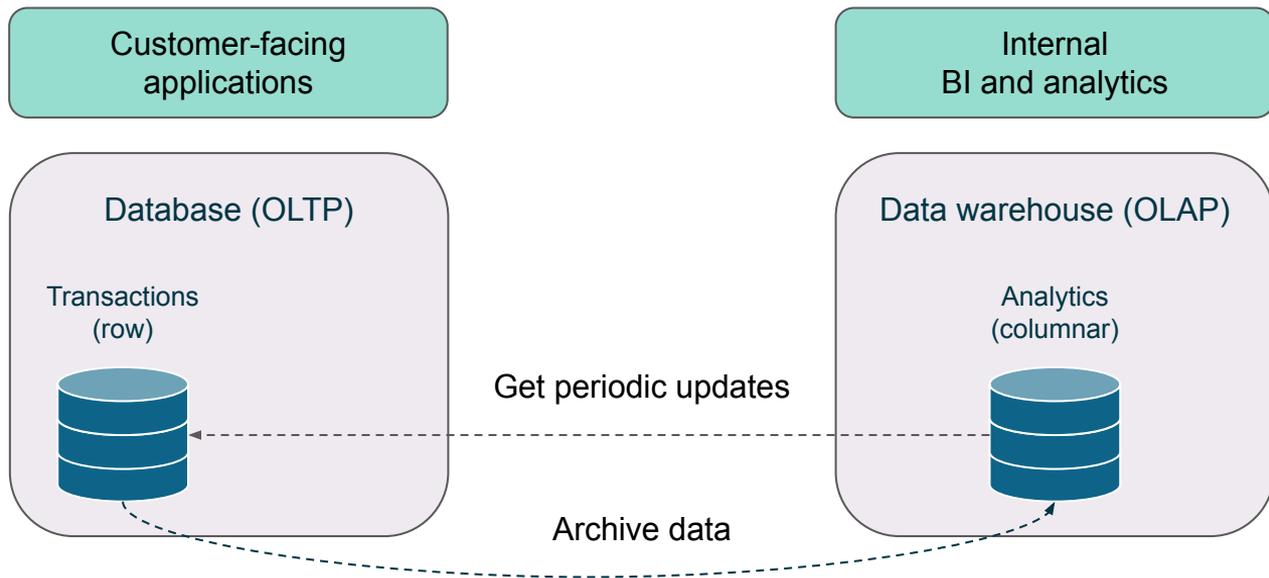
Greenplum
Vertica
Infobright DB

Smart OLTP

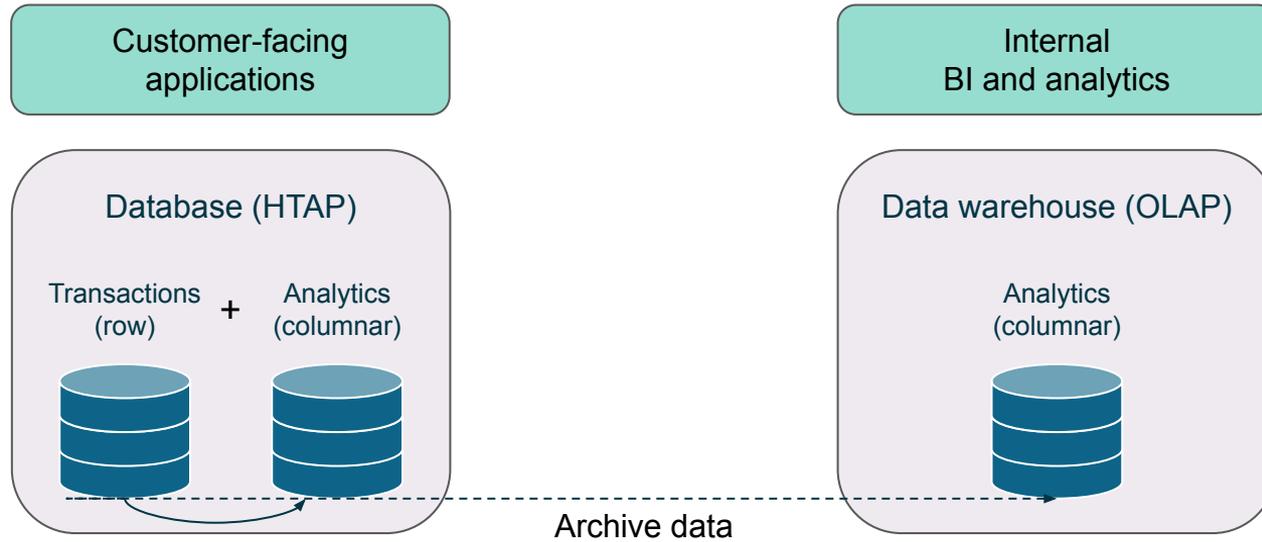


Oracle In-Memory
Microsoft ColumnStore indexes
IBM Db2 Shadow Tables
MemSQL

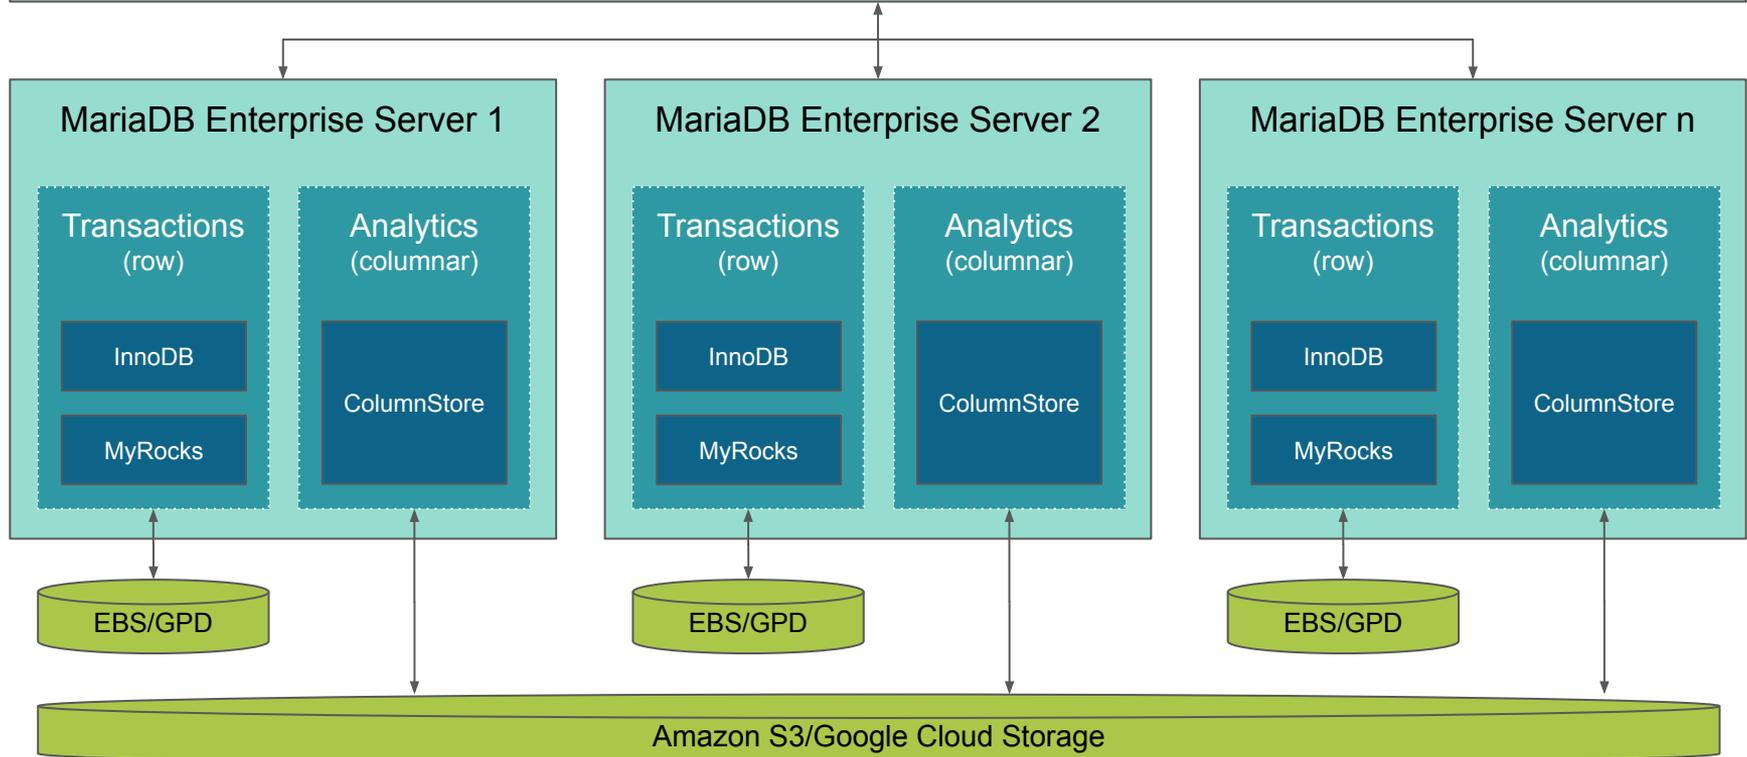
Cause



Solution



MariaDB MaxScale



Center for Information Management(CIM), Detroit, Michigan

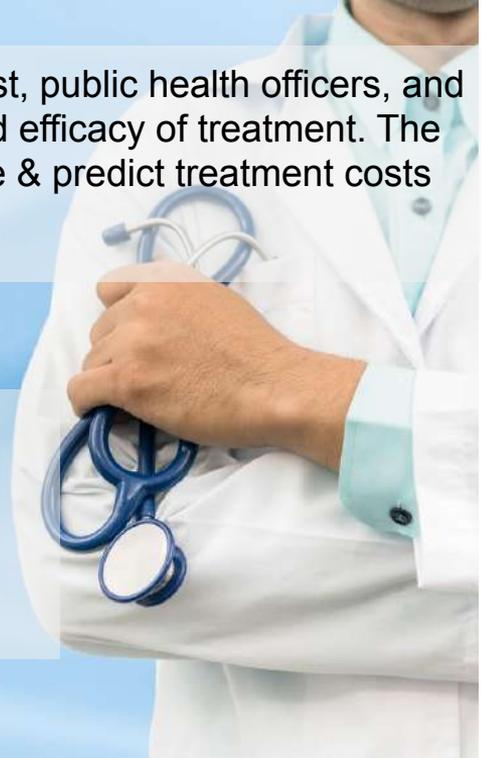
Challenges Faced:

- Take 2+ Days to Load Data
- Unsustainable Oracle Licensing Model
- Historically would need to create/manage indexes to get query performance

Health population data analysis to equip epidemiologist, public health officers, and insurance providers with evidence based outcome and efficacy of treatment. The insurance providers further use this insight to measure & predict treatment costs and population health over several years.

Benefits Realized:

- Data Loads now take minutes
- Can query 200+ columns with fast performance
- Easy administration - no indexes
- High Performance Data Visualization with Tableau



Single | Explore | Compare ▾

Settings Use basic settings

Display Cause | Risk
Etiology | Impairment

Metric Deaths ▾

Location Global ▾

Year 2015

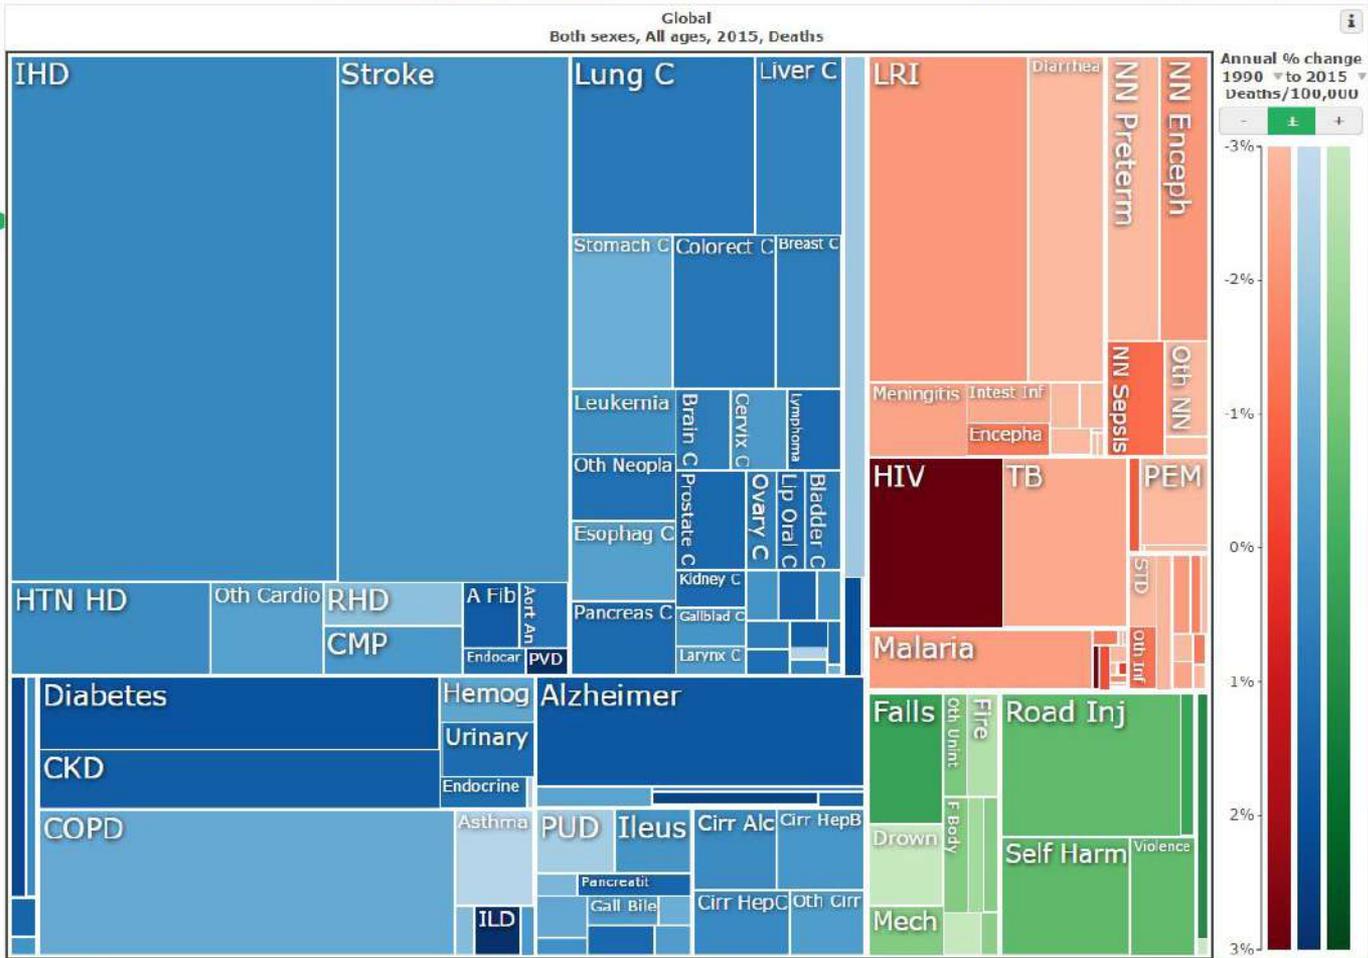
Age All ages ▾

Sex Male | Female | **Both**

Level 3

Value **Observed** | Expected

[Take tour ▶](#)





Tens of Millions of Users
Worldwide Trust Critical
Business Data to MariaDB

Q/A

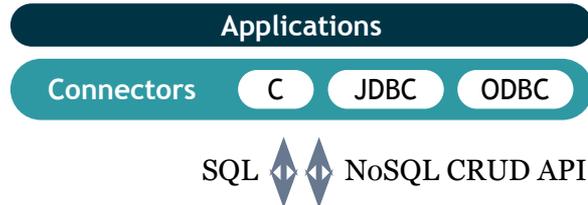


Annexe

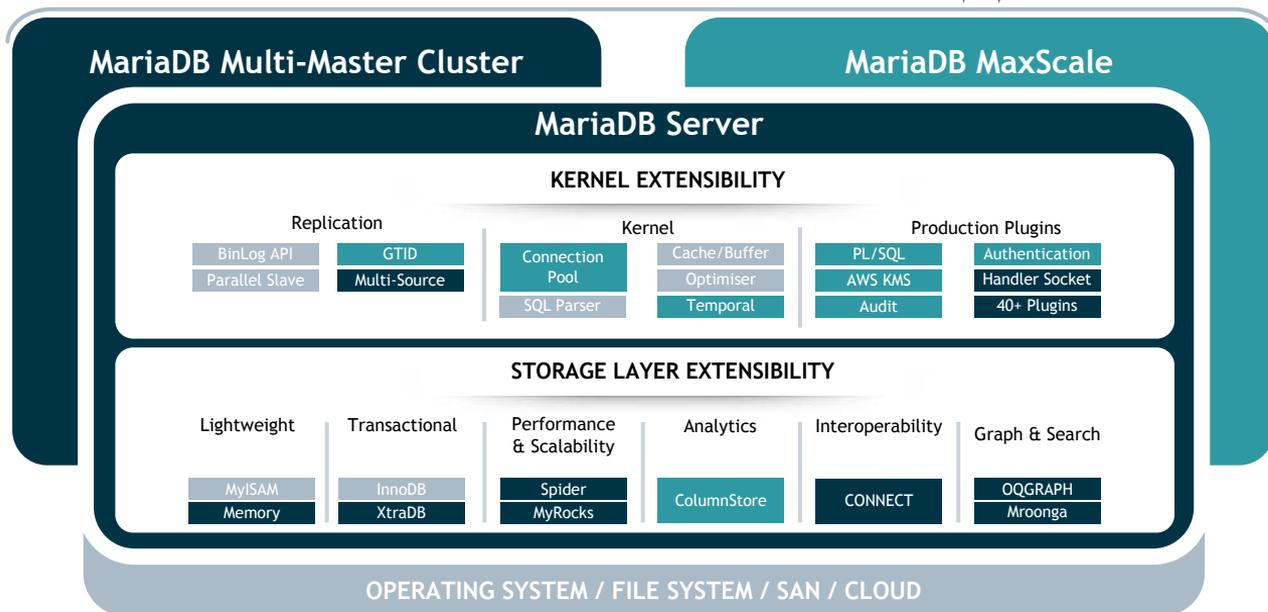
- Architecture extensible
 - MariaDB ColumnStore
 - Moteur de stockage Xpand
 - Architecture de SkySQL
 - Méthodologie de migration
 - MariaDB vs MySQL
 - MariaDB vs Oracle
-

Extensible Architecture

- Original Core
- MariaDB Engineering
- Community Contribution



MariaDB
Replicas
Supporting
Asynchronous,
Semi-Sync &
Synchronous
replication



Deep dive, what is MariaDB



MariaDB ColumnStore

What is ColumnStore

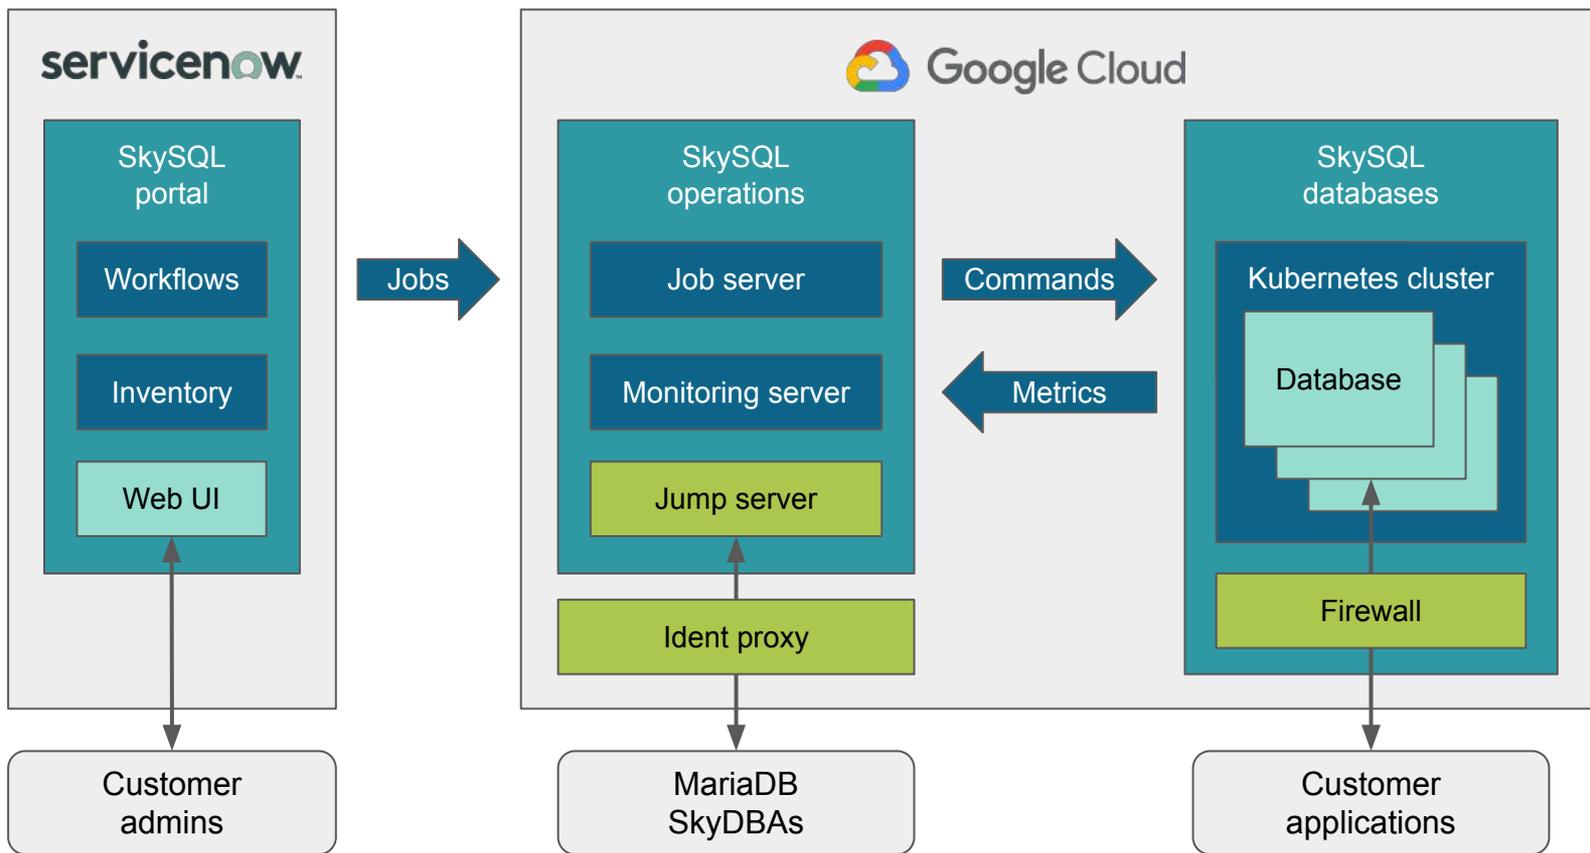
- Reduce disks operation to increase request speed
 - Using another algorithm to access only needed data
 - Compress data
- Administration operation simplified
 - No index
 - Automated failover and failback by design
- Request huge amount of data very quickly
 - Basis infrastructure on 5 servers (2 UM + 3 PM)
 - Store around 1TB per PM (storage server)
 - Request goes to UM in charge of request splitting

Xpand storage engine

The ladder for very high scaling

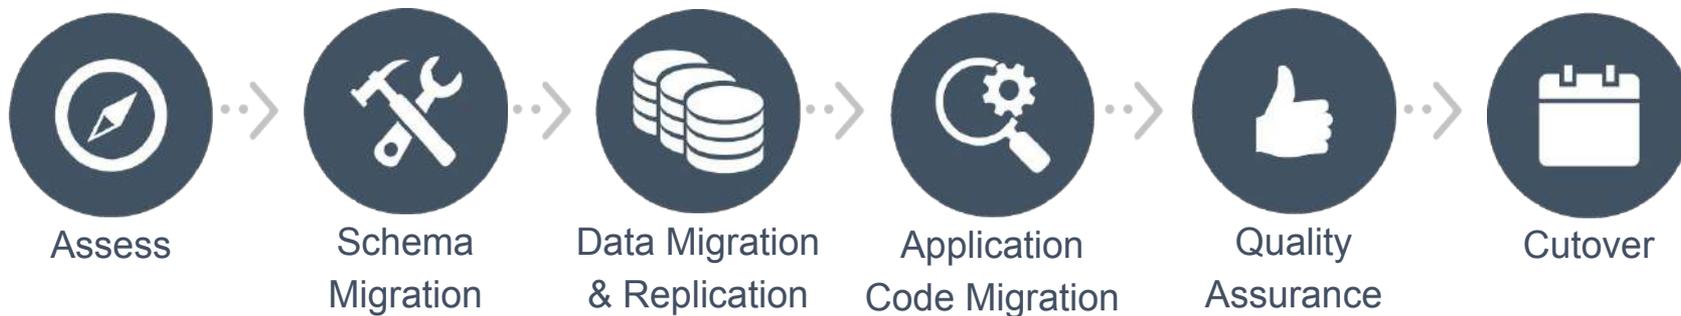
- Already used by Samsung cloud
- Can perform more than 5M requests per second (check OpenWork 2019 videos)
- Highly scalable
 - For reads AND writes requests
 - Designed for huge workloads
- Available on MaxScale
 - Allow HA
 - Allow dynamic schema routing

SkySQL (MariaDB cloud offer) : architecture

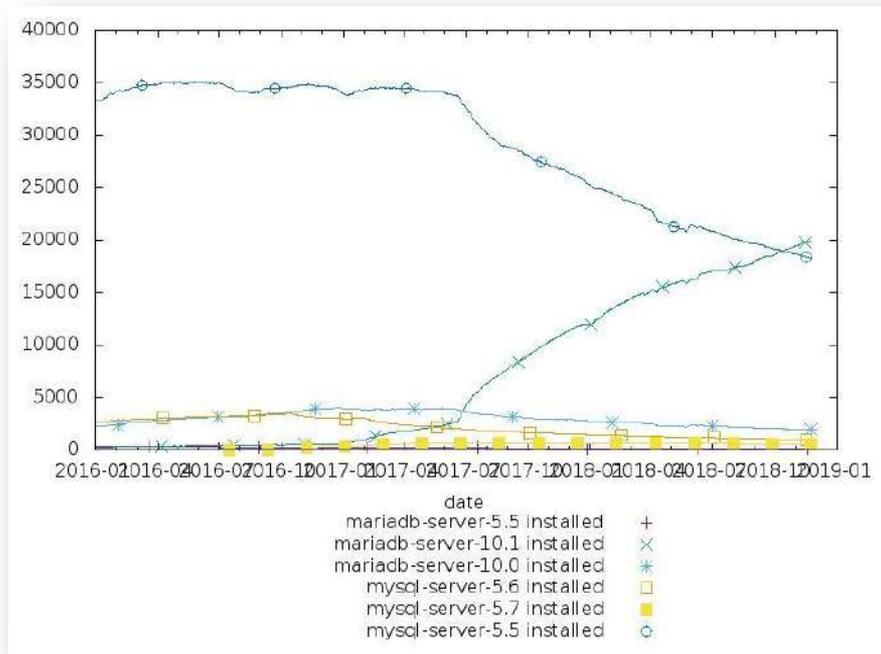


Migration methodology

Moving to MariaDB from other databases



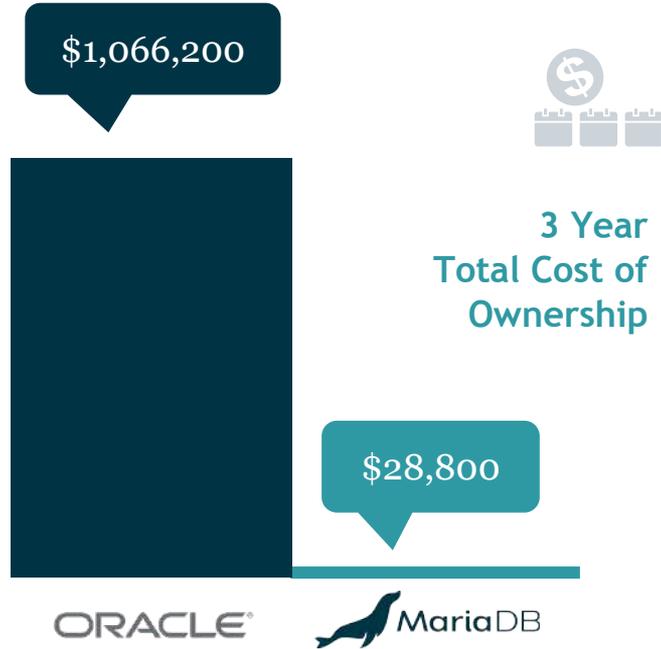
MariaDB adoption



Differentiators against MySQL

- Schema – instant changes (included in MariaDB)
- MaxScale (security, HA, caching, routing, replication)
- Data masking and obfuscation
- GPL/ BSL licence (OpenSource)
- Oracle compatibility
 - Packages
 - PL/SQL
- Multi engine (evolving quickly Xpand)
- Partnership with Codership (Galera)

MariaDB vs Oracle



For one node (2 sockets, 12 core each)

- Oracle costs **40x** more
- Organizations can save more than **\$1 million**

Source: "State of the Open-Source DBMS Market, 2018", **Gartner**, Feb 2018 and MariaDB, July 2018



MariaDB // Oracle TCO

	MariaDB	Oracle
Database	\$9,600	\$47,500
Clustering	\$2,400	\$23,000
Partitioning	Included	\$11,500
Compression	Included	\$11,500
Security	Included	\$15,000
Firewall	Included	\$6,000
TOTAL	\$12,000 (Server)	\$114,500 (Core)

MariaDB	Oracle
Database Servers 3	Database Server 3
	Processors/Server 2
	Cores/Processor 16
	Total Cores 96
	Core Factor 0.5
	Adjusted Cores 48
TOTAL \$36,000	TOTAL \$5,496,000

Add values with MariaDB against Oracle

Saving Millions with MariaDB

The following costs of each database are calculated, using published list prices, based on a minimal configuration capable of meeting standard enterprise requirements for a database.

After three years running on three on-premises servers, each with two, 16-core processors:

- The total cost of Oracle is 84x higher than MariaDB Platform
- The annual cost of Oracle is 33x higher than MariaDB Platform
- Organizations can save over \$9 million after three years by choosing MariaDB Platform
- Organizations can save \$1.1 million annually by replacing Oracle

After three years running on three cloud instances, each with 16 cores (32 hyperthreads):

- The total cost of Oracle is 72x to 151x higher than MariaDB Platform
- The annual cost of Oracle is 60x to 72x higher than MariaDB Platform
- Organizations can save \$7.6 to \$16.3 million after three years by choosing MariaDB Platform
- Organizations can save \$2.1 to \$2.5 million annually by replacing Oracle



Festival Online de la Data



PROCHAINE DATE

Jeudi 9 juillet de 15h à 15h30

Sensibilisation à la DataVisualisation

Franck Nguyen, Synaltic