



**OPENWORKS**

**BE UNSTOPPABLE**



# OPENWORKS

## USING MARIADB COLUMNSTORE WITH MICROSOFT POWER BI FOR VISUALIZATION AND REPORTING

TODD STOFFEL, DIRECTOR PRODUCT MANAGEMENT, MARIADB

# Agenda

1. About ColumnStore
2. Power BI Demo
3. Q&A



# Data Management Technology Challenges for Business Analysts and Other Power Users

- Reliance on older, siloed databases, warehouses, and lakes
- Excel spreadsheet security blanket
- Stale and incomplete datasets
- One-size-fits-all, complex data preparation tools

# Data Management Resource Challenges for Business Analysts and Other Power Users

- IT Approval and Support needed for new projects
- Data Engineers needed for data preparation
- DBAs needed for initial and ongoing data use
- Cloud operations needed to start up projects and avoid runaway costs

## High-performance Engine Supporting a Variety of Analytical Use Cases



### **Better Performance**

Power of SQL  
Freedom of Open Source  
Big Data Analytics



### **Easier Enterprise Analytics**

Single Interface  
OLTP and OLAP  
Interactive, Ad Hoc Analytics  
Easy to Manage and Scale

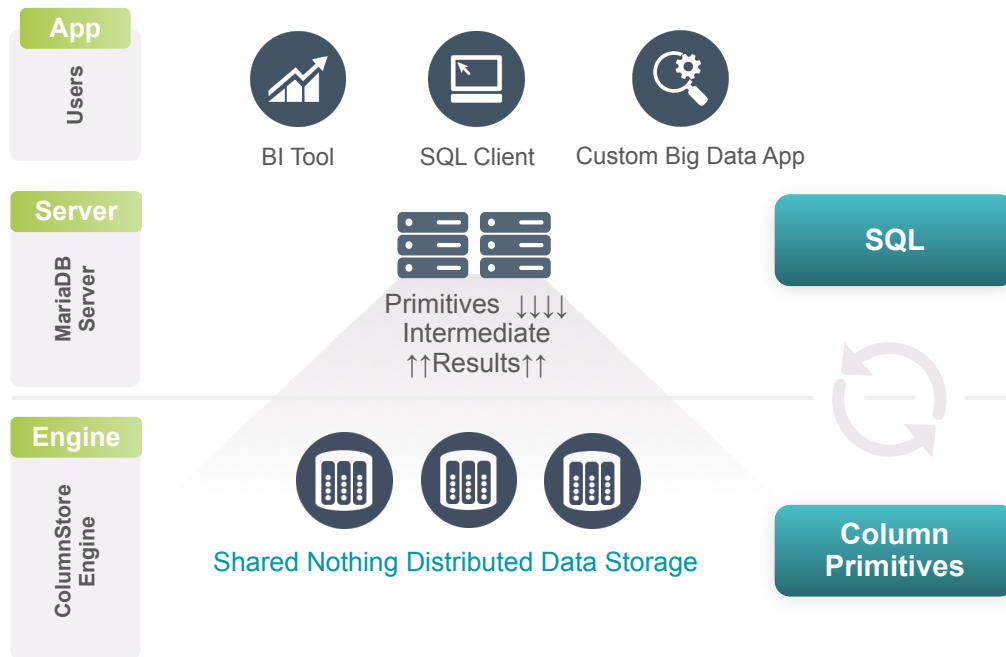


### **Faster, More Efficient Queries**

Scales out Massively Parallel  
Processing (MPP) for  
Distributed Environments

# Cloud-Ready Architecture

Massively Parallel, Shared Nothing



# Row-Oriented v. Column-Oriented

```
SELECT Fname FROM TABLE1 WHERE State = 'NY';
```

## ROW-ORIENTED

- Rows stored sequentially in a file
- Scans through every record row by row

ID	FNAME	LNAME	STATE	ZIP	PHONE	AGE	GENDER
1	Bugs	Bunny	NY	11217	(718) 938-3235	34	M
2	Yosemite	Sam	CA	95389	(209) 375-6572	52	M
3	Daffy	Duck	NY	10013	(212) 227-1810	35	M
4	Elmer	Fudd	ME	04578	(207) 882-7323	43	M
5	Witch	Hazel	MA	01970	(978) 744-0991	57	F

## COLUMN-ORIENTED

- Each column is stored in a separate file
- Scans the only relevant column

ID	FNAME	LNAME	STATE	ZIP	PHONE	AGE	GENDER
1	Bugs	Bunny	NY	11217	(718) 938-3235	34	M
2	Yosemite	Sam	CA	95389	(209) 375-6572	52	M
3	Daffy	Duck	NY	10013	(212) 227-1810	35	M
4	Elmer	Fudd	ME	04578	(207) 882-7323	43	M
5	Witch	Hazel	MA	01970	(978) 744-0991	57	F



# Extent Elimination

```
SELECT Item, sum(Quantity) FROM Orders
WHERE ShipDate
BETWEEN '2016-01-01' AND '2016-01-31'
GROUP BY Item;
```

**EXTENT 1:**  
ShipDate: 2016-01-12 - 2016-03-05



**Horizontal Partition:**  
8 Million Rows  
Extent 1

**EXTENT 2:**  
ShipDate: 2016-03-05 - 2016-09-23



**Horizontal Partition:**  
8 Million Rows  
Extent 2

**EXTENT 3:**  
ShipDate: 2016-09-24 - 2017-01-06



**Horizontal Partition:**  
8 Million Rows  
Extent 2

ID	ORDERID	LINE	ITEM	QUANTIT Y	PRICE	SUPPLIER	SHIPDATE	SHIPMOD E
1	1	1	Laptop	5	1000	Dell	2016-01-12	G
2	1	2	Monitor	5	200	LG	2016-01-13	G
3	2	1	Mouse	1	20	Logitech	2016-02-05	M
4	3	1	Laptop	3	1600	Apple	2016-01-31	P
...	...	...	...	...	...	...	...	...
8M							2016-03-05	
8M+1							2016-03-05	
...	...	...	...	...	...	...	...	...
16M							2016-09-23	
16M+1							2016-09-24	
...	...	...	...	...	...	...	...	...
24M							2017-01-06	

**Eliminated Extent**

**Eliminated Extent**

# Latest Enhancements

- Cloud Infrastructure Improvements:  
Use AWS Graviton and ARM A1 chips
- Performance Improvements
  - Insert Cache for write-heavy workloads accelerates up to 4700X faster
  - Fast Deletes acceleration up to 6X faster
  - `cpimport` function for bulk high-speed data loads from S3 buckets
- High-Availability Improvements
  - Automated sync of metadata between nodes in a cluster
  - This provides HA without the need for third-party shared storage tools and eliminates the need for NFS or EFS



# Demonstration



**THANK YOU**



**OPENWORKS**

**BE UNSTOPPABLE**