

# ORACLE EXADATA STORAGE SERVER X3-2

## KEY FEATURES AND BENEFITS

### FEATURES

- 12 x 3.5 inch High Performance or High Capacity disks
- 1.6 TB of Exadata Smart Flash Cache
- 12 CPU cores dedicated to SQL processing in storage
- 64 GB memory
- Dual InfiniBand ports
- Redundant power supplies
- Oracle Exadata Storage Server software
- Oracle Linux or Solaris based Database Machines

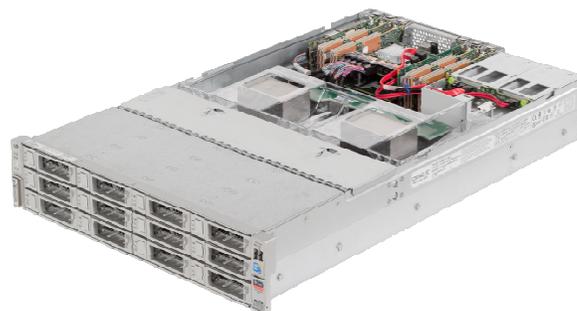
### BENEFITS

- Uncompressed usable capacity of up to 3.25 TB per server when using High Performance disks
- Uncompressed usable capacity of up to 16 TB per server when using High Capacity disks
- Hybrid Columnar Compression delivers 10X-15X compression ratios
- Effective data bandwidth of up to 7.25 GB/second per server with Flash and uncompressed data
- Effective data bandwidth of up to 72.5 GB/second per server with Flash and compressed data
- Software pre-installed

*The Oracle Exadata Storage Server X3-2 is the storage building block of the Oracle Exadata Database Machine, Oracle SuperCluster and Oracle Exadata Storage Expansion Rack. The Exadata Storage Server is highly optimized for use with the Oracle Database and employs a massively parallel architecture and Exadata Smart Flash Cache to dramatically accelerate Oracle Database processing and speed I/O operations. It may be attached to an Exadata Database Machine or SuperCluster to build out the storage capacity and I/O bandwidth and is ideal for Online Transaction Processing (OLTP), Data Warehousing (DW) and consolidation of mixed workloads. Simple to deploy and manage, the Oracle Exadata Storage Server provides linear I/O scalability and mission-critical reliability.*

### Oracle Exadata Storage Server

The Oracle Exadata Storage Server is a fast, reliable, high capacity, industry-standard storage server. Each server comes preconfigured with: 2 x six-core Intel® Xeon® E5-2630L Processors, 64 GB memory, 1.6 TB of Exadata Smart Flash Cache, 12 disks connected to a storage controller with 512MB battery-backed cache, and dual port InfiniBand connectivity. All software is preinstalled and comes complete with management interface for remote access, dual-redundant hot-swappable power supplies and takes up 2U in a 19-inch rack.



The Oracle Exadata Storage Server comes with either twelve 600 GB 15,000 RPM High Performance disks or twelve 3 TB 7,200 RPM High Capacity disks. The High Performance based Exadata Storage Servers provide up to 3.25 TB of uncompressed usable capacity, and up to 1.8 GB/second of raw data bandwidth. The High Capacity based Exadata Storage Servers provide up to 16 TB of uncompressed usable capacity, and up to 1.3 GB/second of raw data bandwidth. When stored in compressed format, the amount of user data and the amount of data bandwidth

**RELATED PRODUCTS AND SERVICES**

## RELATED PRODUCTS

- Oracle Exadata Database Machine X3-2
- Oracle Exadata Database Machine X3-8
- Oracle Exadata Storage Expansion Rack X3-2
- Oracle SuperCluster
- Oracle Database 11g and 12c
- Real Application Clusters
- Partitioning
- Multitenant
- Advanced Compression
- Advanced Security
- Active Data Guard
- Real Application Testing
- OLAP
- Advanced Analytics
- Business Intelligence
- Enterprise Manager
- Oracle Linux
- Oracle Solaris

## RELATED SERVICES

The following services are available from Oracle:

- Advanced Customer Services
- Oracle Premier Support for Systems
- Oracle Infrastructure as a Service On-Premise (IaaS)
- Oracle Platinum Services
- Oracle PlatinumPlus Services
- Consulting Services
- Oracle University courses

delivered by each cell increases up to 10 times.

**Extreme Performance by Offloading Data Intensive Processing**

The Oracle Database and the Exadata Storage Server Software includes a unique technology that offloads data intensive SQL operations from the database servers into the Exadata Storage Servers. By pushing SQL processing to the Exadata Storage Servers, data filtering and processing occurs immediately and in parallel across all storage servers as data is read from disk. Exadata storage offload reduces database server CPU consumption and greatly reduces the amount of data moved between storage and database servers. The CPUs in Exadata Storage Servers do not replace database CPUs. Instead they accelerate data intensive workloads similar to how graphics cards accelerate image intensive workloads.

**Extreme Performance from Exadata Smart Flash Cache**

Exadata systems use the latest PCI flash technology rather than flash disks. PCI flash delivers ultra-high performance by placing flash directly on the high speed PCI bus rather than behind slow disk controllers and directors. Each Exadata Storage Server includes 4 PCI flash cards with a total capacity of 1.6 TB of flash memory. A full rack Exadata Database Machine includes 56 PCI flash cards providing 22.4 TB of flash memory.



Sun Flash Accelerator F40 PCIe Card

The Exadata Smart Flash Cache automatically caches frequently accessed data in PCI flash while keeping infrequently accessed data on disk drives. This provides the performance of flash with the capacity and low cost of disk. The Exadata Smart Flash Cache understands database workloads and knows when to avoid caching data that will never be reused or will not fit in the cache. The Oracle Database and Exadata storage optionally allow the user to provide SQL directives at the database table, index and segment level to ensure that specific data is retained in flash. Tables can be retained in flash without the need to move the table to different tablespaces, files or LUNs like you would have to do with traditional storage and flash disks.

The combination of scale-out storage, InfiniBand networking, database offload, and PCI flash allows a single rack Exadata Database Machine rack to achieve up to 100 GB per second of data scan bandwidth, and up to 1,500,000 random 8K read I/O operations per second (IOPS) when running database workloads. This performance is orders of magnitude faster than traditional database architectures. It is important to note that these are real-world end-to-end performance figures measured running SQL workloads inside a single rack Exadata system. They are not component level

measurements based on low level IO measurement tools.

The Exadata Smart Flash Cache also caches database block writes. Write caching eliminates disk bottlenecks in large scale OLTP and batch workloads. The flash write capacity of a single Exadata rack exceeds 1,000,000 8K write I/Os per second. The Exadata write cache is transparent, persistent, and fully redundant. The I/O performance of the Exadata Smart Flash Cache is comparable to dozens of enterprise disk arrays with thousands of disk drives.

The Exadata Smart Flash Cache also implements algorithms to reduce the latency of log write I/Os. The time to commit user transactions or perform critical updates is very sensitive to the latency of log writes. Smart Flash Logging takes advantage of the flash memory in Exadata storage combined with the high speed RAM memory in the Exadata disk controllers to speed up log writes and bypass the latency spikes that can occur in flash solutions. The Exadata Smart Flash Logging algorithms are unique to the Exadata system.

Flash and RAM memory are central to the architecture of the Exadata Database Machine X3. Prior database systems were disk-centric with flash memory used to accelerate database execution. The Exadata Database Machine X3 systems are flash-centric, with large RAM memory footprints used to further accelerate workload execution. In order to realize the highest level of performance at the lowest cost, the Exadata X3 system implements a mass memory hierarchy that automatically moves active data into flash memory and the most active data into RAM memory while keeping less active data on low-cost disks.

Exadata uses only enterprise grade flash that is designed by the flash manufacturer to have high endurance. Exadata is designed for mission critical workloads and therefore does not use consumer grade flash that can potentially degrade or fail unexpectedly after a few years of usage. The enterprise grade flash chips used in Exadata X3 have an expected endurance of 10 years or more for typical database workloads.

The automatic data tiering between RAM, flash and disk implemented in Exadata provides tremendous advantages over other flash-based solutions. When third-party flash cards are used in a traditional database server, the data placed on flash is only available on that server since flash cards cannot be shared between servers. This precludes the use of RAC and limits the database deployment to the size of single server handicapping performance, scalability, availability, and consolidation of databases. Any component failure, like a flash card, in a single server can lead to a loss of database access. Local flash can become crippled without failing when it hangs, experiences intermittent poor performance, or data loss on power cycle, and may not trigger failover or alerts. In addition it will not automatically be protected by high-level mirroring like Exadata. There is no Hybrid Columnar Compression so the expensive flash resource is underutilized and does not provide I/O resource management for prioritizing bandwidth. Third party flash lacks the storage hierarchy and tiering automatically provided in Exadata and is much more complex to administer.

When considering third party flash arrays (storage arrays comprised solely of flash)

they have other limitations. Flash arrays will be data bandwidth limited since they have to transfer all the data from SQL queries to the database server as they do not provide offload functionality like Exadata Smart Scans. They require much more administrative overhead than Exadata storage since LUNs, mount points and tablespaces must be created. No special backup procedures are required for Exadata flash while backup and recovery procedures are required for flash arrays. With flash arrays there is no automatic dynamic tiering of the data between disk, flash and memory, as in Exadata. And of critical importance, as data usage patterns change the administrator must monitor and reconfigure the data placed on flash arrays to maintain good performance since the data is statically tied to the flash array; unlike Exadata. None of this extra administrative overhead is required for Exadata.

### **Optimizing Storage Use and I/O Through Compression**

Compressing data provides dramatic reduction in the storage consumed for large databases. The Exadata Storage Server provides a very advanced compression capability called Hybrid Columnar Compression (HCC). Hybrid Columnar Compression enables the highest levels of data compression and provides tremendous cost-savings and performance improvements due to reduced I/O. Storage savings range from 5x to 20x with typical storage savings of 10x. On conventional systems, enabling high data compression has the drawback of reducing performance. Because the Exadata Database Machine is able to offload compression overhead into large numbers of processors in Exadata storage, most workloads run faster using Hybrid Columnar Compression than they do without it. Hybrid Columnar Compression combines the compression and analytic performance benefits of column storage while avoiding the dramatic slowdown that pure columnar stores experience for drilldown operations.

Two modes of Hybrid Columnar Compression are available. Query optimized compression mode is suitable for read intensive workloads such as Data Warehouses and provides large storage savings while maintaining excellent performance. Archive compression mode provides the highest degree of compression and is targeted at rarely accessed data that is kept online.

### **Intelligent Scalable Storage Grid**

Oracle Exadata Storage Servers are installed in to a customer supplied 19-inch rack and is connected to an Exadata Database Machine or SuperCluster via InfiniBand. Exadata Storage Servers have dual 40 Gigabit InfiniBand links that provide connectivity many times faster than traditional storage or server networks. Further, Oracle's interconnect protocol uses direct data placement to ensure very low CPU overhead by directly moving data from the wire to database buffers with no extra data copies.

Exadata Storage Servers are architected to scale-out easily. To achieve higher performance and greater storage capacity, additional Exadata Storage Servers can be connected to an Oracle Database Machine or SuperCluster. This, combined with faster InfiniBand interconnect, Exadata Smart Flash Cache and the reduction of data transferred due to offload processing and Hybrid Columnar Compression, yields very large performance improvements. A 10x improvement in query performance

compared to traditional database storage architectures is common, with much greater improvement possible.

An alternative to the purchase of individual Exadata Storage Servers with the requisite rack, InfiniBand switches, cables and other infrastructure built and configured by the customer, is the Exadata Storage Expansion Rack. The Exadata Storage Expansion Rack enables you to easily grow the Exadata storage capacity and bandwidth of any X3-2 and X3-8 Exadata Database Machine or SuperCluster. It is designed for database deployments that require very large amounts of data including: historical or archive data; backups and archives of Exadata Database Machine data; documents, images, file and XML data, LOBs and other large unstructured data. Available in Full Rack, Half Rack and Quarter Rack versions it connects to the Exadata Database Machine or SuperCluster using the integrated InfiniBand fabric to easily scale the system to any capacity. The expansion rack is extremely simple to configure as there are no LUNs or mount points to configure. Storage is configured and added to a database with a few simple commands, completed in minutes. ASM dynamically and automatically balances the data across Exadata Storage Servers, online, evenly spreading the I/O load across the racks fully utilizing all the hardware and easily integrating the expansion rack into the configuration.

#### **Mission Critical High Availability**

The Oracle Exadata Storage Server has complete redundancy built in to support the demands of mission critical applications. Each Exadata Storage Server has dual port InfiniBand connections and dual-redundant, hot-swappable power supplies for high availability. Automatic Storage Management, a feature of the Oracle Database, provides disk mirroring. Hot swappable Exadata disks ensure the database can tolerate disk drive failures. In addition, data is mirrored across storage servers to ensure that storage server failure will not cause loss of data, or inhibit data accessibility.

Oracle Exadata Storage Server X3-2 Hardware	
The Oracle Exadata Storage Server comes preconfigured with:	
<b>Processors</b>	2 x Six-Core Intel® Xeon® E5-2630L (2.0 GHz) Processors
<b>Exadata Smart Flash Cache</b>	1.6 TB
<b>System Memory</b>	64 GB
<b>Disk Controller</b>	Disk Controller HBA with 512MB Battery Backed Write Cache
<b>InfiniBand Connectivity</b>	Dual-Port QDR (40Gb/s) InfiniBand Host Channel Adapter
<b>Power Supplies</b>	Dual-redundant, hot-swappable power supply
<b>Disk Drives</b>	12 x 600 GB 15,000 RPM High Performance or 12 x 3 TB 7,200 RPM High Capacity For raw disk capacity, 1 GB = 1 billion bytes. Actual formatted capacity is less.
<b>Remote Management</b>	Integrated Lights Out Manager (ILOM) Ethernet port

Oracle Exadata Storage Server X3-2 Key Capabilities	
<b>High Performance disks</b>	<ul style="list-style-type: none"> <li>• Up to 1.8 GB/second of uncompressed raw disk bandwidth per cell</li> <li>• Up to 7.25 GB/second of uncompressed Flash data bandwidth per cell</li> <li>• 7.2 TB of raw disk data capacity</li> <li>• Up to 3.25 TB of uncompressed usable capacity per cell</li> </ul>
<b>High Capacity disks</b>	<ul style="list-style-type: none"> <li>• Up to 1.3 GB/second of uncompressed raw disk bandwidth per cell</li> <li>• Up to 6.75 GB/second of uncompressed Flash data bandwidth per cell</li> <li>• 36 TB of raw disk data capacity</li> <li>• Up to 16 TB of uncompressed usable capacity per cell</li> </ul>

Oracle Exadata Storage Server X3-2 Environmental Specifications	
<b>Dimensions and Weight</b>	<ul style="list-style-type: none"> <li>• Height: 3.5 in. (87.6 mm)</li> <li>• Width: 17.5 in. (445.0 mm)</li> <li>• Depth: 29.0 in. (737.0 mm)</li> <li>• Weight: 70 lbs. (31.8 kg)</li> </ul>
<b>Environment</b>	<ul style="list-style-type: none"> <li>• Operating temperature: 5 °C to 35 °C (41 °F to 95 °F)</li> <li>• Non-operating temperature: -40 °C to 70 °C (-40 °F to 158 °F)</li> <li>• Operating relative humidity 10% to 90% non-condensing</li> <li>• Non-operating relative humidity: up to 93%, non-condensing</li> <li>• Operating altitude: Up to 3,000 m, maximum ambient temperature is derated by 1° C per 300 m above 900 m</li> <li>• Non-operating altitude: Up to 12,000 m</li> <li>• Acoustic noise: 7.61 B operating</li> </ul>
<b>Power</b>	<ul style="list-style-type: none"> <li>• Dual-redundant, hot-swappable power supply</li> <li>• Maximum output power: 1000 W</li> <li>• Maximum AC input current at 100 V AC and 1,000 W output: 12.0 A</li> <li>• Specified power supply efficiency at 1,000 W (100%) load: 91%</li> </ul>
<b>Regulations</b> <sup>1</sup>	<ul style="list-style-type: none"> <li>• Safety: EN 60950-1:2006 + A11:2009, IEC 60950-1:2005, 2nd Edition (Evaluated to all CB countries ), UL 60950-1, 2nd Edition, CSA C22.2 No. 60950-1-07</li> <li>• RFI/EMI: EN55022:2006 + A1:2007 Class A, EN61000-3-2:2006 + A1:2009 + A2:2009, EN61000-3-3:2008, 47CFR15 Subpart B (FCC) Class A, ICES-003 Class A, AS/NZS CISPR22:2006 Class A, CISPR22:2008: Class A</li> <li>• Immunity: EN 55024:2010</li> </ul>
<b>Certifications</b> <sup>1</sup>	<ul style="list-style-type: none"> <li>• Safety: UL/cUL, CE, BSMI, GOST R, S-Mark, CSA C22.2 No. 60950-1-07 2nd Ed.</li> <li>• EMC: CCC, CE, FCC, VCCI, ICES, C-Tick, KCC, GOST R, BSMI Class A</li> <li>• Other: Complies with WEEE Directive (2002/96/EC) and RoHS Directive 2011/65/EU (2002/95/EC)</li> </ul>
<sup>1</sup> In some cases, as applicable, regulatory and certification compliance were obtained at the component level.	

Oracle Exadata Storage Server X3-2 Software
<ul style="list-style-type: none"> <li>• Oracle Exadata Storage Server Software 11g Release 2 (release 11.2.3.2 or later)</li> <li>• Oracle Linux 5 Update 8 with the Unbreakable Enterprise Kernel</li> <li>• Requires Oracle Database 11g Release 2 Enterprise Edition or Oracle Database 12c Enterprise Edition for the database accessing Exadata storage</li> </ul>
High-Availability Features
<ul style="list-style-type: none"> <li>• Redundant power supplies</li> <li>• Redundant InfiniBand ports</li> <li>• Hot swappable disk drives</li> <li>• Oracle Automatic Storage Management: All database files either double or triple mirrored; Disk failures do not abort queries or transactions</li> <li>• Oracle Exadata Storage Server Software: Storage server failure can be tolerated without data loss or aborting queries or transactions</li> </ul>
Manageability Features
<ul style="list-style-type: none"> <li>• Oracle Embedded Integrated Lights Out Manager (ILOM)</li> <li>• Oracle Enterprise Manager Cloud Control 12c</li> </ul>
Support Services Provided By Oracle
<ul style="list-style-type: none"> <li>• Hardware Warranty: 1 year with a 4 hour web/phone response during normal business hours (Mon-Fri 8AM-5PM), with 2 business day on-site response/Parts Exchange</li> <li>• Oracle Premier Support for Systems: Oracle Linux and Solaris support and 24x7 with 2 hour on-site hardware service response (subject to proximity to service center)</li> <li>• Oracle Premier Support for Operating Systems</li> <li>• Oracle Customer Data and Device Retention</li> <li>• System Installation Services</li> <li>• Software Configuration Services</li> <li>• Oracle Exadata Start-Up Pack</li> <li>• System Upgrade Support Services including hardware installation and software configuration</li> <li>• Oracle Auto Service Request (ASR)</li> </ul>

## Contact Us

For more information about the Oracle Exadata Storage Server, please visit [oracle.com](http://oracle.com) or call +1.800.ORACLE1 to speak to an Oracle representative.



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All trademarks are used under license and are trademarks or registered trademarks of International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0110