

# Greenplum: Who Are We?

Greenplum Database is the industry's first massively parallel processing (MPP) database server based on open-source technology. It is explicitly designed to support business intelligence (BI) applications and large, multi-terabyte data warehouses

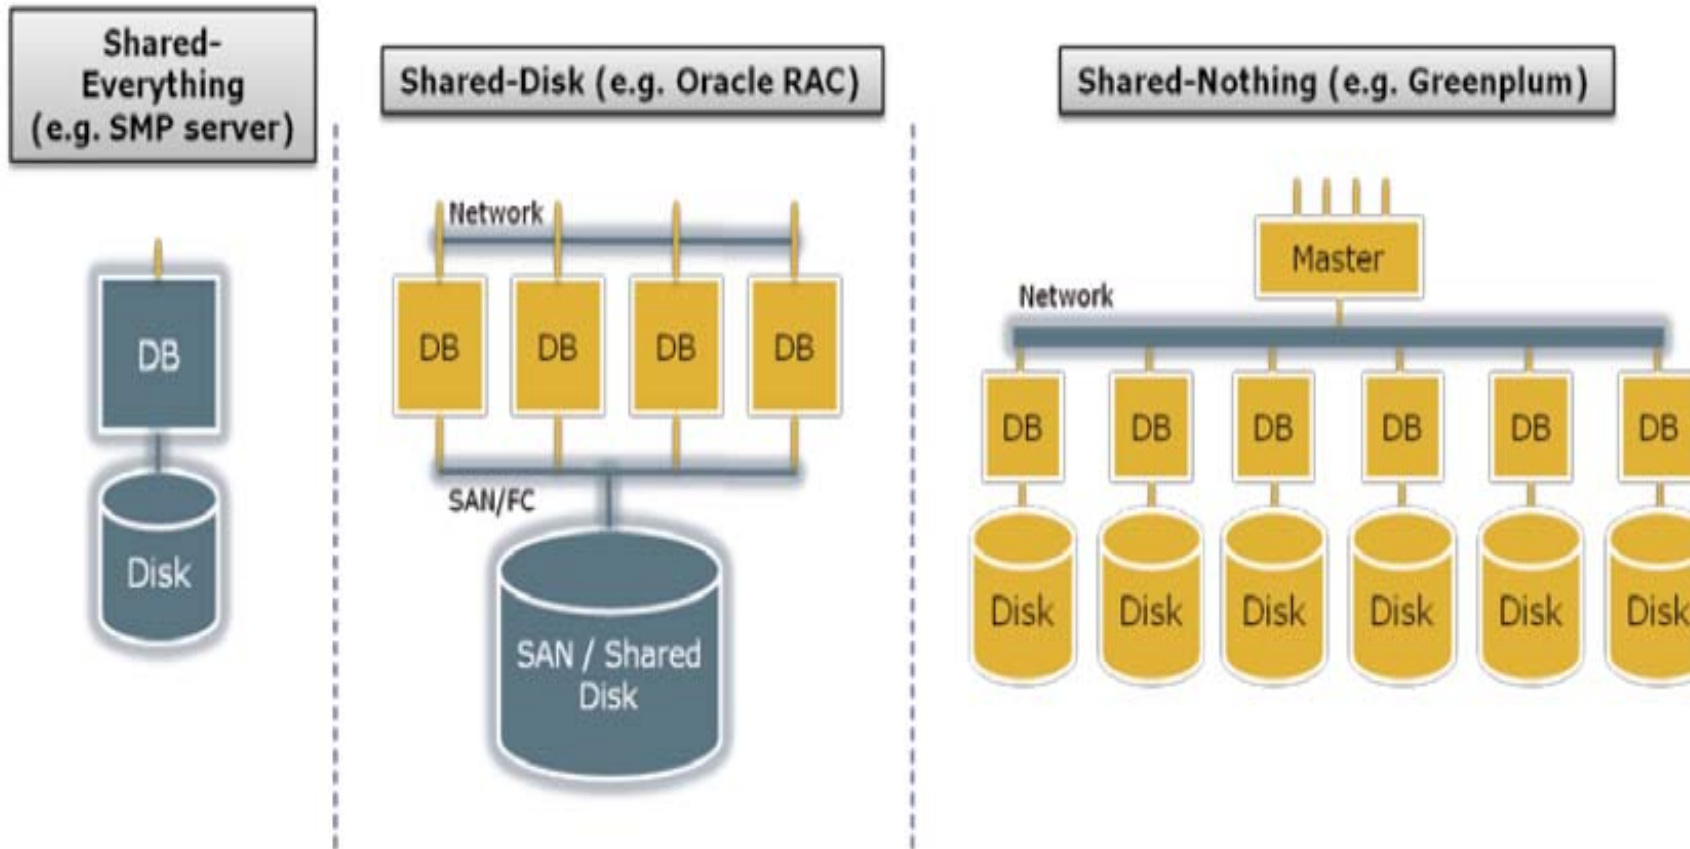
## **Features of MPP**

1. *MPP (also known as a 'shared nothing' architecture) refers to systems with two or more processors, which cooperate to carry out an operation - each processor with its own memory, operating system and disks.*
2. *Each subsystem communicates with the others via a high-speed interconnect.*
3. *Good at complex ad-hoc queries using huge volumes of data.*
4. *Perform faster when faced with BI applications that require full-table scans, many table joins, sorting, or aggregation against very large volumes of data.*
5. *Greenplum leverages this high-performance system architecture to distribute the load of large volume of data warehouses, and is able to use all of a system's resources in parallel to process a query.*

## **Features of SMP**

1. *General-purpose relational database management system designed for Online Transaction Processing (OLTP) applications.*
2. *Shared Everything Architecture*
3. *Good at running complex queries up to a tera byte data with a high level query and database tuning.*
4. *Perform poorly when faced with BI applications that require full-table scans, many table joins, sorting, or aggregation against very large volumes of data.*
5. *Loss in hardware efficiency as the number of jobs being processed increases and also limit to number of processes*

# DB Clustering: Shared Everything vs Shared Nothing



*Note: Shared resources highlighted in grey in all diagrams*

# Greenplum Customers: +60 and growing . . .

