



POZNAN UNIVERSITY OF TECHNOLOGY

Main Memory Data Warehouses

Robert Wrembel

Poznan University of Technology

Institute of Computing Science

Robert.Wrembel@cs.put.poznan.pl

www.cs.put.poznan.pl/rwrembel



Lecture outline

Teradata Data Warehouse Appliance




SAP Hana



Oracle Exadata



IBM Netezza



Targit XBone Server

EMC Greenplum Appliance





DW Appliances

- The slides about IBM Netezza were prepared based on the official IBM materials:
 - "IBM Pure Data Systems for Analytics" - workshop
 - Netezza technical documentation
 - IBM Netezza Database User's Guide. IBM Netezza 7.0.x, Oct 2012
 - IBM Netezza System Administrator's Guide. IBM Netezza 7.0 and Later, Oct 2012
 - IBM Netezza Getting Started Tips. IBM Netezza 7.0, Oct 2012
- The slides about Oracle Exadata were prepared based on:
 - Oracle Exadata Database Machine X4-2 (Oracle data sheet)
 - The Teradata Data Warehouse Appliance. Technical Note on Teradata Data Warehouse Appliance vs. Oracle Exadata



Definition

⇒ DW Appliance:

- **self-contained integrated solution stack of hardware, operating system, RDBMS software and storage, optimized for data warehouse workloads**
- **comes out of the "box" preconfigured and tuned**
- **hardware is designed to work with a particular software whereas the software is tuned to work with this hardware**



IBM Netezza (1)


➔ **The key hardware components include the following:**

- **snippet blades (S-Blades = Snippet Processing Units - SPUs)**
 - **each S-Blade owns several disks which reside in a storage array within the same rack**
- **hosts (servers)**
- **storage arrays (disks)**

IBM Netezza (2)

➔ S-Blade


- for processing data from disks
- **CPU + Netezza Database Accelerator card contains the FPGA query engines, memory, and I/O**



IBM Netezza (3)

⇒ S-Blade tasks

- decompression
- data filtering
- data projection
- SQL operations
- joins
- aggregations
- sorts
- analytical algorithms (data mining, prediction)





IBM Netezza (4)

➔ Host

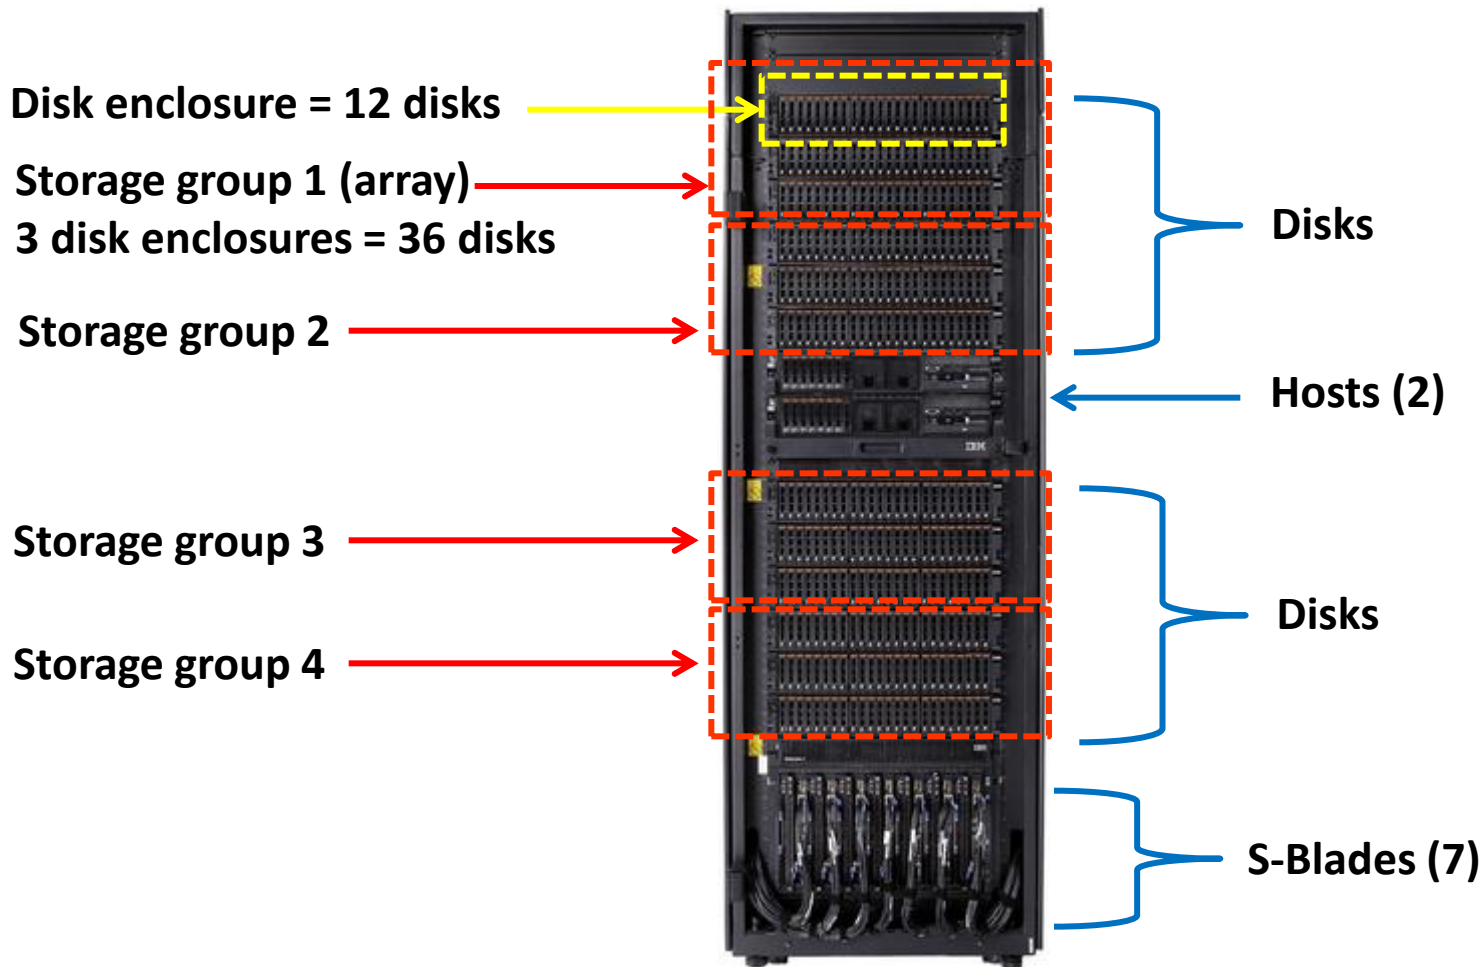
- **Linux OS**
- **administration and security**
- **workload management**
- **query optimization**
- **data loading**
- **data distribution to disks**
- **consolidating and returning query results**
- **system monitoring**
- **active**
- **spare (backup)**



IBM Netezza (5)

- ⇒ **Storage array = storage group**
 - **composed of n disk enclosures**
 - **disk enclosure = 12 disks**
 - **one appliance includes 1 to 4 storage groups**

IBM PureData System for Analytics N2001





IBM PureData System for Analytics N1001




1-rack

Model	N1001-002	N1001-005	N1001-010
S-Blades	4	7	14
Snippet Processors	24	48	96
Disks [Spares]	24 [2]	48 [2]	96 [4]
User Data Storage (raw)	8TB	16TB	32TB

Netezza C1000 Systems

**2-racks and
4-racks**




Model	C1000-8	C1000-16
S-Blades	8	16
Snippet Processors	64	128
Disks [Spares]	288 [16]	576 [32]
User Data Storage (Raw)	288 TB	576 TB

Netezza TwinFin™ 12

8 disk enclosures
12 disks/enclosure
disk capacity: 1TB
8[de] * 12[d] * 1TB = 96TB

12 S-Blades
1 blade includes:
CPU: 2 Intel Quad -Core 2GHz
4 125MHz FPGA
16GB DDR2 RAM
Linux Kernel 64-bit



Disk enclosure

Hosts (2)
one active, one passive
CPU: 2 Intel Quad-Core 2.6GHz
7x146GB SAS Drives
24 GB RAM
Red Hut Linux 5 64-bit

Data load speed: 1TB/h


Netezza TwinFin™ 24

2 * 8 disk enclosures
12 disks/enclosure
disk capacity: 1TB
 $2 * 8[de] * 12[d] * 1TB = 192TB$

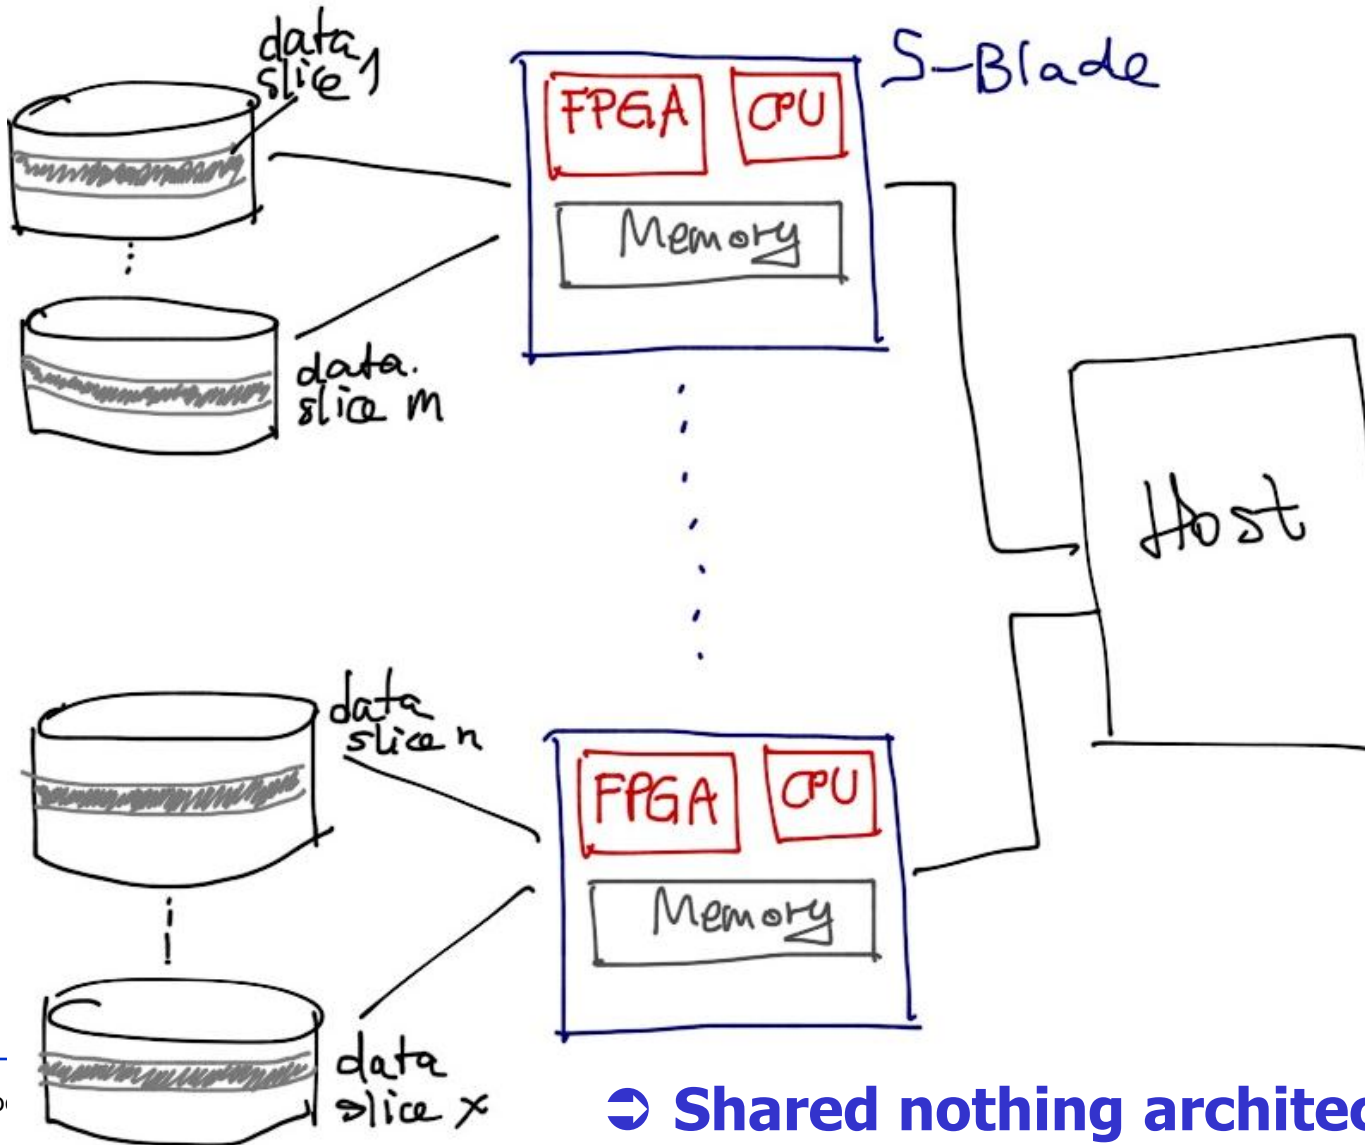
2 hosts

2 * 12 S-Blades

Data load speed: 2TB/h



Netezza Architecture



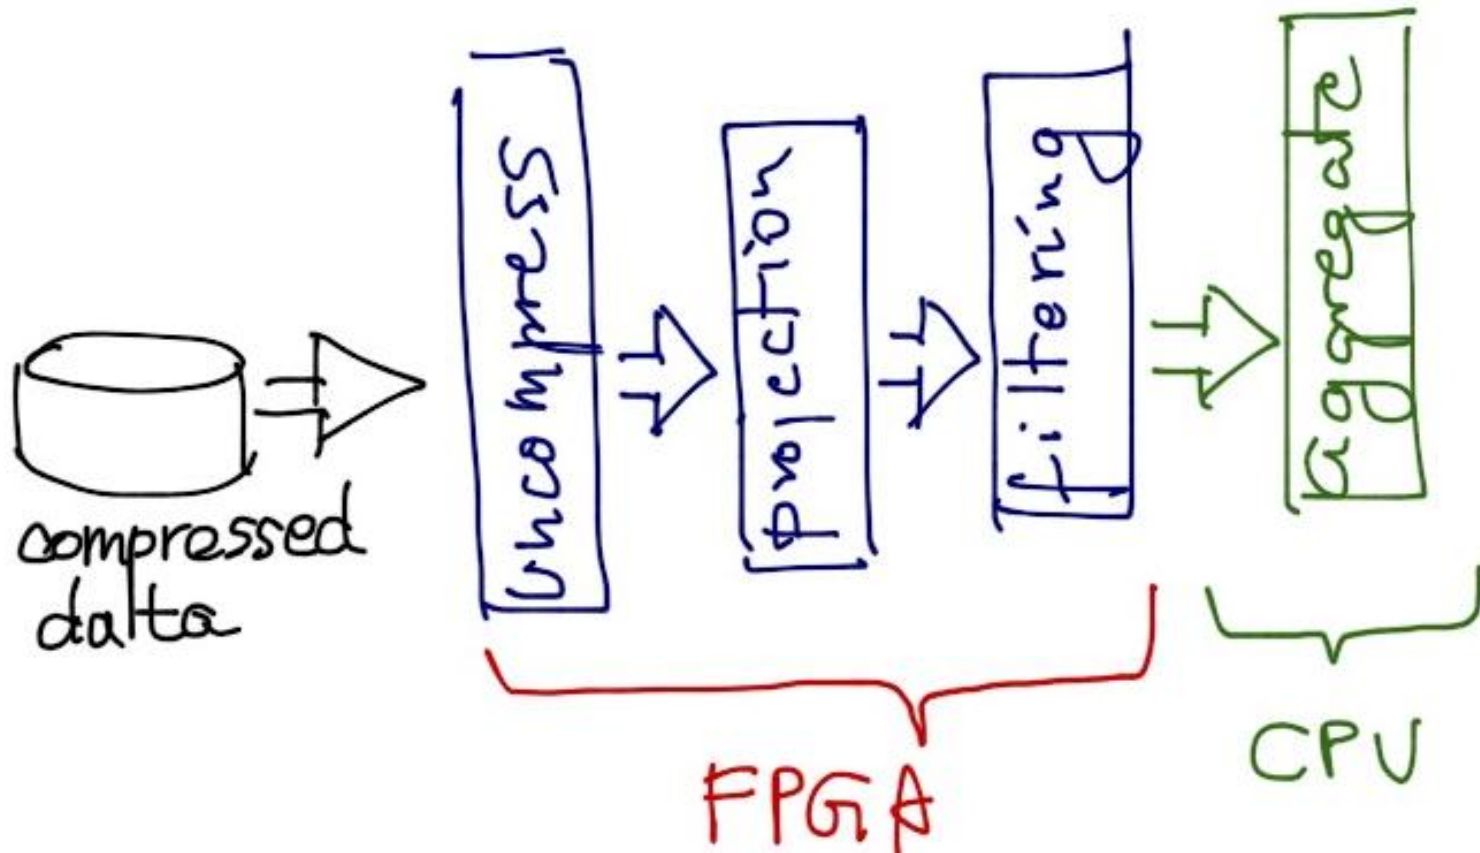


Netezza Architecture

- ➔ **Data slice** ⇒ a disk zone allocated for storing data of one table
- ➔ **Table data are distributed into data slices**
- ➔ **Table data distribution**
 - **hashing**
 - **random (round-robin)**

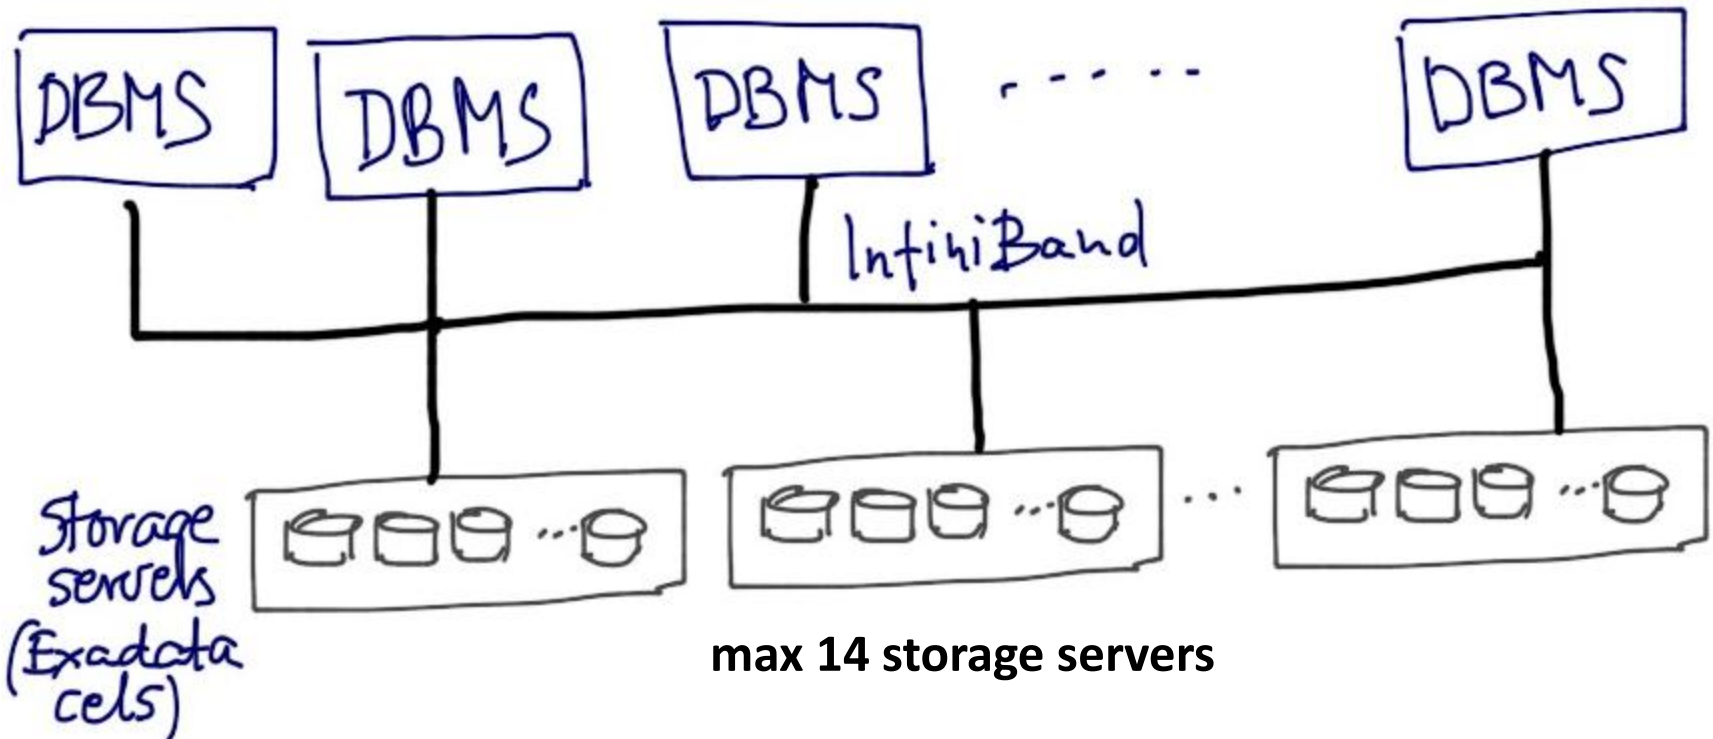
```
CREATE TABLE tab-name  
(...)  
DISTRIBUTE ON {(col1, ...) | RANDOM}
```


Data Processing in S-Blade



Oracle Exadata - Architecture

max 8 DB servers



➔ **Shared disk architecture**

Oracle Exadata - Features (1)

- ⇒ **Suitable for OLTP and OLAP**
- ⇒ **Storage server**
 - **2 CPU Intel Xeon**
 - **Smart Scan module** ⇒ similar to Netezza's S-Blade
 - parallel reads from disks
 - uncompressing
 - filtering
 - **flash memory** ⇒ used as cache for query intensive data
 - each storage server includes 4PCI flash cards of total capacity 3.2TB
 - max flash capacity $14 * 3.2 = 44.8\text{TB}$ (X4-2 series)
 - **data compression**
 - **data distribution to all disks**



Oracle Exadata - Features (2)

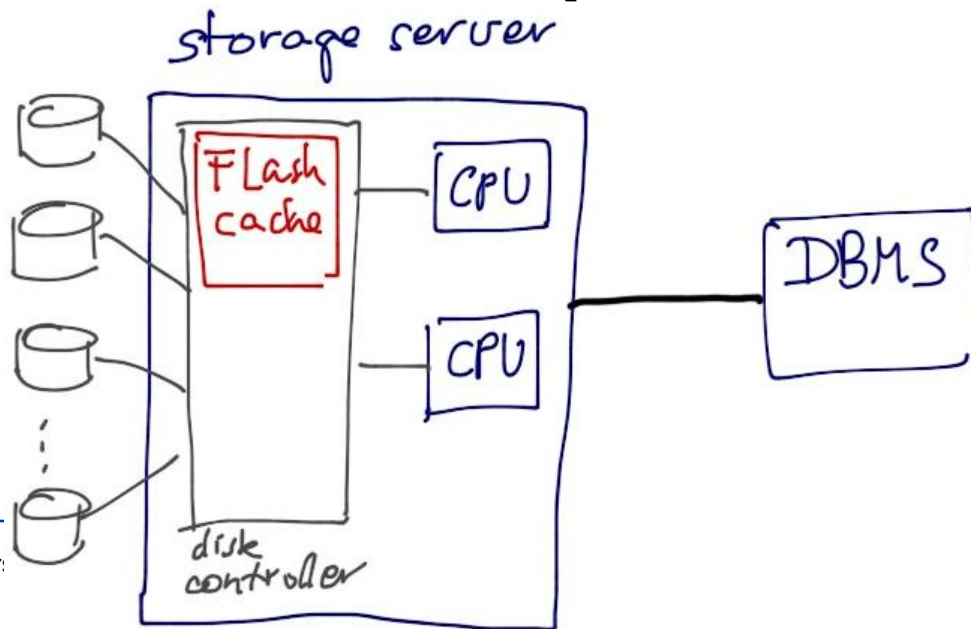
⇒ DB servers

- run under Oracle Linux or SUN Solaris
- process prefiltered data by Smart Scan modules

⇒ InfiniBand switches connect DB servers and storage servers

- 40GB/s

⇒ Max data load rate 20TB/h (full rack X4-2 series)





Oracle Exadata - Models (1)

Exadata Database Machine X3-2 Full Rack	Exadata Database Machine X3-8 Hardware
8 x Database Servers	2 x Database Servers
Each with: <ul style="list-style-type: none"> - 2 x Eight-Core Intel® Xeon® E5-2690 Processors - 256GB Memory - Disk Controller HBA with 512MB Battery Backed Write Cache - 4 x 300GB 10,000 RPM SAS Disks 	Each with: <ul style="list-style-type: none"> - 8 x Ten-Core Intel® Xeon® E7-8870 Processors (2.40 GHz) - 2 TB Memory - Disk Controller HBA with 512MB Battery Backed Write Cache - 8 x 300GB 10,000 RPM SAS Disks
14 x Exadata Storage Servers X3-2	14 x Exadata Storage Servers X3-2
With: <ul style="list-style-type: none"> - 12 x 600GB 15,000 RPM High Performance SAS disks or 12 x 3TB Performance SAS disks or 12 x 2TB 7,200 RPM High Capacity SAS disks Includes: <ul style="list-style-type: none"> - 168 CPU cores for SQL processing - 22.4TB Exadata Smart Flash Cache 	With: <ul style="list-style-type: none"> - 12 x 600GB 15,000 RPM High Performance SAS disks or 12 x 3TB 7,200 RPM High Capacity SAS disks Includes: <ul style="list-style-type: none"> - 168 CPU cores for SQL processing - 22.4TB Exadata Smart Flash Cache

Oracle Exadata - Models (2)

- ⇒ **Quarter rack (X4-2)**
 - 2 DB servers
 - 3 storage servers
- ⇒ **Half rack (X4-2)**
 - 4 DB servers
 - 7 storage servers
- ⇒ **Full rack (X4-2)**
 - 8 DB servers
 - 14 storage servers
- ⇒ **Disk types (X4-2)**
 - high performance (1.2TB)
 - high capacity (4TB)

7 Storage servers

8 DB servers

7 Storage servers





Oracle Exadata - Summary

Exadata Database Machine X4-2 Hardware			
Full Rack	Half Rack	Quarter Rack	Eighth Rack
8 x Database Servers, each with:	4 x Database Servers, each with:	2 x Database Servers, each with:	2 x Database Servers, each with:
<ul style="list-style-type: none"> • 2 x Twelve-Core Intel® Xeon® E5-2697 v2 Processors (2.7 GHz) • 256GB Memory (expandable to 512GB) • Disk Controller HBA with 512MB Battery Backed Write Cache • 4 x 600 GB 10,000 RPM Disks • 2 x QDR (40Gb/s) InfiniBand Ports 		<ul style="list-style-type: none"> • 4 x 1/10 Gb Ethernet Ports (copper) • 2 x 10 Gb Ethernet Ports (optical) • 1 x ILOM Ethernet Port • 2 x Redundant Hot-Swappable Power Supplies 	
192 CPU cores and up to 4TB memory for database processing (24 CPU cores and up to 512 GB memory per Database Server)	96 CPU cores and up to 2 TB memory for database processing (24 CPU cores and up to 512 GB memory per Database Server)	48 CPU cores and up to 1TB memory for database processing (24 CPU cores and up to 512 GB memory per Database Server)	24 CPU cores and up to 1 TB memory for database processing (12 CPU cores per Database Server are enabled with up to 512 GB memory per Database Server)
14 x Exadata Storage Servers X4-2: <ul style="list-style-type: none"> • 168 CPU cores for SQL processing • 56 PCI flash cards with 44.8 TB (raw) Exadata Smart Flash Cache • 168 x 1.2 TB 10,000 RPM High Performance disks or 168 x 4 TB 7,200 RPM High Capacity disks 	7 x Exadata Storage Servers X4-2: <ul style="list-style-type: none"> • 84 CPU cores for SQL processing • 28 PCI flash cards with 22.4 TB (raw) Exadata Smart Flash Cache • 84 x 1.2 TB 10,000 RPM High Performance disks or 84 x 4 TB 7,200 RPM High Capacity disks 	3 x Exadata Storage Servers X4-2: <ul style="list-style-type: none"> • 36 CPU cores for SQL processing • 12 PCI flash cards with 9.6 TB (raw) Exadata Smart Flash Cache • 36 x 1.2 TB 10,000 RPM High Performance disks or 36 x 4 TB 7,200 RPM High Capacity disks 	3 x Exadata Storage Servers X4-2: <ul style="list-style-type: none"> • 36 CPU cores for SQL processing (18 cores enabled) • 6 PCI flash cards with 4.8 TB (raw) Exadata Smart Flash Cache (6 more flash cards reserved for use on upgrade to quarter rack) • 18 x 1.2 TB 10,000 RPM High Performance disks or 18 x 4 TB 7,200 RPM High Capacity disks (18 more reserved for use on upgrade to quarter rack)



Teradata

⇒ **Shared nothing architecture**

⇒ **5 Models**

- **Data Mart Edition ⇒ up to 6TB**
- **Data Mart Appliance ⇒ up to 8TB**
- **Extreme Data Appliance ⇒ up to 234PB**
- **Data Warehouse Appliance ⇒ up to 21PB**
- **Active Enterprise Data Warehouse ⇒ up to 61PB**

⇒ **SUSE Linux**

⇒ **Data compression**