



OPENWORKS

BE UNSTOPPABLE



OPENWORKS

PLUG INTO ANALYTICS: CONNECTING COLUMNSTORE TO SOURCE DATABASES WITH SPIDER ENGINE

KATHRYN SIZEMORE, SR SOLUTIONS ARCHITECT, MARIADB

OVERVIEW

- Scaling Architecture
 - Transactional Source: MariaDB Server & Xpand
 - Analytics: Columnstore
 - Bridging the Gap: Spider
- Plugging in Spider
- Transactional to Analytics Demonstration

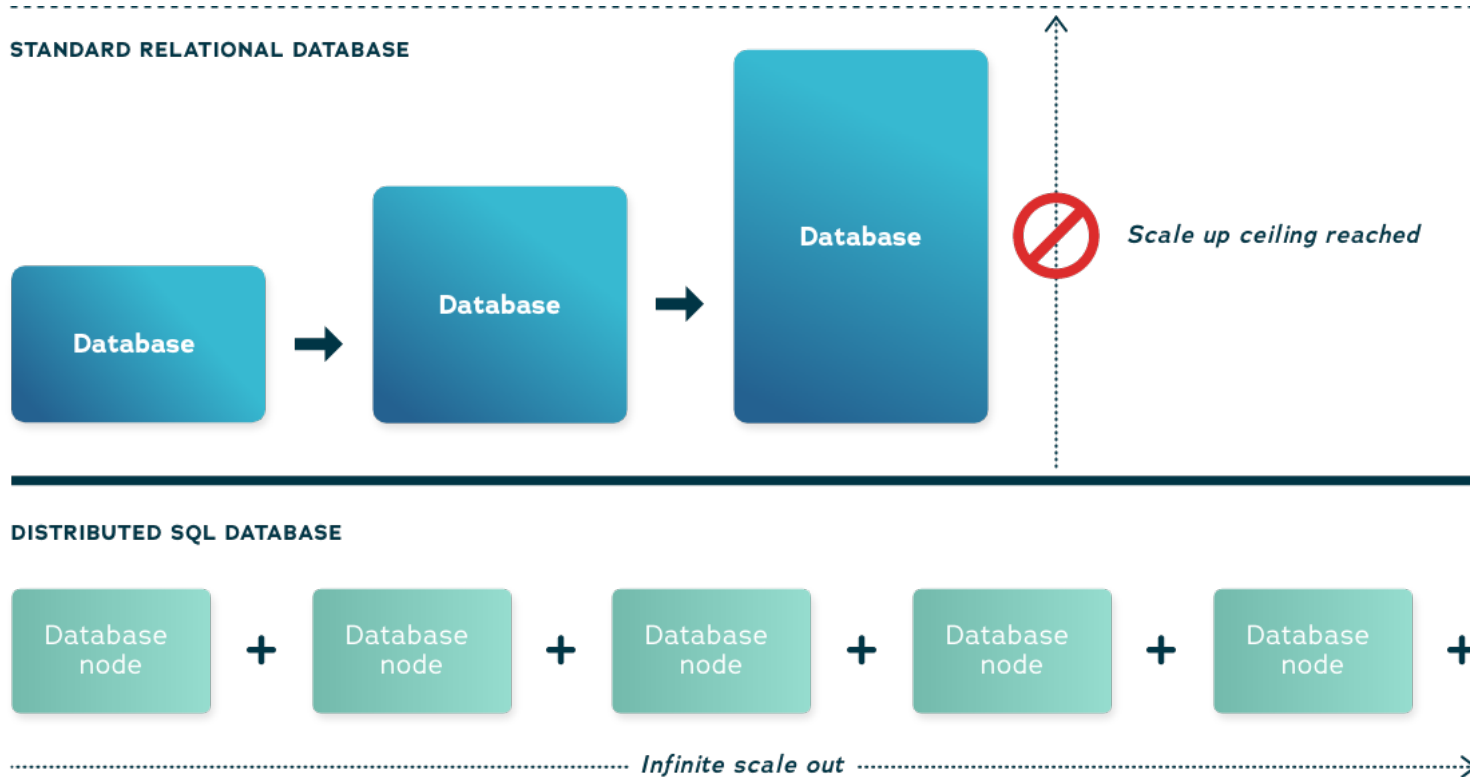
FEEDING YOUR ANALYTICS VIA SPIDER ENGINE

Source Databases

- MongoDB
- IBM DB2
- Amazon Redshift
- Snowflake
- Google BigQuery
- Databricks
- MariaDB Server
- MariaDB Xpand
- MySQL
- Oracle
- Microsoft SQL Server
- PostgreSQL

Spider's ODBC Connects to your Existing Database

THE SCALE CHALLENGE



MARIADB FOR ANY WORKLOAD

Transactional

TRANSACTIONAL APPLICATIONS

MARIADB MAXSCALE

Enterprise Server

InnoDB

Cluster

Run the Business

Analytical

REPORTING & ANALYTICS

MARIADB MAXSCALE

Enterprise Server

ColumnStore

Observe the Business

Distributed

DISTRIBUTED SQL (PERFORMANCE)

MARIADB MAXSCALE

Xpand

Excellent for web-scale applications, OLTP, and real-time analytics

Scale the Business

MARIADB XPAND DISTRIBUTED SQL

The image shows a screenshot of the MariaDB Xpand website. The top navigation bar includes links for Knowledge Base, Contact Us, Log In, and Register. The main heading is "Distributed SQL Powered by MariaDB Xpand". Below this, there is a section titled "MARIADB XPAND" with a brief description: "MariaDB Xpand is a distributed SQL database for businesses who need to run web, mobile and IoT applications at a scale beyond that which standard relational databases are capable of reaching. These tools with a singular, underlying focus on scaling SQL to unimagined levels without sacrificing the relational data model and ACID transactions. It's a database for every business in the digital era, from digitally native disruptors to industry titans undergoing digital transformation." A small video player is visible below the text. To the right of the website screenshot is the MariaDB logo (a seal) and the text "MariaDB". Below the website screenshot is a diagram showing an "Application" box connected to three database icons, each labeled "RW". At the bottom left of the screenshot is a "2022 AWARDS TECHNOLOGY OF THE YEAR" badge, and at the bottom right is the "InfoWorld" logo.

- MariaDB's Distributed SQL Database
- A single logical database distributed across multiple nodes
- Scale anywhere, anytime
- Use any MySQL or MariaDB ecosystems of tools
- Never worry about availability

Xpand. Purpose built for scale.

XPAND USE CASES

Mission Critical, High Write Load, Strict Consistency, Scalability



Global Media/IoT

Global Devices



Gaming Platforms

Gambling & Social
Gaming



Financial Services

FinTech & Investment



Hospitality Industry

Hotels & Amusement
Parks

MARIADB COLUMNSTORE

High-performance Engine Supporting a Variety of Analytical Use Cases



Incredible Analytic Query Speed

Faster than other column-oriented analytics databases



Native MariaDB Plugin Engine

Easier to use with MariaDB Language and Installation



Blazing Fast Data Ingestion

Bypass the SQL layer for incredible speed with cimpport

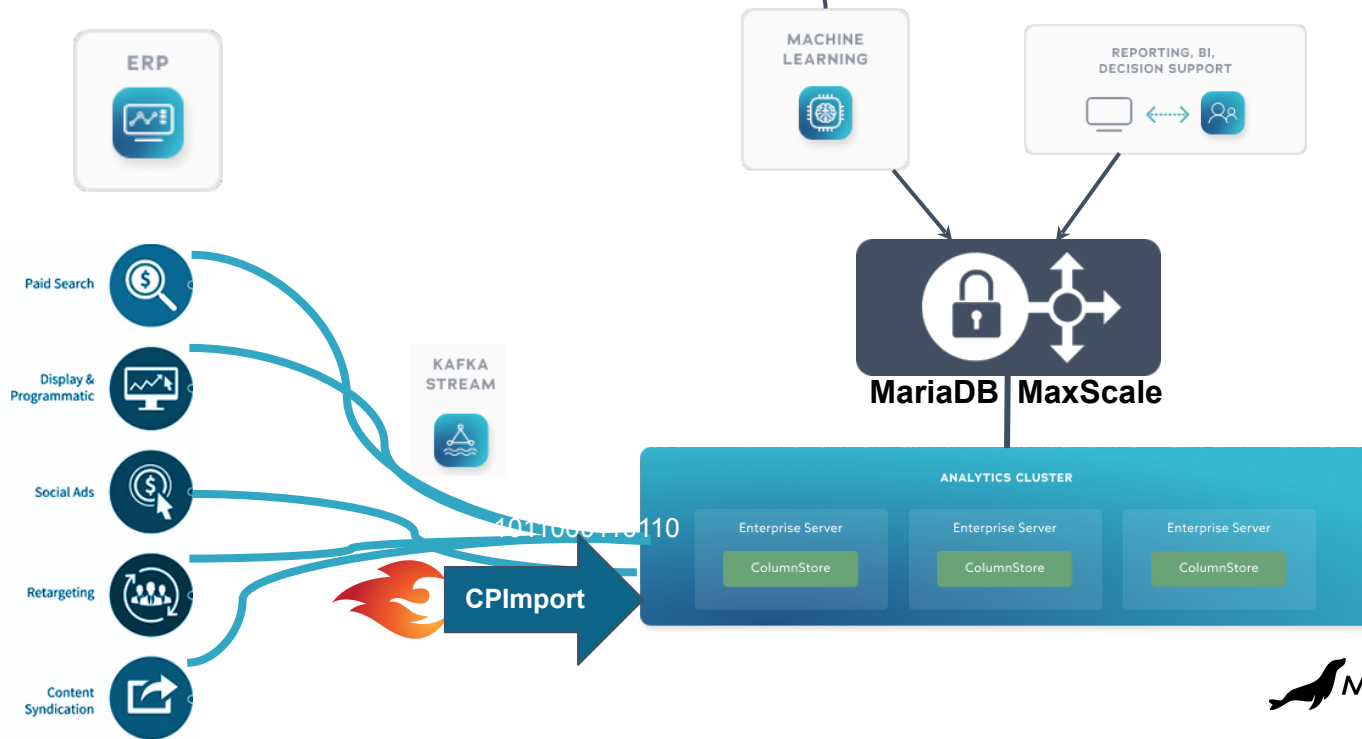


Hardware & Cost Efficient

Store massive amounts of data with low cost object storage

COLUMNSTORE USE CASES

Observe Your Business
- Intelligently Adapt



COLUMNSTORE USE CASES

Ingest Data and Dynamically Analyze

Customer Dashboards

WEBAPPS AND MOBILEAPPS

REPORTING, BI, DECISION SUPPORT

MACHINE LEARNING



Customer Data

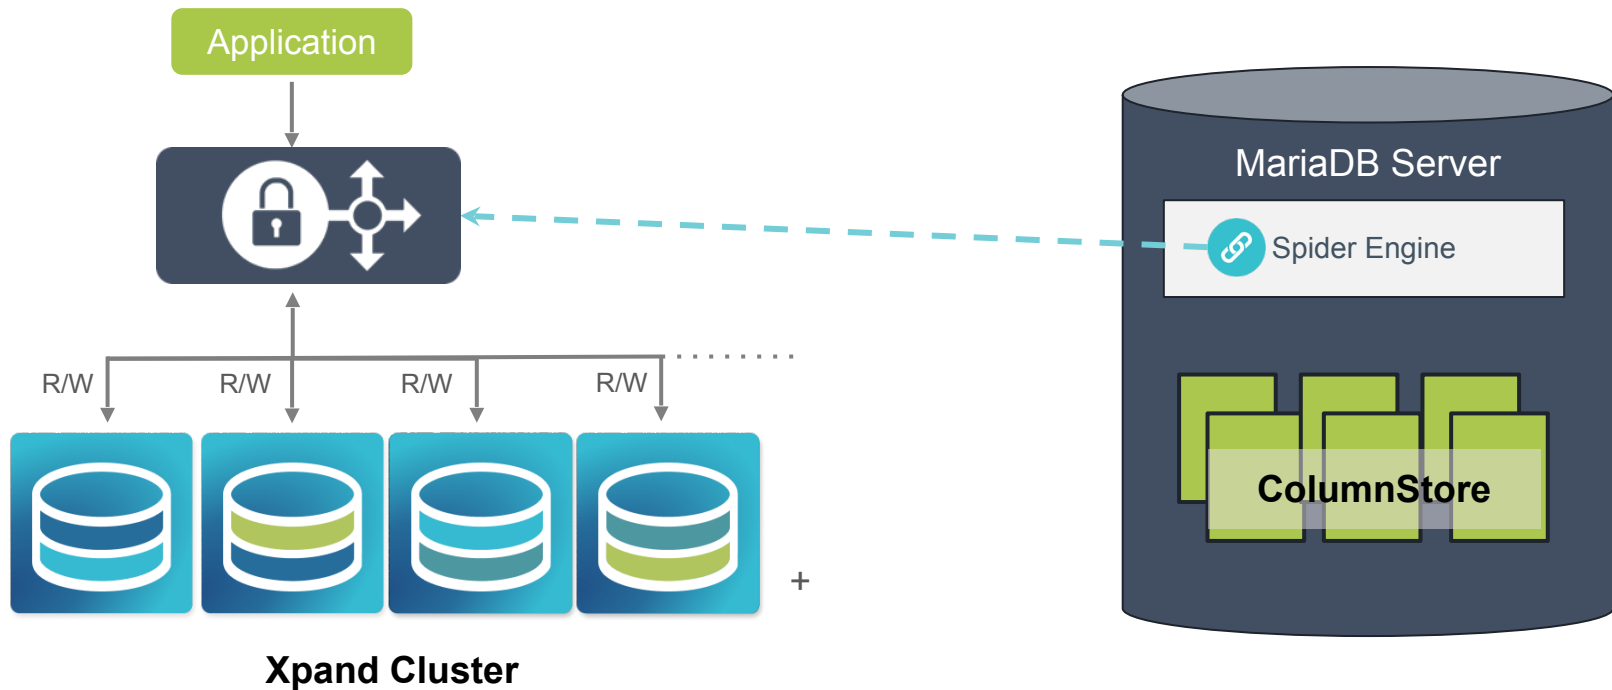


BRIDGE THE GAP WITH SPIDER

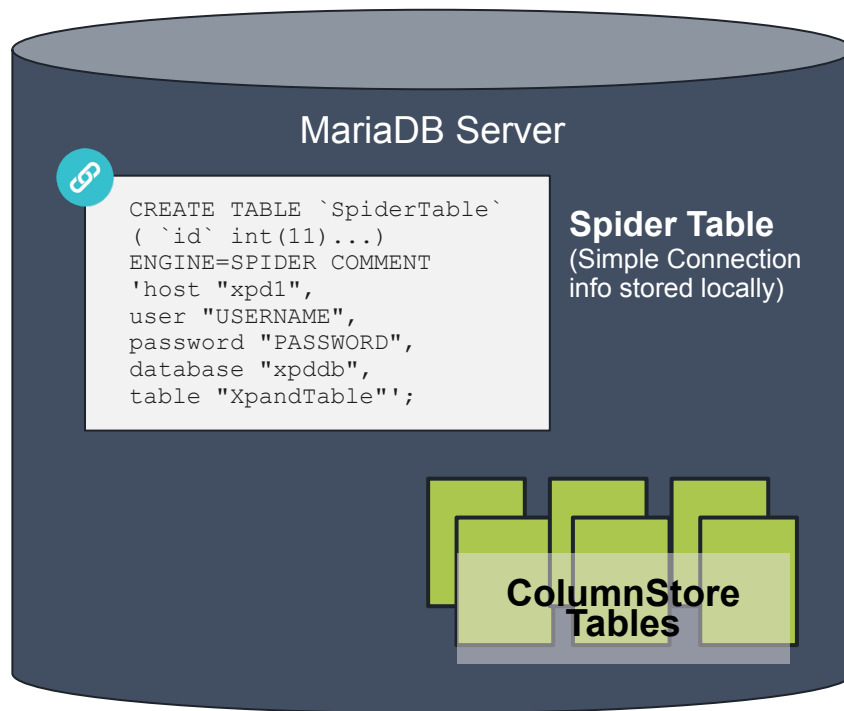

PLUGIN WITH SPIDER ENGINE

- Create a Spider Table with ODBC connection to Source Database
- Queries and Applications see this table as local
- Use Spider Table
 - Create Fly Queries - Joining to source
 - Create a Stored Procedure to Select Data From source
 - Use to ETL data to Target Columnstore Tables

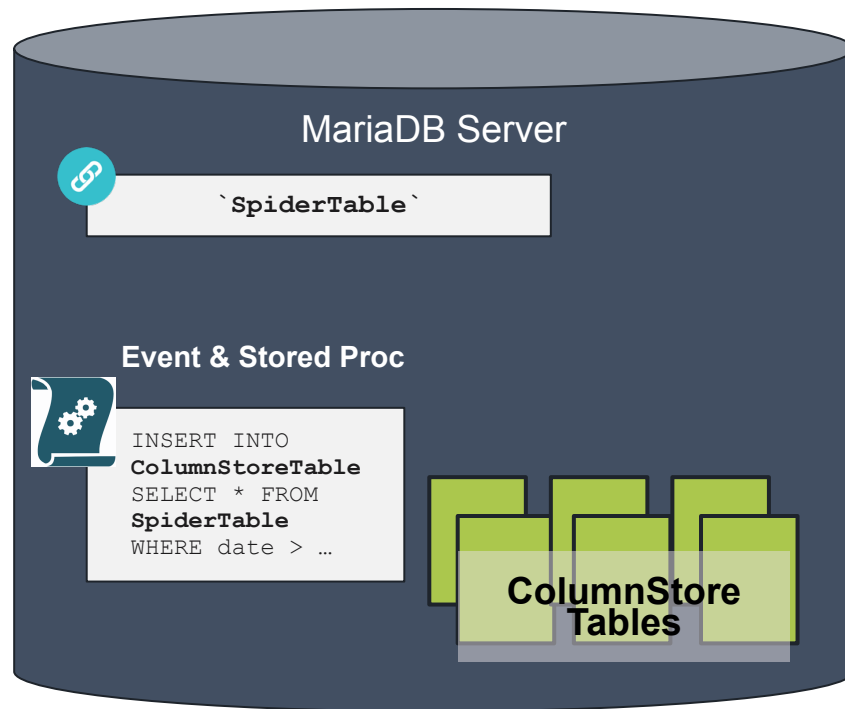
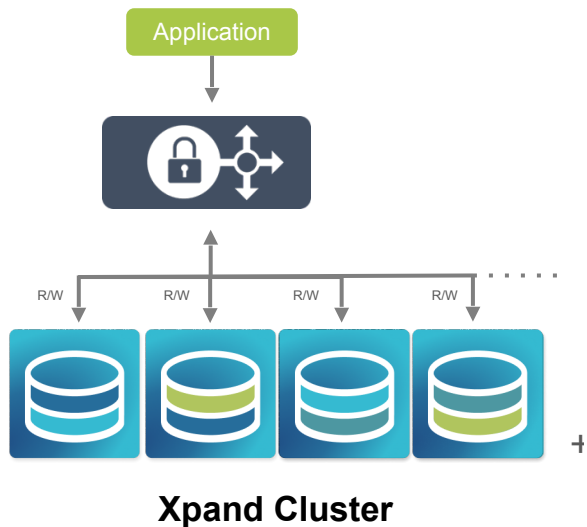
TRANSACTIONAL TO ANALYTICS ARCHITECTURE



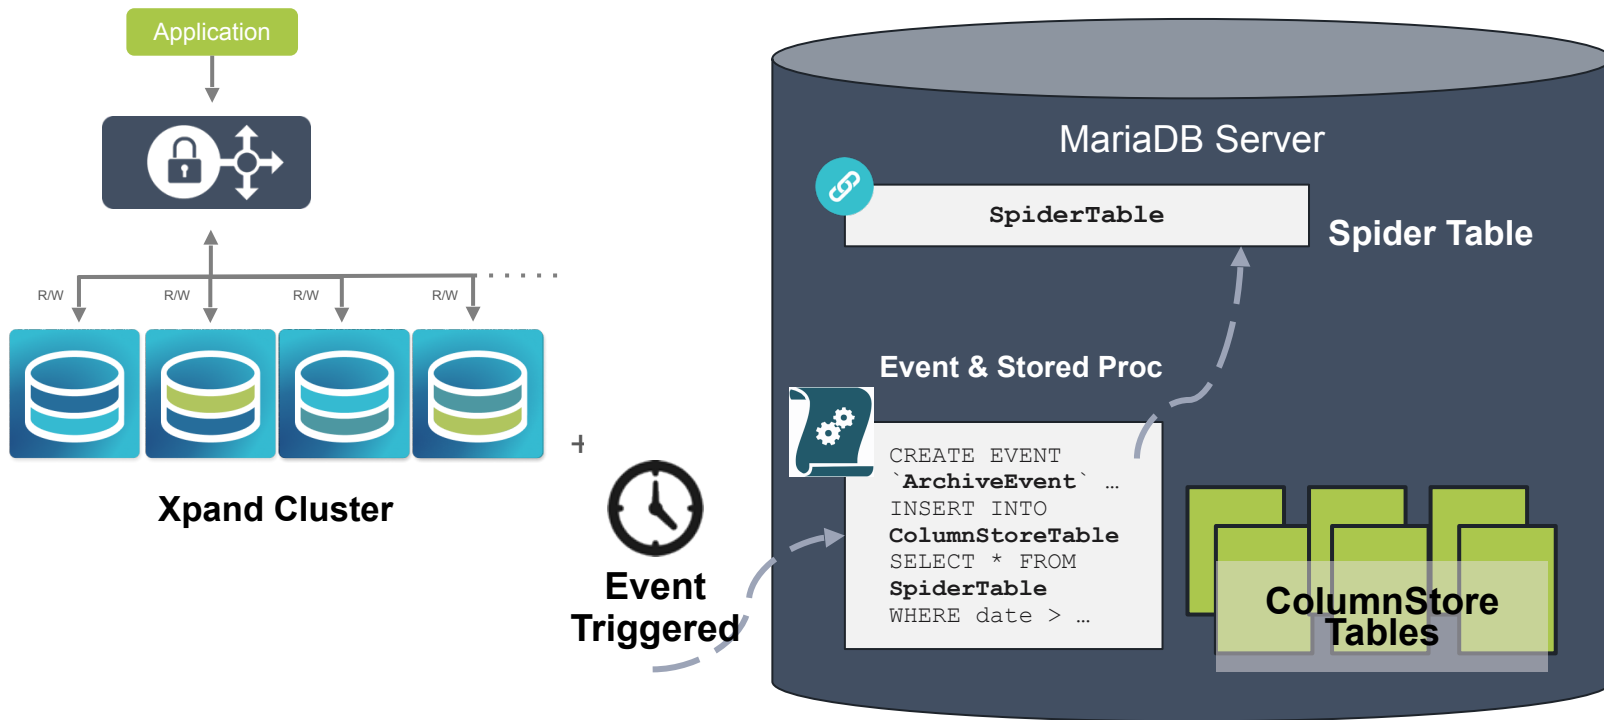
TRANSACTIONAL TO ANALYTICS ARCHITECTURE



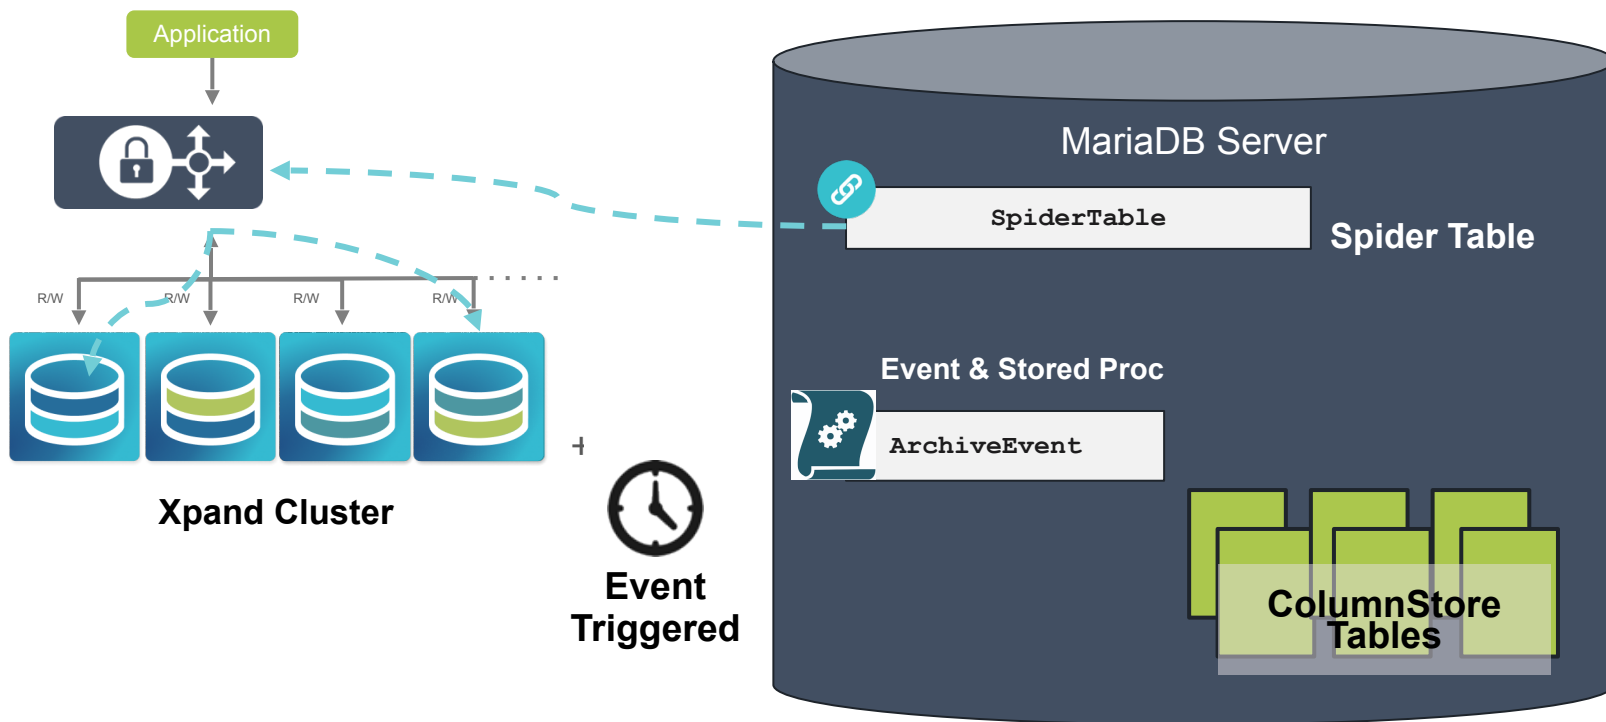
TRANSACTIONAL TO ANALYTICS ARCHITECTURE



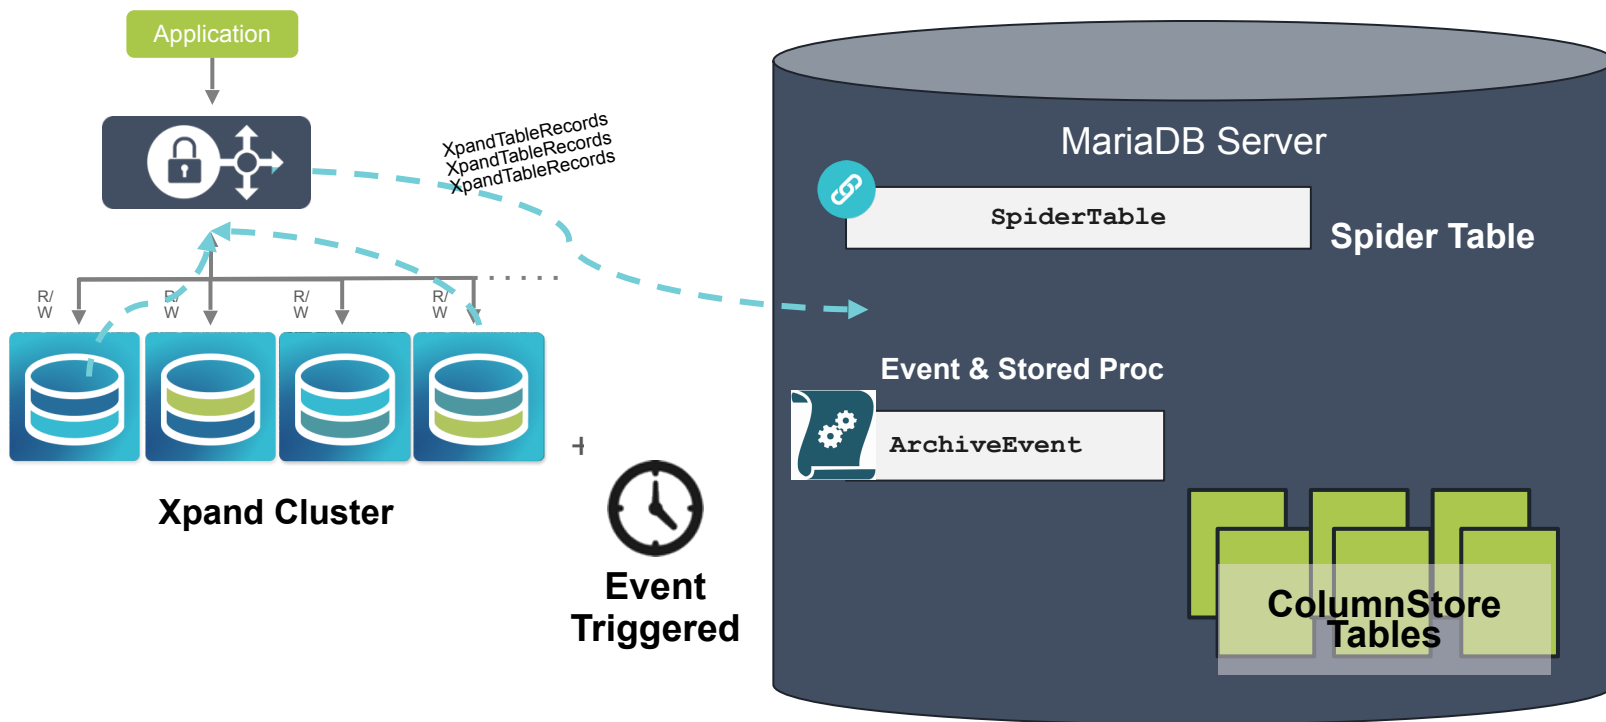
TRANSACTIONAL TO ANALYTICS ARCHITECTURE



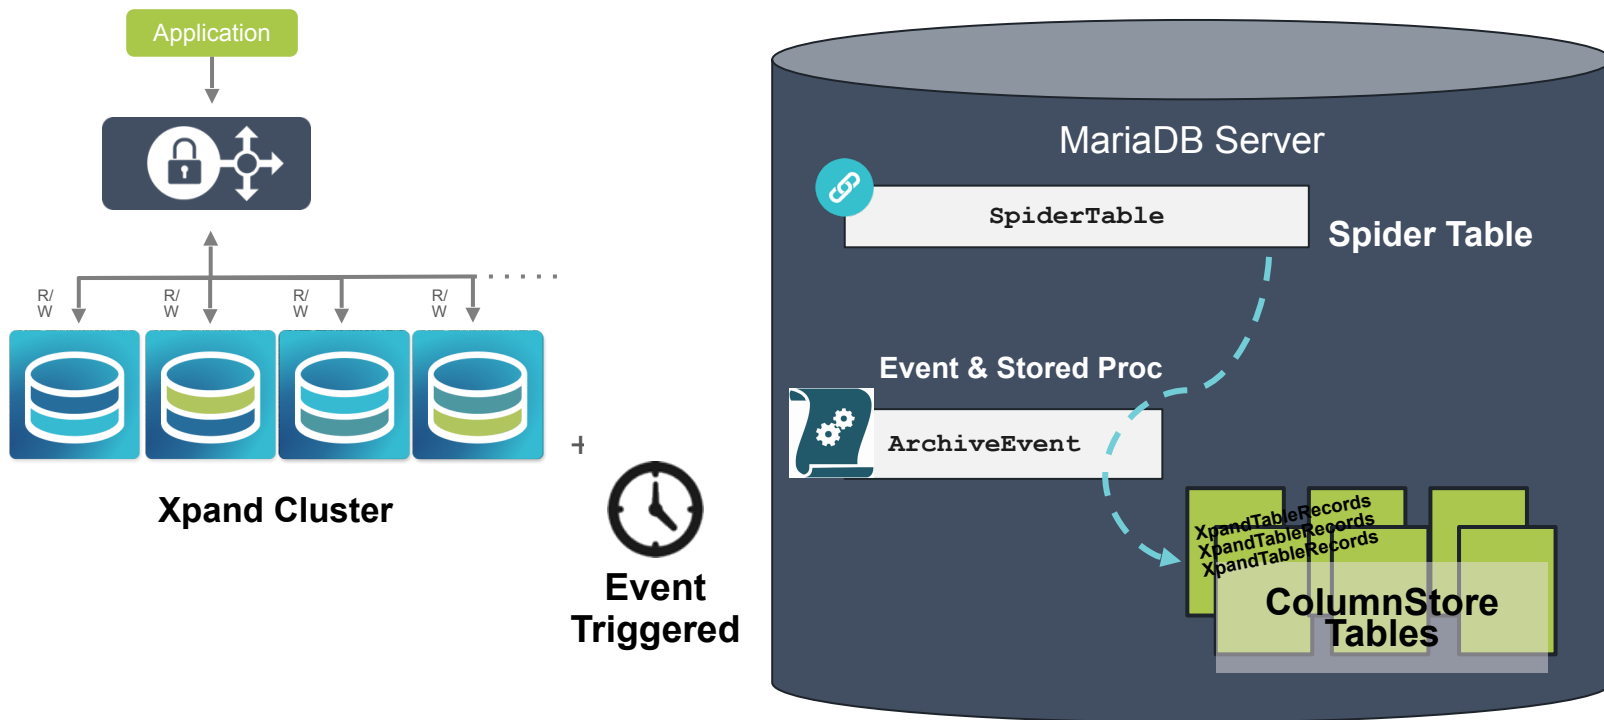
TRANSACTIONAL TO ANALYTICS ARCHITECTURE



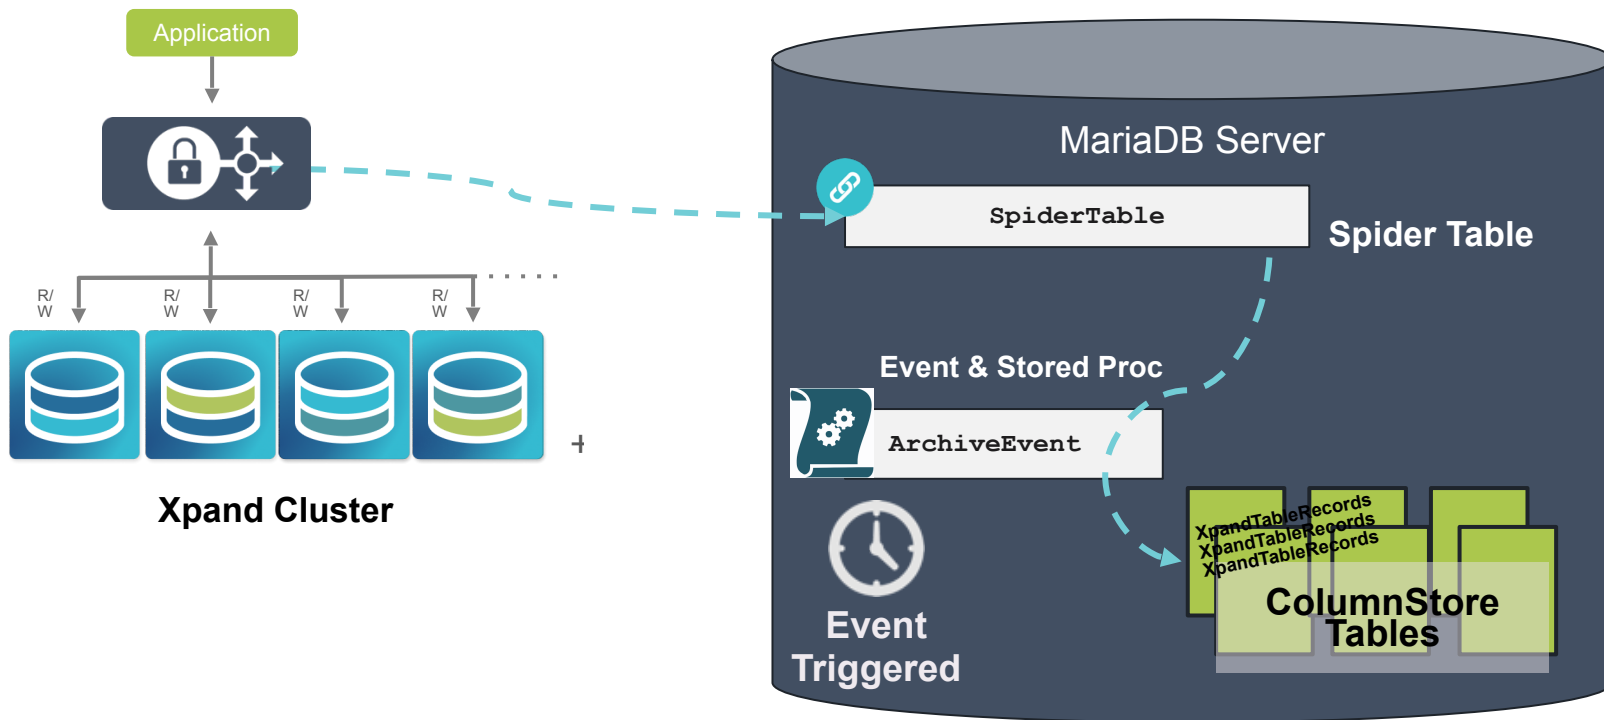
TRANSACTIONAL TO ANALYTICS ARCHITECTURE



TRANSACTIONAL TO ANALYTICS ARCHITECTURE




TRANSACTIONAL TO ANALYTICS ARCHITECTURE




DEMO

DEMO ARCHITECTURE

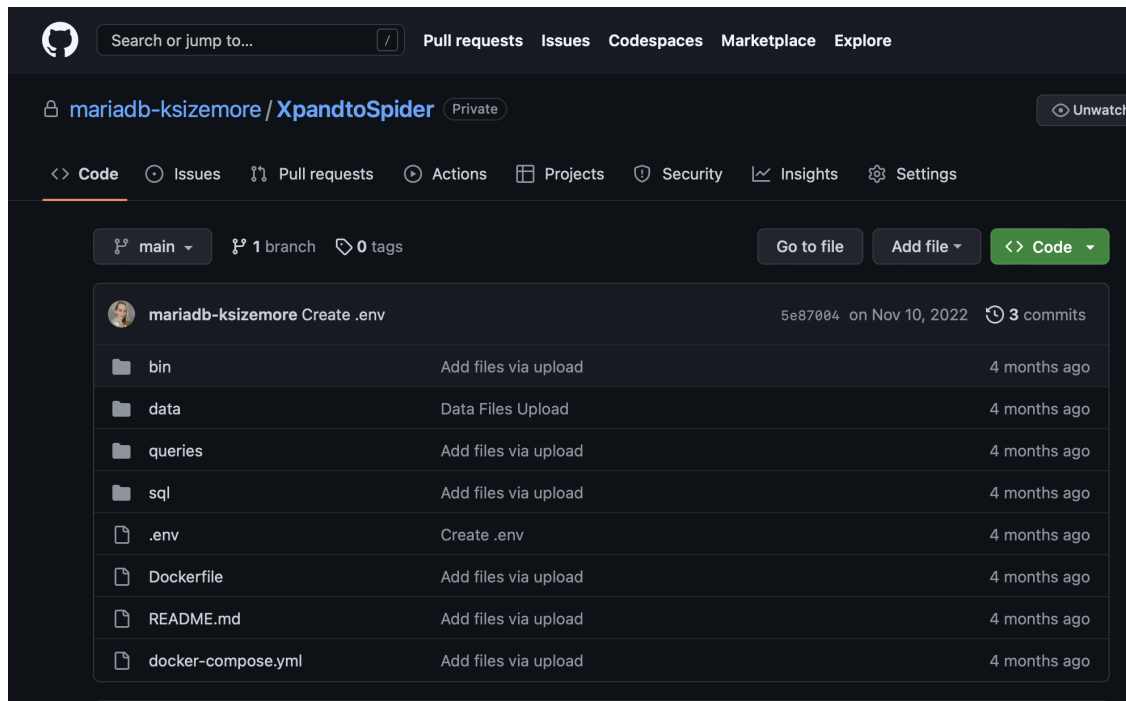


Xpand Docker



**MariaDB Docker
+ Columnstore
+ Spider Engine**

GITHUB DOCKER DEMO



The screenshot shows the GitHub interface for the repository `mariadb-ksizemore / XpandtoSpider`. The repository is private and has 3 commits. The main branch is selected. The repository structure is as follows:

Item	Action	Time
bin	Add files via upload	4 months ago
data	Data Files Upload	4 months ago
queries	Add files via upload	4 months ago
sql	Add files via upload	4 months ago
.env	Create .env	4 months ago
Dockerfile	Add files via upload	4 months ago
README.md	Add files via upload	4 months ago
docker-compose.yml	Add files via upload	4 months ago

<https://github.com/mariadb-ksizemore/XpandtoSpider>

SPIDER TABLE SETUP

Example of Connection Credentials Written into Create Table

```
CREATE TABLE airlines_spdr
(
    iata_code VARCHAR(2),
    airline VARCHAR(30)
) ENGINE=SPIDER COMMENT
'host "xpd1",
database "bts",
table "airlines",
user "importuser",
password "importuserpasswd",
port "3306";
```

SPIDER TABLE SETUP

Example of Connection Credentials Saved in SERVER

```
CREATE SERVER xpand1
FOREIGN DATA WRAPPER mariadb
OPTIONS (
  HOST "xpd1",
  PORT 3306,
  USER "importuser",
  PASSWORD "importuserpasswd",
  DATABASE "bts"
);
```

```
CREATE TABLE airports_spdr
( iata_code VARCHAR(3),
  airport VARCHAR(80),
  city VARCHAR(30),
  state VARCHAR(2),
  country VARCHAR(30),
  latitude DECIMAL(11, 4),
  longitude DECIMAL(11, 4)
) ENGINE=SPIDER COMMENT
'server "xpand1",
database "bts",
table "airports";
```

NEXT STEPS

Check out these sources to learn more about MariaDB

- Attend this session next:
Xpand Your Data Possibilities & Pipeline with MariaDB Xpand & Qlik Data Integration
- Watch this session On-Demand tomorrow:
Using MariaDB ColumnStore with Power BI for Visualization and Reporting
- Check out this **Demo**: <https://github.com/mariadb-ksizemore/XpandtoSpider>
- Try Columnstore on SkySQL with **\$500 free credit**

QUESTIONS?



 kathryn.sizemore@mariadb.com

 <https://www.linkedin.com/in/kathryn-sizemore>



MariaDB

OPENWORKS

THANK YOU