


## Is there any alternative to Exadata X5?

April 2015

# Contents


---




- 1 About Benchware Ltd.**
- 2 Licensing
- 3 Scalability
- 4 Storage Performance
- 5 Server Performance
- 6 Costs
- 7 Myths
- 8 Gartner says . . .

# Contents

---

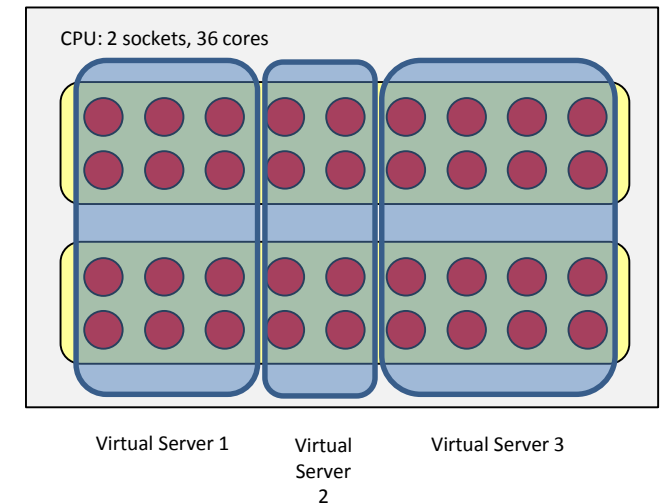


- 1 About Benchware Ltd.
- 2 Licensing**
- 3 Scalability
- 4 Storage Performance
- 5 Server Performance
- 6 Costs
- 7 Myths
- 8 Gartner says . . .




- Exadata now supports Oracle VM with trusted partitions
  - Separation of network zones
  - Separation of database classes, e.g. with different service level (prod, test) or different load profiles (OLTP, DWH)
  - Separation of applications, e.g. SAP, Siebel, ...
  - Limited licensing of Oracle options, e.g. Database In-Memory
- Benchware recommendation
  - Use only a few virtual server
  - Use Oracle Multitenant Option<sup>[1]</sup> instead of server virtualization
  - Use RAC One Node to move database services between server

X5-2 Server System




[1] Oracle Multitenant on SuperCluster T5-8: Scalability Study; An Oracle White Paper April 2014

# Capacity on Demand



- A minimum of 40% of all server cores must be licensed
- Example
  - 1/8 rack has 18 cores per server
  - Minimum of 8 cores per server must be licensed
- Entry-level Exadata now available

X5-2 Server System Eighth Rack




[2] IBM POWER supports fine granular and sophisticated server partitioning (accepted by Oracle licensing model).

[3] Oracle VM for SPARC (logical domains) and x86.


# Contents

---



- 1 About Benchware Ltd.
- 2 Licensing
- 3 Scalability**
- 4 Storage Performance
- 5 Server Performance
- 6 Costs
- 7 Myths
- 8 Gartner says . . .

# Scale-Up versus Scale-Out




- Exadata provides limited scale-up features
  - Memory capacity default 256 GByte
  - Memory extension to 512 GByte
  - Memory extension to 768 GByte




[4] For comparison IBM P8 (2-socket server) supports 1 TByte, 2 TByte planned for 2015; SPARC T7 (2-socket server) will support 4 TByte.

[5] Benchware recommends to use maximum memory capacity. High memory capacities provide more flexibility for database consolidation platforms, more KEEP cache in OLTP systems and more PGA for DWH systems and in-memory column store for analytic applications. **Block access in memory is factor 5'000 faster than block access on flash storage (100 ns versus 500 µs).**

# Scale-Up versus Scale-Out




- Some applications need
  - Large scale-up server
  - More than 2 sockets and 36 cores
  - More than 768 GByte memory
  - But not certified for RAC
- X4-8 with Oracle VM and Capacity on Demand (CoD)
  - Minimum 48 cores per server (out of 120)
- There are some drawbacks with Intel E7 processors . . .





# Workload dependent Elastic Configurations



## ■ Building Blocks


- Foundation is 1/8 or 1/4 rack
  - » 2 Exadata Database Servers
  - » 3 Exadata Storage Servers
- Add Exadata Database Server as needed
- Add Exadata Storage Server as needed
- **Even from different Exadata generations!**
- Oracle maintains interfaces, patch levels, security, etc.




X5-2 Database Server



X5-2 Storage Server HC




X5-2 Storage Server EF




# Contents

---

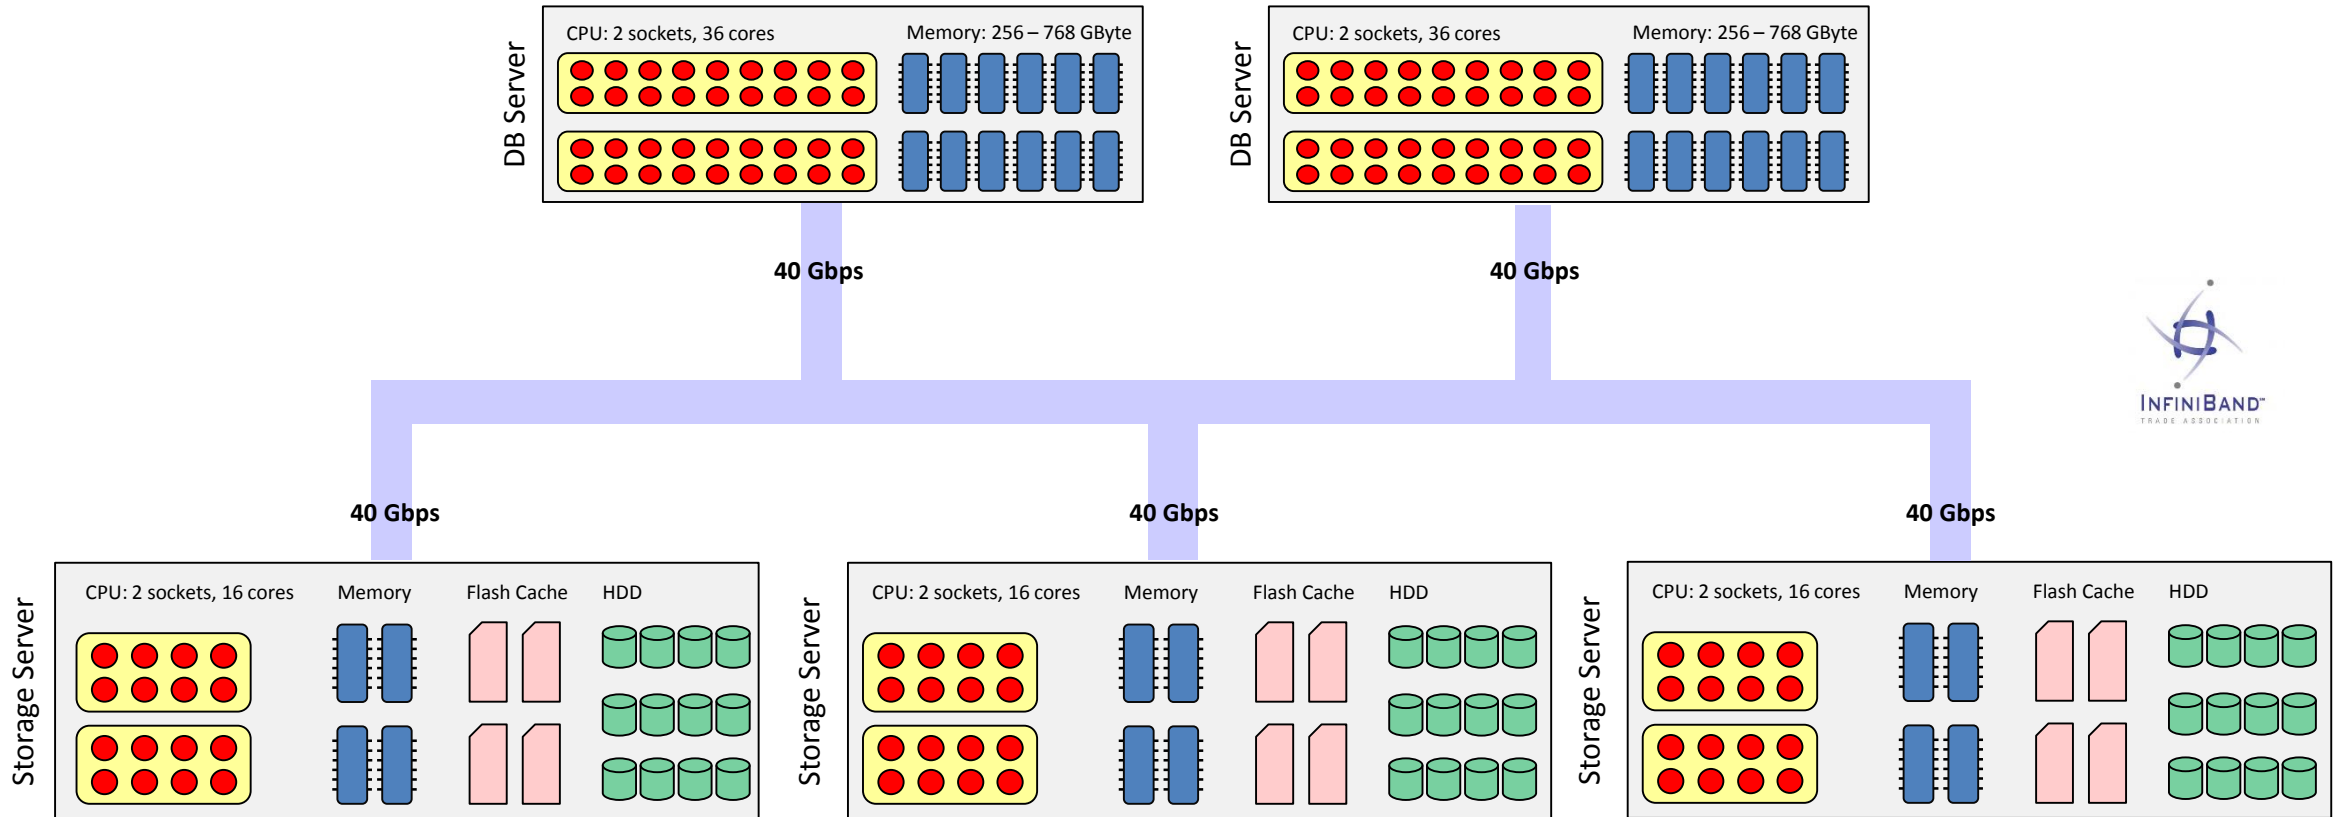


- 1 About Benchware Ltd.
- 2 Licensing
- 3 Scalability
- 4 Storage Performance**
- 5 Server Performance
- 6 Costs
- 7 Myths
- 8 Gartner says . . .


# Grid Architecture with Offload Functions




Transfer performance for each component limited by IB HCA of 40 Gbps or 4 GBps




# Grid Architecture with Offload Functions



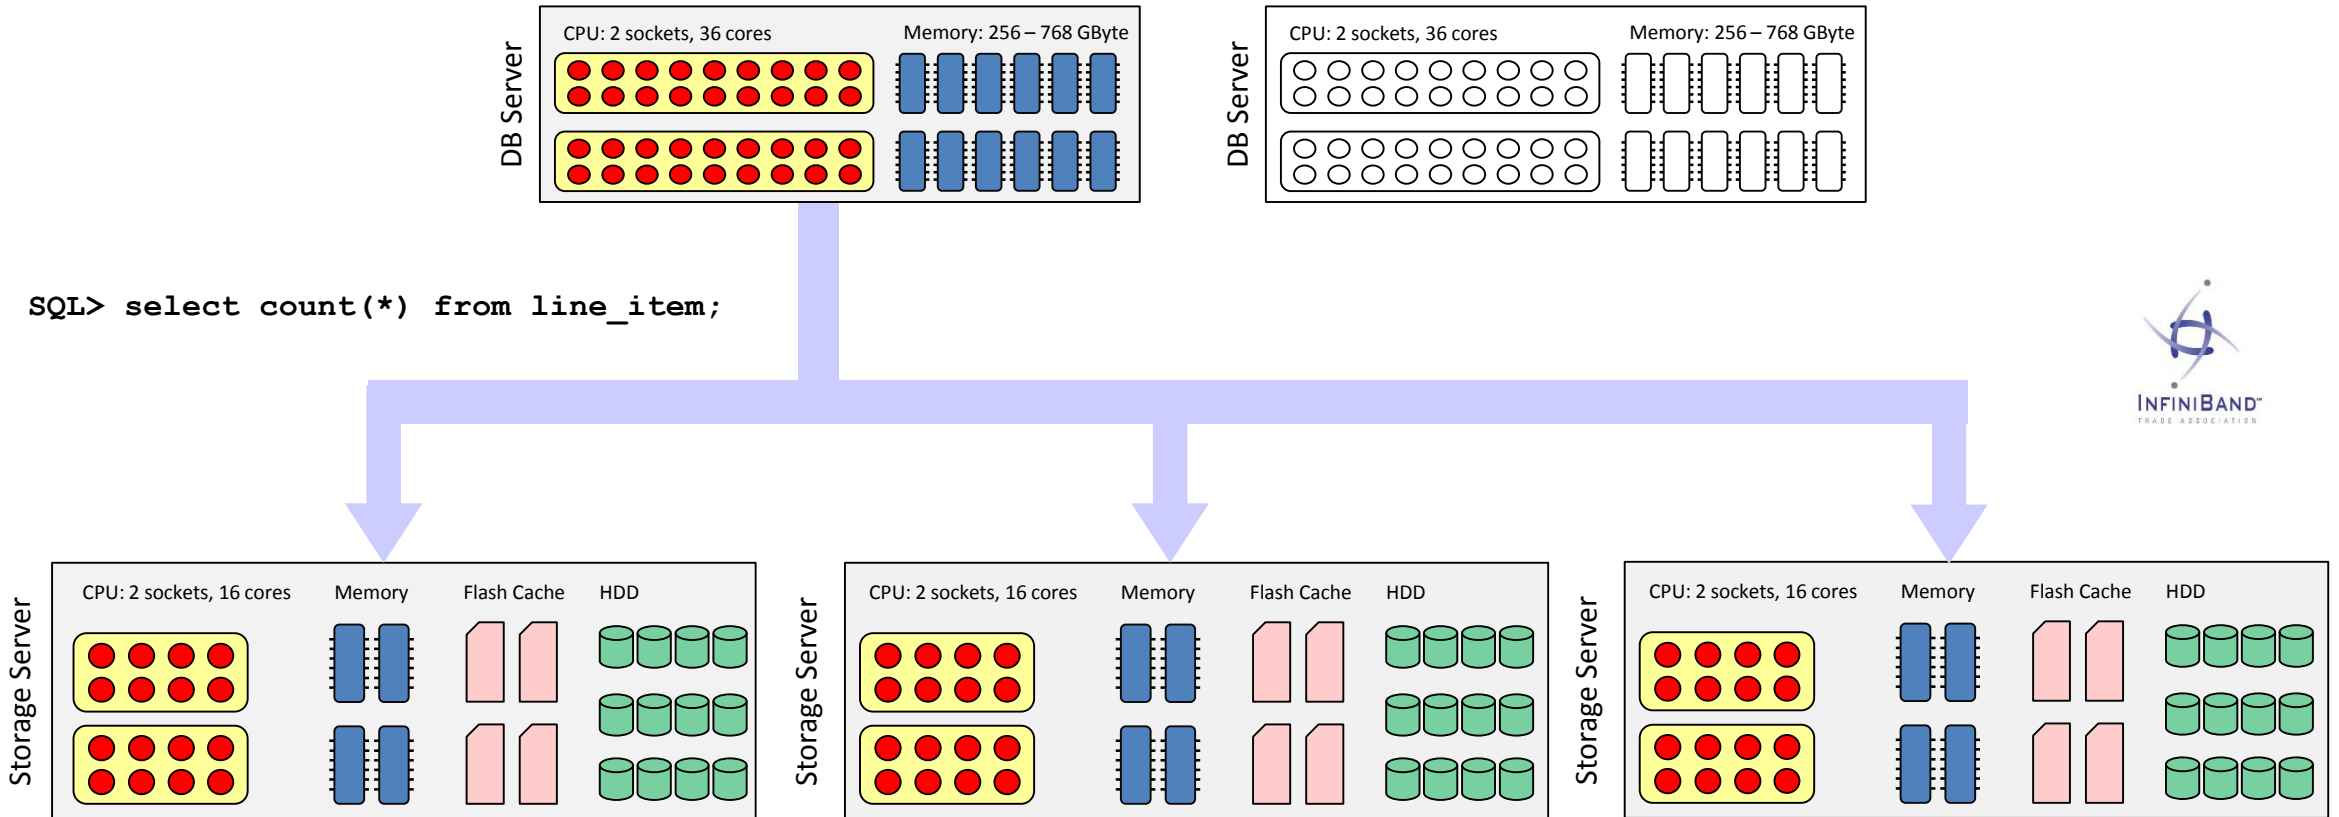
I/O offloaded from DB Server (~ 25% more overall database performance per core)



# Grid Architecture with Offload Functions




## Offload Function Smart Scan




# Flash Storage, Flash Cache and Auto-Tiering

---




- Exadata V2, X2, X3, X4
  - HDD with flash cache
  - Uses auto-tiering mechanism
- Auto-tiering algorithms work fine in most situations - but not all situations
  - Also true for shared storage systems
- Exadata X5
  - Replacement for “High Performance” Storage Server
  - Flash only, called “Extreme Flash”
  - Requires software license, 20k USD per flash module
- Benchware recommendation
  - EF for active OLTP data for predictable Oracle performance
  - HC for DWH systems or inactive data
  - Instead of auto-tiering use Oracle 12 ADO


# Performance Example DWH: Smart Scan




Oracle storage performance: sequential read, n processes, Oracle internal DOP = 4



# Performance Example OLTP: Random I/O




Oracle storage performance: random read, n processes





# Compression and Encryption




## Exadata versus competition


- Compression
  - Hybrid Columnar Compression
    - » Different options dependent on workload profile
  - Part of the Exadata storage software license
    - » 10'000 USD per hdd disk
    - » 20'000 USD per flash module
- Encryption
  - Advanced Security Option
  - License on database server
    - » 15'000 USD per core
- Some flash storage vendors offer inline **deduplication, compression and encryption** for free
- PureStorage FA-420, 2 RU, 5 TByte
  - Compression rate factor 3 for dense Oracle tables
  - Compression rate factor 15 for sparse Oracle tables
  - Additional functionality increases I/O service time, but still < 1 ms per 8 kByte

# Contents

---



- 1 About Benchware Ltd.
- 2 Licensing
- 3 Scalability
- 4 Storage Performance
- 5 Server Performance**
- 6 Costs
- 7 Myths
- 8 Gartner says . . .




- Extract from Oracle Marketing presentation

## Exadata X5-2 Database Servers

**Faster Processors, More Cores, More Memory**



- 18-core Xeon Haswell-EP - 50% faster than existing Ivy Bridge (X4-2) processor
- Faster and larger memory - DDR4 upgradeable from 256 GB (8x32GB) to 768 GB




## Other benchmark results

- CPU performance has two components <sup>[4]</sup>
  - Single-thread performance
  - Throughput

	#Cores	SPECint_base2006 (speed)	SPECint_base_rate2006 (throughput)	Oracle CPU calibration sys.aux_stats\$ (speed)
X2-2	12	40.8	367	2'795
X3-2	16	54.3	630	2'751
X4-2	24	88.1	806	3'074
X5-2	36	65.2	1'380	2'890

[4] [www.spec.org](http://www.spec.org)



## Benchmark results

### ■ PL/SQL performance

- Native compiled SIMPLE\_INTEGER
- Arithmetic addition

	#Cores	Speed [Mops]	Throughput [Mops]
X2-2	12	452	5'414 (451)
X3-2	16	423	8'465 (530)
X4-2	24	434	13'300 (554)
X5-2	36	409	21'070 (585)

[5] in brackets throughput per core

### ■ SQL performance

- Cached table (no I/O)
- Light OLTP


	#Cores	Speed [tps@servicetime]	Throughput [tps]
X2-2	12	19'160@58μs	270'600 (22'550)
X3-2	16	18'380@48μs	373'400 (23'337)
X4-2	24	18'150@55μs	545'700 (22'738)
X5-2	36	42'900@23μs	1'215'000 (33'750)

[6] for comparison IBM P8 (3.425 GHz) achieves 41'250 tps @24μs speed and 69'665 tps throughput per core


[7] X5-2 uses DDR-4 RAM, X4-2 and previous systems DDR-3 RAM

# Contents

---




- 1 About Benchware Ltd.
- 2 Licensing
- 3 Scalability
- 4 Storage Performance
- 5 Server Performance
- 6 Costs**
- 7 Myths
- 8 Gartner says . . .



Similar to conventional platforms

- Pro Exadata
  - Extreme efficient platform [8]
  - I/O Performance
    - » Sequential read (smart scan)
  - High availability included (MAA ready)
  - Uniform management tool for Database Machine Administrator
  - Comprehensive Resource Manager included
  - Investment protection
  - . . .
- Pro conventional platform
  - Free choice of number of sockets and number of cores for database servers
  - Scale-Up and Scale-Out
  - Oracle SE possible
  - Free choice of combination of hdd and flash technology
  - Some flash storage vendors deliver transparent deduplication, compression and encryption for free

[8] Logistics company (450+ Oracle databases) migrated from 120 2-socket blades to 12 2-socket Exadata X4-2 database server



## Advantages for Exadata


- Pro Exadata
  - Time to market
    - » Nearly no engineering
    - » Savings around 3 – 6 months elapsed time
    - » No interface issues
  - Maturity and stability (based on best practices and standards)
  - Support
  - Patching
  - Agility (not possible with silo organization)
- Pro conventional platform
  - ?

[9] The Exadata case; SwissRe internal study of Exadata X2-2 Full Rack over 2 years




# Contents

---




- 1 About Benchware Ltd.
- 2 Licensing
- 3 Scalability
- 4 Storage Performance
- 5 Server Performance
- 6 Costs
- 7 Myths**
- 8 Gartner says . . .



- Exadata is only for large companies
- Early Warning Systems GmbH (Zurich)
  - Subsidiary of FIFA
  - 10 employees, 3 people in IT
  - Running Exadata X3-2 QR
  - External support for patching
  - Oracle 12c since January 2015


# Contents

---



- 1 About Benchware Ltd.
- 2 Licensing
- 3 Scalability
- 4 Storage Performance
- 5 Exadata Performance
- 6 Costs
- 7 Myths
- 8 Gartner says . . .**

# Gartner says . . .



---

## Ten Trends and Technologies to Impact IT Over the Next Five Years

- No. 4 – Software-Defined Infrastructures
  - New way to operate, orchestrate and automate
  - Allows configuration from one place
  - Organizationally disruptive
  
- No. 5 – Integrated Systems Evolution
  - Stack sale versus component sale
  - Optimized for special purpose
  - Best of brand versus best of breed

---

[10] Ten Trends and Technologies to Impact IT Over the Next Five Years; Gartner Webinar 15<sup>th</sup> of April 2015, David Cappucio

**BENCHWARE**

*swiss precision in performance measurement*

*[www.benchmarkware.ch](http://www.benchmarkware.ch)*

*[info@benchmarkware.ch](mailto:info@benchmarkware.ch)*