

ORACLE EXADATA DATABASE MACHINE X8-2

ORACLE®
Exadata

The Oracle Exadata Database Machine is engineered to deliver dramatically better performance, cost effectiveness, and availability for Oracle databases. Exadata features a modern cloud-enabled architecture with scale-out high-performance database servers, scale-out intelligent storage servers with state-of-the-art PCI Flash, and an ultra-fast InfiniBand internal fabric that connects all servers and storage. Unique algorithms and protocols in Exadata implement database intelligence in storage, compute, and InfiniBand networking to deliver higher performance and capacity at lower costs than other platforms. Exadata runs all types of database workloads including Online Transaction Processing (OLTP), Data Warehousing (DW), In-Memory Analytics as well as consolidation of mixed workloads. Simple and fast to implement, the Exadata Database Machine powers and protects your most important databases. Exadata can be purchased and deployed on premises as the ideal foundation for a private database cloud, or it can be acquired using a subscription model and deployed in the Oracle Public Cloud or Cloud at Customer with all infrastructure management performed by Oracle.



Engineered System for Fast and Reliable Deployment

The Exadata Database Machine is an easy to deploy system that integrates all the hardware needed for running Oracle Database. The database servers, storage servers and network are pre-configured, pre-tuned and pre-tested by Oracle experts, eliminating weeks or months of effort typically required to deploy mission critical systems. Extensive end-to-end testing ensures that all components including database software, operating system, firmware, drivers, etc., work seamlessly together and that there are no performance bottlenecks or single points of failure that can affect the complete system.

Because all Exadata Database Machines are identically configured, customers benefit from the experience of thousands of other users that have deployed the Exadata Database Machine for their mission critical applications. Customer machines are also identical to the machines Oracle Support uses for problem identification and resolution, the machines Oracle Development uses for development and testing of the Oracle Database, and the machines Oracle uses to implement its own SaaS and PaaS public cloud services. Hence, **Exadata is the most thoroughly tested and tuned platform for running Oracle Database, and is also the most supportable platform.**

KEY FEATURES

- Up to 912 CPU cores and 28.5TB memory per rack for database processing
- Up to 576 CPU cores per rack dedicated to SQL processing in storage
- From 2 to 19 Database Servers per rack
- From 3 to 18 Storage Servers per rack
- Up to 920 TB of Flash capacity (raw) per rack
- Up to 3.0 PB of disk capacity (raw) per rack
- Hybrid Columnar Compression often delivers 10X-15X compression ratios
- 40 Gb/second (QDR) InfiniBand Network
- Complete redundancy for high availability

KEY BENEFITS

- Pre-configured, pre-tested system optimized for all database applications
- Uncompressed I/O bandwidth of up to 350 GB/second per full rack from SQL
- Ability to perform up to 4.8M 8K database read I/O operations, or 4.3M 8K Flash write I/O operations per second per full rack
- Easily add compute or storage servers to meet the needs of any size application
- Scale by connecting multiple Exadata Database Machine X8-2 racks or Exadata Storage Expansion Racks. Up to 18 racks can be connected by simply adding InfiniBand cables and internal switches. Larger configurations can be built with external InfiniBand switches

"I would recommend Exadata as a platform for performance, reliability and ease of support. Those three things mean very much to me in my daily life."

- Richard Ewald
Senior Technical Architect
Sprint

The Oracle Exadata Database Machine runs the standard Oracle Database. Therefore, **any application that uses the Oracle Database today can be easily and seamlessly migrated to use the Exadata Database Machine with no changes to the application.** Databases can also be easily migrated off Exadata, eliminating "lock-in" concerns.

Customers thinking of deploying databases on the Public Cloud can be confident that **Exadata provides 100% compatibility between on-premises and public cloud**, enabling seamless migration to the public cloud as well as hybrid cloud deployments.

Extreme System Scalability and Growth with Elastic Configurations

The Exadata Database Machine uses a **scale-out architecture for both database servers and storage servers**. As workloads grow, database CPUs, storage, and networking are added to grow an Exadata Database Machine in a balanced fashion to scale without bottlenecks. The architecture allows elastic, seamless expansion from small to extremely large configurations to accommodate evolving workloads of any size.

A high-bandwidth low-latency 40 Gb/second **InfiniBand network** connects all the components inside an Exadata Database Machine. Specialized database networking protocols over InfiniBand provide **much lower latency and higher bandwidth** than is possible using generic communication protocols. This ensures both faster response time for OLTP operations, and higher throughput for analytic workloads. External connectivity to the Exadata Database Machine is via standard 10 Gigabit or 25 Gigabit Ethernet.

Exadata Database Machine is the most versatile database platform. The Exadata X8-2 Database Machine uses powerful database servers each with two 24-core x86 processors and 384 GB of memory (expandable up to 1.5 TB). Exadata also uses scale-out, intelligent storage servers available in two configurations – High Capacity (HC) or Extreme Flash (EF). HC Storage Servers have four NVMe PCI Flash cards each with 6.4 TB (raw) Exadata Smart Flash Cache and twelve 14 TB 7,200 RPM disks. EF Storage Servers have an all-Flash configuration with eight NVMe PCI Flash drives, each with 6.4 TB (raw) storage capacity. The minimal configuration of an Exadata Database Machine consists of two database servers and three storage servers, which **can be expanded into elastic configurations adding more database and/or storage servers within the same rack**. Elastic configurations provide a flexible and efficient mechanism to meet any size business need.

In addition to upgrading within a rack, **multiple racks can be connected using the integrated InfiniBand fabric** to form even larger configurations. For example, a system composed of four racks is simply four times as powerful as a single rack system: it provides quadruple the I/O throughput, quadruple the storage capacity, and quadruple the processing power. It can be configured as a large single system or logically partitioned for consolidation of multiple databases. Scaling out is easy, as Oracle Real Application Clusters (RAC) can dynamically add more processing power, and Automatic Storage Management (ASM) can dynamically add more storage capacity.

RELATED PRODUCTS

- Oracle Database Exadata Cloud Service
- Oracle Database Exadata Cloud at Customer
- Oracle Exadata Database Machine X8-8
- Oracle Exadata Storage Expansion Rack X8-2
- Oracle Exadata Storage Server X8-2 Plus InfiniBand Infrastructure
- Oracle Exadata Database Server X8-2 Plus InfiniBand Infrastructure
- Oracle SuperCluster
- Oracle Database 11g, 12c, 18c and 19c
- Real Application Clusters
- Partitioning
- Multitenant
- Database In-Memory
- Advanced Compression
- Advanced Security
- Active Data Guard
- GoldenGate
- Real Application Testing
- OLAP
- Advanced Analytics
- Business Intelligence
- Enterprise Manager
- Oracle Linux
- Oracle Virtual Machine

RELATED SERVICES

The following services are available from Oracle:

- Advanced Customer Services
- Oracle Premier Support for Systems
- Oracle Platinum Services
- Consulting Services
- Oracle University courses




Fig 1: Elastic Scale-out to Multi-rack Exadata

When even larger storage capacity is required, the **Oracle Exadata Storage Expansion Rack** is available. The Exadata Storage Expansion Rack enables customers to grow the storage capacity and bandwidth of any Exadata Database Machine. It is designed for database deployments with very large amounts of data, including historical or archive data, backups, documents, images, XML, JSON and LOBs. The storage expansion rack connects to the Exadata Database Machine using the integrated InfiniBand fabric and is very simple to configure with a few simple commands, as there are no LUNs or mount points. The starting configuration of the Oracle Exadata Storage Expansion Rack consists of four storage servers and can be expanded with additional storage servers.

Exadata Database Machines protect your investment as **newer generation servers and storage can be incorporated seamlessly into existing Exadata Database Machines**. Similarly, new software releases are compatible with most previous generation Exadata Database Machines.

Extreme Flash Storage Server: Record-breaking I/O Performance

Exadata **Extreme Flash (EF) Storage Server**, first introduced with Exadata X5, is the foundation of a database-optimized all-Flash Exadata Database Machine. Each EF Storage Server contains eight 6.4 TB state-of-the-art Flash Accelerator F640 NVMe PCI Flash drives, offering 51.2TB raw Flash capacity per EF Storage Server. Exadata X8 uses state-of-the-art Flash memory technology for improved speed, power efficiency, and endurance compared to previous generations of Flash. The enterprise grade Flash used in Exadata X8 have an expected endurance of 8 years or more for typical database workloads. This is very different from consumer-grade Flash that can potentially experience performance degradations or fail unexpectedly after a few years of usage. In addition, Exadata delivers ultra-high performance by placing the Flash devices directly

on the high speed PCI bus rather than behind slow disk controllers and directors. Finally, Exadata Flash uses state-of-the-art NVMe (Non-Volatile Memory Express) Flash to achieve extremely low latency and CPU overhead.

Flash performance is often limited and bottlenecked in traditional storage architectures. In contrast, Exadata uses a combination of scale-out storage, InfiniBand networking, database offload, and PCI Flash to deliver extremely high performance from Flash. A traditional full rack Exadata Database Machine X8-2, with 8 database servers and 14 Extreme Flash storage servers, can achieve up to **350 GB per second of analytic scan bandwidth from SQL**, and **0.25 ms Database I/O latency at 3.5 Million Flash IOPS** when running database workloads. A slightly different full rack combination, with 11 database servers and 11 Extreme Flash storage servers, can achieve up to **6.57 Million random 8K database read and 5.72 Million random 8K Flash write I/O operations per second (IOPS)**, which is an industry record for database workloads.

This performance is orders of magnitude faster than traditional storage array architectures, and it is also much faster than current all-Flash storage arrays. These are real-world end-to-end performance figures measured running SQL workloads with standard 8K database I/O sizes inside a single rack Exadata system, unlike storage vendor performance quotes that are usually based on small I/O sizes and low-level I/O tools and therefore are many times higher than can be achieved from SQL workloads.

“Oracle Exadata enabled seamless and rapid migration of our cloud-based, integrated business applications, required no major modifications, and helped us reduce costs. With Oracle, we dramatically improved our batch processing speed by 241x and overall system performance by 3x, enabling us to improve operating efficiency for our daily, critical tasks and to enhance customer service.”

- Kyoji Kato
Executive Officer and GM
Daiwa House




Fig 2: Flash Accelerator PCIe Card

“Oracle Exadata Database Machine is helping to transform our business. Our SAP environment, one of the world’s largest, can now support twice as much throughput with improved stability.”

- Milt Simonds
Director, Enterprise Platform Delivery
AmerisourceBergen Corporation

High Capacity Storage Server: Tiered Disk and Flash Deliver Cost of Disk with Performance of Flash

The second Exadata storage option is the Exadata X8-2 **High Capacity (HC) Storage Server**. This server includes twelve 14 TB SAS disk drives with 168 TB total raw disk capacity. It also has four Flash Accelerator F640 NVMe PCIe cards with a total raw capacity of 25.6 TB of Flash memory. Exadata Flash in a High Capacity Storage Server can be used directly as Flash disks, but is almost always configured as a Flash cache (**Exadata Smart Flash Cache**) in front of disk storage to deliver the best performance.

Exadata Smart Flash Cache automatically caches frequently accessed data while keeping infrequently accessed data on disk, delivering the high I/O rates and fast response times of Flash with the large capacity and low cost of disk. The Exadata Smart Flash Cache uniquely understands database workloads and **knows when to avoid caching data that the database will rarely access or is too big to fit in the cache.** For example, Exadata doesn't cache I/Os caused by backups, large table scans, or temporary results that will be quickly deleted. In addition to automatic caching, administrators can optionally provide SQL directives to ensure that specific tables, indexes, or partitions are preferentially retained in the Flash cache. A single full rack Exadata Database Machine X8-2, with 8 database servers and 14 High Capacity storage servers can achieve up to **350 GB per second of analytic scan bandwidth from SQL**, and up to **4.8 Million random 8K read I/O operations per second (IOPS) from SQL**, and **0.25 ms I/O latency at 2.7 Million Flash IOPS** when running database workloads.

It is common for hit rates in the Exadata Smart Flash Cache to be over 95%, or even 99% in real-world database workloads. Such high Flash cache hit rates mean that Exadata Smart Flash Cache provides an **effective Flash capacity** that is many times larger than the physical Flash. For example, a full rack Exadata Database Machine X8-2 with 8 database servers and 14 High Capacity Storage Servers often has an effective Flash capacity equal to the usable disk capacity of 700 TB.

The Exadata Smart Flash Cache also caches database block writes using Exadata Write Back Flash Cache technology. Write caching eliminates disk bottlenecks in large scale OLTP and batch workloads. The Flash write capacity of a single full rack Exadata Database Machine X8-2 with 8 database servers and 14 High Capacity Storage Servers exceeds **4.3 Million 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 automatic data tiering between RAM, Flash and disk implemented in Exadata provides tremendous advantages over other Flash-based solutions. Many storage vendors have recognized that the architecture of their traditional storage arrays inherently bottlenecks the performance of Flash and therefore have developed new Flash-only arrays. These Flash-only arrays deliver higher performance than traditional arrays but give up the cost advantages of smart tiering of data between disk and Flash, as the overall size of data that can benefit from Flash is limited to the size of expensive Flash. And these Flash arrays are unable to benefit from Exadata's unique storage optimization technologies. Data deduplication provided by some Flash arrays is very effective for Virtual Desktop Infrastructure (VDI) environments, but not for databases.

Exadata not only delivers much more capacity than generic all-Flash arrays, it also delivers better performance. Flash-only storage arrays cannot match the throughput of Exadata's integrated and optimized architecture with full InfiniBand based scale-out, fast PCI Flash, offload of data intensive operations to storage, and algorithms throughout that are specifically optimized for databases.

"We chose Oracle Exadata because it offered a complete solution ... we've created daily financial reports 4x faster and liquidity risk reports 7x faster to consistently meet our service-level agreement, improved credit risk management, and reduced our data center footprint."

- Vaibhav Samant
Senior Vice President, IT
HDFC Bank Ltd.

"None of the reports takes more than 10 minutes. It was taking 3-4 hours before, now it completes in 3 minutes. It sounds like unreal but it is real."

- Finance User
Turkcell

Extended Capacity Storage Server: Much Lower Cost Exadata Storage for Low Use Data

New for Exadata X8, a third Exadata storage option has been introduced - the Exadata X8-2 **Extended (XT) Storage Server**.

Each Exadata XT Storage Server includes twelve 14 TB SAS disk drives with 168 TB total raw disk capacity. To achieve a lower cost, Flash is not included, and storage software is optional.

This new storage option enables customers to extend the operational and management benefits of Exadata to rarely accessed data that must be kept online. This lower-cost addition to the Exadata Storage Server lineup delivers Exadata class benefits:

- **Efficient** – The XT server offers the same high capacity as the HC Storage server, including Hybrid Columnar Compression
- **Simple** – The XT server adds capacity to Exadata while remaining transparent to applications, transparent to SQL, and retains the same operational model
- **Secure** – The XT server enables customers to extend the same security model and encryption used for online data to low-use data, because it is integrated within the same Exadata
- **Fast and Scalable** – Unlike other low-access data storage solutions, the XT server is integrated to the Exadata fabric, for fast access and easy scale-out
- **Compatible** – The XT server is just another flavor of Exadata Storage server – you can just add XT servers to any Exadata rack

With Exadata X8-2 Extended (XT) Storage Server enterprises particularly in Finance, Insurance, and Telecommunications can meet their long-term data retention compliance requirements with the same trusted and continually validated Exadata solution, avoiding the operational risks and costs of managing information lifecycle on two or more platforms.

Accelerating Database Processing with Smart System Software

As data volumes continue their fast growth, conventional storage arrays struggle to quickly transfer data from disk and Flash to database servers at a rate that keeps the CPUs busy. Modern servers with dozens of CPU cores can consume data at many tens to hundreds of gigabytes a second. This is far faster than conventional storage arrays can deliver data through their storage controllers and the storage network.

What enables Exadata's unparalleled performance without any of the bottlenecks of traditional storage arrays is **Exadata System Software**. This software powers the Exadata storage servers, providing a **unique** highly efficient database-optimized storage infrastructure. Each Exadata Storage Server has two 16-core x86 processors that are used to offload database processing, and a rack of Exadata Database Machine can have a total of up to 576 processor cores in the storage servers able to offload the database servers. The CPUs in the storage servers do not replace database CPUs. Instead they

"The heart and soul of our stack right now is Oracle Exadata Database Machine. With Oracle Exadata, we've been able to reduce queries from days to minutes, and those that used to take minutes to seconds."

- Chris Wones
Enterprise Architect
Dunnhumby

"[With Exadata] We can more quickly process 65 billion daily transactions for data charging, while providing real-time information for customer inquiries, increasing customer satisfaction, and reducing costs."

- Jin Hyung Lee
ICT Team Manager, Networking
Engineering
SK Telecom

accelerate database intensive workloads similar to how graphics cards accelerate image intensive workloads.

One of the many unique features of Exadata System software is **Smart Scan** technology, which **offloads data intensive SQL operations from the database servers directly into the storage servers**. By pushing SQL processing to the storage servers, data filtering and processing occur immediately and in parallel across all storage servers, as data is read from disk and Flash. **Only the rows and columns that are directly relevant to a query are sent to the database servers.**

For example, if a query is executed to identify the customers who placed sales orders over \$1000 in the month of March, an Exadata system will offload the scanning of the table to the Exadata storage, filter out all sales orders that are less than \$1000, filter out sales orders not in March, and extract just the relevant customer information. This reduces the data transferred to the database servers by orders of magnitude. Thus Smart Scan greatly accelerates query execution, eliminates bottlenecks, and significantly reduces the CPU usage of the database servers.

“Exadata is the heart of the booking engine, and we cannot operate as a business, we cannot sell tickets without it.”

- James Callaghan
Chief Technologist
Westjet

Storage Index is another powerful **unique** capability of Oracle Exadata System software that helps avoid unnecessary I/O operations and improves overall performance. The storage index, maintained in-memory at the storage server, tracks summary information for table columns contained in a storage region on that storage server. When a query specifies a WHERE clause, Exadata System software examines the storage index using a Bloom filter to determine if rows with the specified column value might exist in a region of disk on the storage server. If the column value doesn't exist in the Bloom filter, then scan I/O in that region for that query is skipped. Storage Indexes make many SQL operations run dramatically faster because large numbers of I/O operations are automatically replaced by a few in-memory lookups.

Besides the intrinsic capabilities of Exadata System software, the combination of Oracle Database software, Exadata System software and Exadata infrastructure enables several additional unique capabilities that offer unparalleled performance levels for **OLTP** workloads. For example, **Exafusion Direct-to-Wire Protocol uniquely** allows database processes to read and send Oracle Real Applications Cluster (Oracle RAC) messages directly over the InfiniBand network using **Remote Direct Memory Access (RDMA)**, bypassing the OS kernel and networking software overhead. This improves the response time and scalability of Oracle RAC OLTP configurations on Oracle Exadata Database Machine, especially for workloads with high-contention updates.

In some OLTP workloads, more than half of remote reads are for Undo Blocks to satisfy read consistency. Exadata uniquely leverages ultra-fast **RDMA to read UNDO blocks** from other database instances, further improving OLTP performance.

The **Smart Fusion Block Transfer** capability **uniquely** improves performance of a RAC OLTP configuration further by eliminating the impact of redo log write latency, especially when hot blocks need to be transferred between sending and receiving nodes. The block is transferred as soon as the I/O to the redo log is issued at the sending node, without waiting for it to complete. Oracle internal tests show that Smart Block Transfer increases throughput (about 40% higher) and decreases response times (about 33% less) for communication-intensive workloads.

“Exadata delivers an amazing 20x compression for our Data Warehouse.”

- Jonathan Walsh
Head of BI & DW
Morrisons, Plc.

To further accelerate OLTP workloads, the Exadata Smart Flash Cache implements a unique algorithm to ensure low latency of database log writes called **Exadata Smart Flash Logging**. The time it takes 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 reduce the average latency of log writes and avoid the latency spikes that occur in other Flash solutions.

In addition, Exadata **uniquely** uses Machine Learning to implement **Automatic Indexing with Oracle Database 19c**. Automatic Indexing continually analyzes executing SQL and creates new indexes to accelerate performance. Automatic Indexing **continuously learns and tunes** the database as the underlying data model or usage patterns change.

Exadata also **uniquely** implements **Real Time Statistics** gathering as DML operations insert, update or delete data. Real Time Statistics allows the SQL optimizer to adapt plans dynamically as the distribution of data changes.

Optimizing Storage Use and I/O Through Compression

The Exadata Storage Server provides a unique compression capability called **Hybrid Columnar Compression (HCC) that enables dramatic reductions in storage for large databases**. Hybrid Columnar Compression technology is an innovative method of organizing data within a database table that uses a combination of both row and columnar methods for storing data. This hybrid approach achieves the compression benefits of columnar storage, while avoiding the performance shortfalls of a pure columnar format.

With Hybrid Columnar Compression, Exadata enables the highest levels of data compression possible with Oracle databases, and provides substantial cost-savings and performance improvements due to reduced I/O, especially for analytic workloads. Storage savings is data-dependent and often ranges from 5x to 20x. Average storage savings is an industry-leading 10x. On conventional systems, enabling high data compression has the drawback of reducing performance as it adds the load of decompression to the CPU. Because the Exadata Database Machine is able to offload decompression to processors in Exadata storage, and in addition there is reduced I/O need because of the high compression achieved, most analytic workloads run faster using Hybrid Columnar Compression than without it.

Two modes of Hybrid Columnar Compression are available. **Warehouse compression** mode is suitable for read-intensive workloads and provides large storage savings and enhanced analytic performance. **Archive compression** mode provides the highest degree of compression and is targeted at data that is seldom accessed but still must be kept online. In addition, this data can now be seamlessly stored on the XT storage server for further cost reduction.

On OLTP systems, Hybrid Columnar Compression can be used to compress older, less active data while newer, more active and update-intensive data can be compressed using Advanced Row Compression. Oracle Database Release 18c and above provides the ability to change the type of compression used by individual table partitions online

(even if there are global indexes on the table), to ensure seamless tiering across different compression types as data ages and becomes less active.

For data analytics, Exadata Smart Flash Cache implements a unique algorithm to accelerate reporting and analytical queries, called **Exadata Columnar Flash Cache**. Columnar Flash Caching implements a dual format architecture in Exadata Flash by automatically transforming frequently scanned Hybrid Columnar Compressed data into a pure columnar format as it is loaded into the Flash cache. Smart scans on pure columnar data in Flash run faster because they read only the selected columns, reducing Flash I/Os and storage server CPU consumption. This accelerates reporting and analytic queries while maintaining excellent performance for OLTP-style single row lookups.

Fault Tolerant and Fastest Database In-Memory Machine for Analytics and Mixed Workloads

Exadata is the ideal platform for running Oracle Database In-Memory. Oracle Database In-Memory on Exadata does not require all data to reside in memory. Data can be stored across multiple tiers of storage, with the hottest data in memory providing extremely high query performance, active data on Flash providing very high I/O throughput, and less active or older data on disk at a very low cost. **A single query can access data from all three tiers: memory, Flash and disk, completely transparently.** This allows Exadata to run faster, support higher capacities and deliver lower costs than competing products.

In addition, Exadata uniquely implements **In-Memory columnar formats in Flash Cache**. This feature extends the Exadata Columnar Flash Cache by automatically transforming data into In-Memory columnar formats as it is loaded into Flash cache. Smart Scans also process multiple column values with a single instruction by leveraging ultra-fast Single Instruction Multiple Data (SIMD) Vector instructions. Smart Scan results are passed back to the database server in Oracle Database In-Memory formats, further reducing the load on database server CPUs. The effect is to seamlessly extend the In-Memory columnar store size from DRAM capacity in the database server to Flash capacity in storage servers. An Exadata X8-2 Full Rack HC has 360 TB of Flash Cache, capable of servicing some of the largest in-memory workloads. Databases not using Oracle Database In-Memory still benefit from Exadata Columnar Flash Cache without the vector processing optimizations.

Exadata uniquely implements **Fault Tolerant memory duplication for Oracle Database In-Memory**. On a generic cluster configuration, when a server node fails, the in-memory data on that node is lost, and it takes many minutes to repopulate the in-memory data on a surviving node. During this time, analytic queries will run orders of magnitude slower. This means generic platforms will fail to meet business SLAs. However, on Exadata, Fault-Tolerant memory duplication can eliminate this slowdown by duplicating any subset of the in-memory data across the clustered database servers. If a database server fails, queries will transparently access the duplicate copy on a surviving database server and continue without interruption.

Exadata uniquely integrates with **Active Data Guard** to allow customers to run In-Memory analytics on a standby database, further improving the return on investment of the standby system, and enhancing availability and overall performance.

In-Memory Acceleration for OLTP and Consolidation

For OLTP workloads Exadata **uniquely** implements **In-Memory OLTP Acceleration**. This feature utilizes the memory installed in Exadata Storage Servers as an extension of the memory cache (buffer cache) on database servers. Specialized algorithms transfer data between the cache on database servers and the in-memory cache on storage servers. This reduces the I/O latency to 100 μ s for all I/Os served from in-memory cache. Exadata's **uniquely** keeps only one in-memory copy of data across database and storage servers, avoiding memory waste from caching the same block multiple times. This greatly improves both efficiency and capacity and is only possible because of Exadata's unique end-to-end integration.

Enhanced Consolidation using Oracle Virtual Machines

Consolidated environments running on Exadata can use Xen-based Oracle Virtual Machines (OVMs) for a high degree of isolation between workloads, especially in cloud, hosted, shared, service provider, and test/dev environments. Using OVM, multiple RAC clusters can be deployed on the same set of database servers of the Exadata Database Machine, for consolidation of applications with specific Clusterware version needs.

Exadata Database Machine is the world's fastest virtualized database platform.

Exadata virtual machines use high speed InfiniBand networking with Single Root I/O Virtualization (SR-IOV) to ensure that performance within a virtual machine is similar to Exadata's excellent raw hardware performance. Exadata Smart Scans greatly decrease virtualization overhead compared to other platforms by dramatically reducing message traffic between virtual machines. Exadata virtual machines can dynamically expand or shrink the use of CPUs based on the workload requirement of the applications running in that virtual machine.

Virtual machines on Exadata are considered Trusted Partitions and therefore software can be licensed at the virtual machine level instead of the physical processor level. Without Trusted Partitions, database options and other Oracle software must be licensed at a server or cluster level even though not all databases running on that server or cluster may require a particular option.

Enterprise-Class Security with Extreme Performance

Exadata Database Machine is the world's most secure database machine. Building on the high security capabilities in the Oracle Database such as Transparent Data Encryption (TDE), Exadata **uniquely moves decryption processing from database server software into the Exadata Storage Server hardware**. Exadata storage leverages hardware decryption and compression together to provide the highest performance secure databases. **Encryption occurs after the data is compressed so that the cost of decryption is decreased by the degree of compression**. By leveraging both technologies, Exadata is able to query fully encrypted and compressed databases with minimal overhead at a rate of hundreds of gigabytes of (original) user data per second. In addition, TDE provides a complete key management solution to keep all data encrypted and secure.

"I don't get the calls in the middle of the night anymore that we have a system down. Exadata is taken as always being available."

- James Callaghan
Chief Technologist
Westjet

"Our critical electronic payments service has been live on Exadata since early 2011 with 100% uptime. The service reliably processes the transfer of billions of Euros per week and achieves subsecond response times for online enquiries."

Martin McGeough
Database Technical Architect
Vocalink

Exadata is designed and delivered as an integrated whole, instead of a collection of components. In traditional database deployments, the customer takes on all the system integration tasks, including the task of ensuring the security of each individual software and hardware component, and ensuring that security is maintained across the full product stack. **Oracle delivers full stack security in the Exadata Database Machine.** Exadata virtual machines provide an added layer of isolation at the operating system level. Additionally, on both physical and virtual deployments, Exadata systems use minimal Oracle Linux distributions to ensure that just the RPMs required to run the Oracle Database are installed and enabled. With this approach, system security is much stronger than default Linux installations and many security vulnerabilities are avoided. In addition, Exadata leverages the ksplice capability of Oracle Linux to apply security updates while the OS stays online.

Exadata security begins at power-up time with Secure Boot, which ensures that the system UEFI firmware only allows the execution of cryptographically signed boot loaders that the system recognizes as trustworthy. With each reboot of the server, every executed component is verified. This prevents malware from hiding embedded code in the boot chain.

In addition, the disk and Flash technologies used in Exadata X8 enable Stored Data Encryption. In Stored Data Encryption, disk and Flash storage devices encrypt all user data as it enters the devices. Exadata's Secure Erase feature leverages this capability when an Exadata is re-purposed or decommissioned to instantly erase all user data present on storage devices by changing the encryption keys used to encrypt the user data. With Secure Erase, because the previous encryption key is deleted, there is no need to worry about latent data left on storage devices due to over-provisioning or sector sparing.

Exadata security has been probed and evaluated by hundreds of leading banks, telecoms, and government organizations worldwide. The security findings of all these evaluations have been incorporated into the Exadata standard configuration. Therefore Exadata benefits from scrutiny both by Oracle Security experts and by hundreds of industry security experts around the world.

Mission Critical High Availability

The Exadata Database Machine is engineered to provide the highest levels of availability. **All types of failures are protected against** from simple failures such as disk, server, or network, to complex site failures and human errors. Each Exadata Database Machine has **completely redundant hardware**, including redundant InfiniBand networking, redundant Power Distribution Units (PDU), redundant power supplies, and redundant database and storage servers. Oracle RAC protects against database server failure. Oracle ASM provides data mirroring to protect against disk or storage server failures. Oracle RMAN provides extremely fast and efficient backups to disk or tape. Oracle's Flashback technology allows backing out user errors at the database, table or even row level. Using Oracle Data Guard, a second Exadata Database Machine can be deployed in a Maximum Availability Architecture (MAA) configuration to transparently maintain a real-time copy of the database at a remote site and provide full protection against primary database failures and site disasters.

"By consolidating 350 database servers and storage systems onto Oracle Exadata, we gained a high-performance, reliable, and scalable mobile billing platform, enabling us to calculate billings data 10x faster, and halve maintenance costs."

- Tomoki Shimamura
Senior Manager Billing Systems Group
NTT DoCoMo, Inc.

"By integrating 20 legacy database servers for our investment trust sales system into four Oracle Exadata Database Machines, we can provide information to customers 136x faster, enhance our competitive advantage, and support transaction growth for the next 10 years at lower costs."

- Tomoshiro Takemoto
Senior Managing Director
Cloud Computing Service Division
Nomura Research Institute Ltd.

Exadata in an MAA configuration is recognized by the analyst firm IDC as a system that **delivers at least 5-nines availability** and is categorized in the IDC AL4 fault-tolerant market segment, along with HP Integrity NonStop and IBM z Systems¹.

The Exadata principle of deep hardware and software integration is also evident in the many ways Exadata **uniquely** assures high availability across several different failure conditions. One such **unique** capability is **Instant Detection of Compute and Storage Server Failures**. On non-Exadata platforms, detecting a server failure requires waiting for a long timeout, leading to extended application brownouts. Exadata leverages InfiniBand integration to quickly determine that a server is not reachable through any network path enabling Exadata to immediately initiate eviction of the failed server from the cluster. This entire operation can be completed in less than 2 seconds, leading to virtual elimination of application brownout conditions.

Disk and Flash devices occasionally exhibit very long latency I/O operations due to internal recovery of failed sectors, internal firmware reboots, or wear leveling. These long I/O operations can cause stalls in mission critical OLTP databases. With Exadata's **unique I/O Latency Capping**, Oracle Exadata System software automatically redirects read I/O operations to an ASM-mirrored copy of the data when the latency of a read I/O is much longer than expected. Similarly, it automatically redirects high latency write I/O operations to a healthy Flash device, eliminating write outliers. Exadata System Software uses Machine Learning techniques to predict components susceptible to failure and takes proactive action to gracefully take such components out of service. If disks do fail, ASM performs a rebalance operation for the data that was resident on the disk. Exadata allows hot swapping of disks, fans, power supplies, and PCIe Flash cards to avoid downtime. Exadata System software takes rebalance one step further by preserving the Flash cache population and storage indexes when moving data between storage servers to maintain consistent application performance. On rare occasions when there are outliers within the networking subsystem, Exadata redirects the I/O issued by the database server to another storage server.

Exadata automates monitoring of CPU, memory, Input/Output, file system, and network. This automation combines machine learning techniques with the deep lessons learned from thousands of mission critical real-world deployments. For example, Exadata can detect that anomalous use of system resources is affecting database performance, identify the process responsible, and issue an alert, without any pre-existing set-up.

Because of its industry leading availability, the Exadata Database Machine has been deployed by leading companies for their most critical applications including interbank fund transfers, online securities trading, real-time call tracking, and web-based retailing. Exadata's Mission Critical availability capabilities are not restricted to OLTP workloads; they also apply to warehousing and analytics.

Ideal Platform for Database as a Service

The Exadata Database Machine can host many databases, enabling massive database consolidation or a sophisticated Database as a Service private cloud. Multi-database environments inherently have diverse, complex, and unpredictable workloads mixing

"With Oracle Platinum Services, we achieved a 40% increase in operational efficiency - virtually eliminating the need for dedicated resources to perform updates, monitor systems, and resolve incidents."

- Roland Schiller
Architect
Swiss Re

¹ Worldwide Fault-Tolerant Servers Market Shares, 2014: *Vendors Are Hearing the Customer — More Bold Moves Needed to Grow the Segment*, IDC, Peter Rutten, Lloyd Cohen, October 2015

OLTP, analytics, and batch operations with sequential and random access patterns. Exadata's ability to **run any type or mix of database workloads with industry leading scalability and performance** makes it an ideal consolidation platform – whether for multi-database workloads, or for pluggable databases with Oracle Multitenant in Oracle Database 12c, 18c and Oracle Database 19c.

Multi-database environments create an inherent risk that one database will consume too many resources and therefore impact the quality of service of other databases. The Exadata Database Machine **uniquely** provides **end-to-end prioritization** from the application to database CPUs, network, and storage. Priorities and resource limits can be specified at the physical database, pluggable database, connection, application, user, or even job level to ensure that each of the consolidated databases or SQL operations receives the necessary resources and achieves the target response times.

Exadata **uniquely** implements **database and I/O resource management**. Fine-grained priorities specified for operations at the database level are automatically communicated to Exadata Storage Servers and applied to each I/O operation to ensure that prioritization of database operations applies to both CPU operations and I/O operations. The same resource management principles can also be applied when multiple databases are deployed on one Exadata rack, as is typical in a consolidated private cloud.

Exadata also **uniquely** implements **database network resource management** to ensure that network intensive workloads such as reporting, batch, and backups don't stall response time sensitive interactive workloads. Latency sensitive network operations such as RAC Cache Fusion communication and log file writes are automatically moved to the head of the message queue in server and storage network cards as well as InfiniBand network switches, bypassing any non-latency sensitive messages. Latency critical messages even jump ahead of non-latency critical messages that have already been partially sent across the network, ensuring low response times even in the presence of large network DMA (Direct Memory Access) operations.

Because of Exadata's unique Consolidation and Database as a Service capabilities, Exadata is the **only** platform that can support up to 4000 Pluggable Databases within a single Oracle Multitenant Container Database.

Fast Deployment of Development and Test Databases with Exadata Snapshots and Containers

Space-efficient database snapshots can be quickly created for test and development purposes directly on Exadata. Exadata database snapshots are integrated with Oracle Multitenant to provide an extremely simple interface for creating new pluggable database (PDB) snapshots.

Snapshots start with a shared read-only copy of the production database (or PDB) that has been cleansed of sensitive information. A hierarchy of read-write snapshots can be created from this shared copy. As changes are made, each snapshot writes the changed blocks to a sparse disk group. Since multiple users can create independent snapshots from the same base database copy, multiple test and development environments can share space while maintaining independent databases for each task.

"We've seen system performance increase 4x and the total cost of storage ownership fall by 30%, even as our data volumes have surged."

- Holger Haun
Managing Director
IDS GmbH – Analysis and Reporting
Service

"On an annual basis Exadata provides at least half a million dollars in savings in operating costs."

- James Callaghan
Chief Technologist
Westjet

All Exadata specific features such as Smart Scan, resource management and Smart Flash Cache work seamlessly on database instances created via Exadata snapshots, hence providing an exact test and development environment while using a fraction of valuable storage resources. Backups of snapshots on Exadata are also space efficient as only the changed information is backed up.

In addition Exadata supports Docker containers, enabling customers to easily and rapidly provision databases for testing and development to foster agile development. Customer can also use Docker containers to host lightweight applications and agents.

Comprehensive System Management

Oracle Enterprise Manager uses a holistic approach to manage the Exadata Database Machine and provides comprehensive capabilities from monitoring and reporting to active lifecycle management. It enables:

- Unified Monitoring: Oracle Enterprise Manager 13c **uniquely** supports a single pane of glass view of all the hardware and software components such as database servers, storage servers, InfiniBand switches, and monitors the operations running on them and their resource utilization. DBAs can drill down from database monitoring screens to the Exadata storage layer to quickly determine root causes of any performance bottlenecks.
- Lights-out monitoring within Enterprise Manager is optimized for Exadata with predefined metrics and thresholds so that administrators receive timely notifications when issues arise, and manage those exceptions. In addition, hardware incidents are automatically detected and service requests logged to reduce problem resolution time.
- The Exachk tool, which is integrated with Enterprise Manager's powerful compliance framework, provides functionality for system administrators to automate the assessment of Engineered Systems for known configuration problems and best practices. Administrators can leverage the Consistency Check functionality to check for deviations in configuration across the racks or among the database servers of a rack.
- Exadata's built in Management Server (MS) processes constantly monitor the health of hardware and software components, and send alerts to both administrators and Oracle support when faulty components are detected.

Highest Level of Service

Oracle offers a complete set of support services for the Exadata family of products including: 24x7 hardware service, system monitoring, software installation and configuration among other standard and custom offerings.

Of particular value is **Oracle Platinum Services** that is available exclusively for Oracle's Engineered Systems. Platinum Services provides fault monitoring, faster response times, and expedited escalation to development. With Platinum Services, Oracle support engineers perform software maintenance and patching remotely. Platinum Services provides a higher level of support than has ever been available before for all software and hardware within an Engineered System including the Oracle Database. Platinum Services is provided at no extra charge to Exadata customers.

IT Agility

Exadata is a complete system for running databases including storage, servers, and internal networks. Management of a traditional database system is typically spread across the management teams of each of the components such as the database team, the storage team, and the system administration team. In contrast, an **Exadata system is typically managed by a unified Database Machine Administration (DMA) team.** Database Machine Administrators have full control of all resources in the Exadata Database Machine including storage resources. New database deployments and configuration changes can be implemented by the Database Machine Administrators without coordination across different component management teams that are often overloaded and have differing priorities. Database Machine Administrators can focus on application and business specific enhancements rather than coordinating across component teams, or tuning and triaging of low level configuration issues.

Dramatically Lower Costs

Because of the extreme performance, high storage capacity, and unique compression capabilities delivered by the Exadata Database Machine, workloads that would require very large traditional hardware systems can be run on much smaller Exadata systems. The hardware needed for an application deployed on an Exadata system is often reduced 2-4X compared to a traditional system.

Exadata provides a huge RAM, Flash, and disk footprint for large data sets. Raw disk storage on an Exadata full rack can exceed 3 Petabytes while raw Flash storage can be up to 920 TB. In addition, Hybrid Columnar Compression often expands storage and memory capacity 10X. By intelligently moving active data across disk, Flash, and memory tiers, Exadata simultaneously delivers the highest performance and the lowest cost.

Exadata has the **unique** ability to consolidate many databases supporting multiple workloads in a single cloud platform. High-end OLTP, analytics, batch, reporting, and backups can all run simultaneously within and across databases with extreme performance. **The extreme performance and capacity of Exadata enables very large numbers of databases and workloads to be consolidated on Exadata.** Consolidating databases on Exadata reduces system hardware cost, software cost, and greatly reduces ongoing operations cost.

The uniformity of Exadata Database Machine configurations results in large cost savings. **Exadata standardizes not just technologies, but also integration, testing, security, hardening, tuning, and support.** Customers deploy Exadata systems much faster and with a lot less labor than traditional systems. Low level tuning, integration, and maintenance is reduced or eliminated. Because all Exadata users run a configuration that is identical to thousands of other users, and is identical to Oracle's internal configurations, it is far less likely that issues will be encountered, and issue resolution is quicker and simpler reducing both operations cost and downtime cost.

Capacity-on-Demand Software Licensing

An X8-2 database server has a substantial amount of compute capacity with two 24-core x86 processors (48 cores in total). The Capacity-on-Demand feature allows a number of cores per database server to be turned off during the hardware installation, leaving at least 14 cores enabled. As your workload grows and more cores are needed, Capacity-on-Demand can be used to re-enable cores and license software 2 cores at a time. This pay-as-you-grow approach to software licensing is another way in which Exadata helps to align costs with business growth.


Exadata in Oracle Public Cloud

With Oracle Database Exadata Cloud Service, customers can now run Oracle databases in the cloud on Exadata, with the same extreme performance and availability experienced by thousands of organizations deploying Exadata on-premises. Exadata Cloud Service combines the world's #1 database – Oracle, and the most powerful database platform – Exadata, with all the simplicity and cost effectiveness of the public cloud.

Oracle databases deployed in the cloud as part of this service include **all Oracle Database options and features** – such as Oracle Multitenant, In-Memory Database, Real Application Clusters (RAC), Active Data Guard, Partitioning, Advanced Compression, Advanced Security, etc. It also includes all Oracle Enterprise Manager packs. Oracle databases deployed in the Exadata Cloud Service are **100% compatible** with those deployed on-premises, ensuring a smooth transition to the cloud, and an efficient hybrid cloud strategy. With pay-as you-grow dedicated Exadata configurations, and infrastructure managed by Oracle experts, Exadata Cloud Service enables business agility and operational flexibility with **zero CapEx**.

To commence an Exadata Cloud Service, customers log in through **Oracle Cloud Portal** (<https://cloud.oracle.com>) and activate the Exadata configuration allocated and dedicated for them. They can start with a Quarter Rack with 2 database servers and 3 storage servers. The database servers have a minimum and maximum number of compute cores (OCPUs) that can be enabled for the chosen configuration, and customers can specify their desired number of compute cores within these limits. Customers subscribe to the Exadata service for a specified term, and the pricing is based on the number of enabled compute cores for that term. As business grows, customers can enable additional compute cores, or opt for larger Exadata configurations, thus paying only for the processing power that is required. Compute cores can be expanded and shrunk online to immediately adapt to peak or unexpected workloads.

Once the Exadata Service is activated, customers can start creating databases on the allocated Exadata system. Customers have root privilege for the Exadata database server VMs and DBA password on the Oracle Databases, so they can configure the system and create/deploy databases as they like, and load additional agent software on the Exadata database servers to conform to business standards or security monitoring requirements. They can connect to the configured databases using standard Oracle Net Services clients such as JDBC and OCI. Customers can also deploy their applications



within Oracle Public Cloud infrastructure, using cloud offerings such as Oracle Cloud Infrastructure or Oracle Java Cloud Service.

Customers perform their own database administration tasks, using the familiar Oracle Enterprise Manager interface. All supporting infrastructure for Exadata Cloud Service is deployed, maintained and managed by Oracle, including datacenter networking, private Exadata InfiniBand networks, physical Exadata database and storage servers, firmware, and Exadata storage software. This allows customers to focus on their business needs and application requirements, and not IT infrastructure management.

Exadata Cloud Service is an ideal fit for:

- Running business-critical production OLTP or analytic databases at almost any scale without incurring the capital expenditure and complexity of maintaining the underlying IT infrastructure.
- Consolidating a variety of workloads in the Cloud using multiple Oracle Databases or Oracle Multitenant.
- Maintaining synchronized Oracle standby or replica databases for disaster recovery in the Cloud and/or query offloading using Oracle Active Data Guard or Oracle GoldenGate.
- Quickly provisioning high-performance Oracle databases for ad-hoc business reasons such as feature development, functionality testing, application certification, proof-of-concept, try-before-buy, Database 12c sandbox testing, etc.
- Executing time-sensitive business applications such as launching a web-based marketing campaign, running loyalty programs, rolling out new business initiatives, etc.

An attractive aspect of all these use cases for existing Oracle Database customers is that their applications and data models **do not have to change**. Their data center simply expands to include the elasticity and flexibility of the Oracle Cloud. They also do not have to invest in multiple cloud platforms for multiple workloads since Exadata provides a unified platform for all workloads – analytics, data warehousing, OLTP, consolidation, in-memory and mixed-workloads.

With a database platform **uniquely** engineered for extreme performance for all workloads, along with fast deployment, simplified management, low operating costs and reduced risks, Exadata Cloud Service is the best Cloud Database platform available today.

Customers that are not able to move their databases to the Public Cloud can get the Oracle Public Cloud implemented in their own data center using the Oracle Database Exadata Cloud at Customer. Exadata Cloud at Customer:

- Is paid for using subscription based pricing that includes all Database and Exadata software, all Exadata hardware, and remote operation and infrastructure management by Oracle.
- Includes the same Web and REST based DB and OS provisioning, management, and orchestration as the Oracle Public Cloud.
- The Oracle Cloud Operations Team Manages the Exadata Infrastructure including Servers, storage, storage software, networking, firmware, hypervisor, etc.

- Exadata is the only platform suitable for deploying Mission Critical databases using Cloud at Customer.

Exadata Business Benefits

Beyond the operational benefits of extreme performance, availability, and security, and deployment flexibilities across on-premises and Cloud, Exadata also directly benefits the business.

Exadata accelerates time to market for new business applications since the time needed for system configuration, tuning, and testing is largely eliminated. Deployment times are reduced from months to days, and the risk of unexpected system level issues after go-live is greatly reduced. When a new application is deployed, it is common for unanticipated application usage patterns to create performance issues. Exadata's huge I/O, network, and compute throughput can absorb spikes created by unanticipated workloads without slowing response times of mission critical workloads. Overall Exadata speeds application deployment and reduces risk, allowing businesses to innovate faster.

Exadata's extreme performance and large memory and Flash capacity enhance employee productivity and customer satisfaction by greatly improving user response times. **Users spend more time doing useful work, and less time waiting** for the system to respond.

Exadata's extreme performance does not just improve business efficiency, it also **enables business users to make smarter decisions, discover growth opportunities, and reduce costs**. Users can analyze data in real-time, explore different possibilities, and perform rapid iteration to find better solutions. Exadata enables:

- Real-time business data analysis
- Faster financial closes
- Better planning and budgeting
- More effective and faster projections

Conclusion

Exadata delivers a fully integrated database platform with the latest hardware technologies and **unique** software to deliver extreme performance, availability, and security. This coupled with cost savings, ease of management, and enhanced supportability result in greater business agility and efficiency. Given what can be achieved with Exadata, it is no surprise it is the new global standard for running Oracle Databases – **whether on-premises, or in the Oracle Public Cloud**.

EXADATA SERVER HARDWARE¹

Server Type	CPU	Memory	Disk	Flash	Network
Database Server	2 x 24-core Intel Xeon 8260 processors (2.4GHz)	384 GB (default) to 1.5 TB (max)	4x 1.2 TB 10,000 RPM disks (Hot-Swappable)	None	2x 10 Gb copper Ethernet ports (client) OR 2 x10/25 Gb optical Ethernet port 1x 1/10 Gb copper Ethernet port (mgmt) 2x 10/25 Gb optical Ethernet ports (client) 2x QDR (40 Gb) InfiniBand ports 1x ILOM Ethernet port 4x 10 Gb copper (client – optional) or 2x 10/25 Gb optical Ethernet (client – optional)
Storage Server HC	2 x 16-core Intel Xeon 5218 processors (2.3GHz)	192 GB	12x 14 TB 7,200 RPM disks	4x 6.4 TB NVMe PCIe 3.0 Flash cards	2x QDR (40 Gb) InfiniBand ports 1x ILOM Ethernet port
Storage Server EF		192 GB	None	8x 6.4 TB NVMe PCIe 3.0 Flash cards	
Storage Server XT	1 x 16-core Intel 5218 Xeon processor (2.3GHz)	96 GB	12x 14 TB 7,200 RPM disks	None	
Eighth Rack Storage Server HC	2 x 16-core Intel Xeon 5218 processors (2.3GHz) 16 cores enabled	192 GB	6x 14 TB 7,200 RPM disks	2x 6.4 TB NVMe PCIe 3.0 Flash cards	2x QDR (40 Gb) InfiniBand ports 1x ILOM Ethernet port

¹ All servers include redundant hot swappable fans and power supplies

EXADATA TYPICAL RACK CONFIGURATIONS²

Rack Size	Database Servers and Cores	Storage Servers and Cores	HC Storage Capacity (raw)	EF Storage Capacity (raw)
Eighth Rack ³	2x servers, 48 cores	3x servers, 48 cores for SQL offload	252 TB disk, 38.4 TB Flash	or 76.8 TB Flash 153.6 TB Flash 358.4 TB Flash 716.8 TB Flash
Quarter Rack	2x servers, 96 cores	3x servers, 96 cores for SQL offload	504 TB disk, 76.8 TB Flash	
Half Rack	4x servers, 192 cores	7x servers, 224 cores for SQL offload	1176 TB disk, 179.2 TB Flash	
Full Rack	8x servers, 384 cores	14x servers, 448 cores for SQL offload	2,352 TB disk, 358.4 TB Flash	

² Each rack is 42 RU (Rack Units) in height, has 2x redundant Power Distribution Units (PDUs), 2x 36-port QDR (40 Gb/s) InfiniBand switches and 1x 48-port Cisco Ethernet switch for administration. Included Spare Parts Kit Contains:

- 1 x 6.4 TB NVMe PCI Flash card and 1 x 14 TB High Capacity disk, or
- 1 x 6.4 TB NVMe PCI Flash card

³ Eighth Rack is the minimum Exadata configuration. Eighth Rack database servers have one processor each with all cores enabled. Maximum memory supported per database server is 768GB. Eighth Rack EF storage servers have half the cores and Flash drives enabled. Eighth Rack HC storage servers have half the cores enabled and half the disks and Flash cards removed. The optional NIC is not available for the Eighth Rack.

EXADATA ELASTIC CONFIGURATIONS⁴

Rack Size	Database Servers and Cores	Storage Servers and Cores	HC Storage Capacity (raw)	EF Storage Capacity (raw)
Quarter Rack	2x servers, 96 cores	3x servers, 96 cores for SQL offload	504 TB disk, 76.8 TB Flash	or 153.6 TB Flash n/a 921.6 TB Flash max per rack
+ Database Servers	Up to 19x servers ⁵ , 912 cores max per rack	n/a	n/a	
+ Storage Servers	n/a	Up to 18x servers ⁵ , 576 cores max per rack	3,024 TB disk, 460.8 TB Flash max per rack	

⁴ Elastic configurations allow adding database or storage servers to a quarter rack to achieve the exact ratio of compute to storage that the application needs. A full rack elastic configuration cannot exceed 22 servers and 39 RU (Rack Units). Database Servers = 1 RU, Storage Servers = 2 RU

⁵ Maximum number of database servers allowed in an elastic configuration is 19. Maximum number of storage servers allowed in an elastic configuration is 18.

OTHER ELASTIC EXPANSION OPTIONS

Multi-Rack Connection	Connect any combination of up to 18 Exadata Database Machine racks or Exadata Storage Expansion Racks via the InfiniBand fabric. Larger configurations can be built with external InfiniBand switches. Connected racks can be any combination of X2, X3, X4, X5, X6, X7 or X8 hardware.
Eighth Rack Expansion Options	<p>. Expand just compute or just storage or both, described as follows:</p> <ul style="list-style-type: none"> • Database Servers can be expanded by installing an additional 24 core CPU per server • EF Storage Servers can be expanded by enabling 10 cores and installing 4 PCI Flash cards per server • HC Storage Servers can be expanded by adding Eighth Rack High Capacity Storage Server

EXADATA CAPACITY AND PERFORMANCE METRICS: INDIVIDUAL SERVERS

Server Type	Maximum SQL Read Bandwidth ²	Maximum SQL Read IOPS ³	Maximum SQL Write IOPS ⁴	PCI Flash Capacity (raw) ⁵	Disk Data Capacity(raw)
Database Server	N/A	597,500	544,000	N/A	4.8 TB
Storage Server HC ¹	25 GB/s	597,000	520,000	25.6 TB	168 TB
Eighth Rack Storage Server HC ¹	12.5 GB/s	298,500	260,000	12.8 TB	84 TB
Storage Server EF ¹	40 GB/s	597,000	520,000	51.2 TB	N/A
Storage Server XT ¹	1.8 GB/s	2,600	2,600	N/A	168 TB

EXADATA TYPICAL RACK CONFIGURATIONS: FLASH CAPACITY AND PERFORMANCE METRICS (HC & EF)

Flash Metrics		Maximum SQL Flash Bandwidth	Maximum SQL Flash Read IOPS	Maximum SQL Flash Write IOPS	PCI Flash Capacity (raw)
Full Rack	HC	350 GB/s	4,776,000	4,352,000	358.4TB
	EF	560 GB/s	4,776,000	4,352,000	716.8 TB
Half Rack	HC	175 GB/s	2,388,000	2,176,000	179.2 TB
	EF	280 GB/s	2,388,000	2,176,000	358.4 TB
Quarter Rack	HC	75 GB/s	1,194,000	1,088,000	76.8 TB
	EF	120 GB/s	1,194,000	1,088,000	153.6 TB
Eighth Rack	HC	38 GB/s	597,000	544,000	38.4 TB
	EF	60 GB/s	597,000	544,000	76.8 TB

EXADATA TYPICAL RACK CONFIGURATIONS: DISK CAPACITY AND PERFORMANCE METRICS (HC)

Disk Metrics	Maximum SQL Disk Bandwidth	Maximum SQL Disk IOPS	Data Capacity (raw)
Full Rack	25 GB/s	36,000	2,352 TB
Half Rack	12.5 GB/s	18,000	1,176 TB
Quarter Rack	5.4 GB/s	7,800	504 TB
Eighth Rack	2.7 GB/s	3,900	252 TB

EXADATA TYPICAL RACK CONFIGURATIONS: COMBINED METRICS (HC & EF)

Combined Metrics		Data Capacity (Usable) – Normal Redundancy ⁶	Data Capacity (Usable) – High Redundancy ⁶	Maximum Data Load Rate ⁷
Full Rack	HC	953 TB	699 TB	35 TB/hour
	EF	282 TB	206 TB	35 TB/hour
Half Rack	HC	477 TB	349 TB	17.5 TB/hour
	EF	141 TB	103 TB	17.5 TB/hour
Quarter Rack	HC	191 TB	150 TB	7.5 TB/hour
	EF	56 TB	44 TB	7.5 TB/hour
Eighth Rack	HC	95 TB	75 TB	3.8 TB/hour
	EF	28 TB	22 TB	3.8 TB/hour

¹HC = High Capacity; EF = Extreme Flash; XT = Extended. Actual system performance varies by application.

²Bandwidth is peak physical scan bandwidth achieved running SQL, assuming no database compression. Effective user data bandwidth is higher when database compression is used.

³Based on 8K I/O requests running SQL. Note that the I/O size greatly affects Flash IOPS. Other products quote IOPS based on smaller I/Os that are not relevant for databases.

⁴Based on 8K I/O requests running SQL. Flash write I/Os measured at the storage servers after ASM mirroring, which usually issues multiple storage I/Os to maintain redundancy.

⁵Raw capacity is measured in standard disk drive terminology with 1 GB = 1 billion bytes.

⁶Usable capacity is measured using normal powers of 2 space terminology with 1 TB = 1024 * 1024 * 1024 bytes. It is the actual space available to create a database after taking into account space needed for ASM redundancy, recovering from a drive failure. Normal redundancy calculations reflect the use of Grid Infrastructure version 12.2.0.1 or later.

⁷Load rates are typically limited by database server CPU, not I/O. Rates vary based on load method, indexes, data types, compression, and partitioning.

EXADATA DATABASE MACHINE COMPONENT ENVIRONMENTAL SPECIFICATIONS

Metric	Database Server X8-2 Plus InfiniBand Infrastructure	Storage Server X8-2 High Capacity Plus InfiniBand Infrastructure	Eighth Rack Storage Server X8-2 High Capacity Plus InfiniBand Infrastructure	Storage Server X8-2 Extreme Flash Plus InfiniBand Infrastructure	Storage Server X8-2 Extended Plus InfiniBand Infrastructure
Height	1.7 in. (42.6 mm)	3.4 in. (86.9 mm)			
Width	17.2 in. (436.5 mm)	17.5 in. (445.0 mm)			
Depth	29.0 in. (737.0 mm)	29.9 in. (759.0 mm)			
Acoustic noise (operating)	7.7 B	8.2 B	8.2 B	8.2 B	8.2 B
Weight	45.6 lbs (20.7 kg)	76.7 lbs (34.8 kg)	67.5 lbs (30.6 kg)	60.6 lbs (27.5 kg)	66.7 lbs (30.2 kg)
Maximum power usage	0.7 kW (0.7 kVA)	0.7 kW (0.7 kVA)	0.5 kW (0.5 kVA)	0.7 kW (0.7 kVA)	0.5 kW (0.5 kVA)
Typical power usage ¹	0.5 kW (0.5 kVA)	0.5 kW (0.5 kVA)	0.4 kW (0.4 kVA)	0.5 kW (0.5 kVA)	0.3 kW (0.3 kVA)
Cooling at maximum usage	2,409 BTU/hour (2,541 kJ/hour)	2,317 BTU/hour (2,444 kJ/hour)	1,714 BTU/hour (1,809 kJ/hour)	2,283 BTU/hour (2,408 kJ/hour)	1,570 BTU/hour (1,656 kJ/hour)
Cooling at typical usage	1,686 BTU/hour (1,779 kJ/hour)	1,622 BTU/hour (1,711 kJ/hour)	1,200 BTU/hour (1,266 kJ/hour)	1,598 BTU/hour (1,686 kJ/hour)	1,099 BTU/hour (1,159 kJ/hour)
Airflow at maximum usage ²	112 CFM	107 CFM	79 CFM	106 CFM	73 CFM
Airflow at typical usage ²	78 CFM	75 CFM	56 CFM	74 CFM	51 CFM

Operating temperature/humidity: 5 °C to 32 °C (41 °F to 89.6 °F), as measured by an industry grade temperature measurement device directed at the front bezel of the servers, 10% to 90% relative humidity, non-condensing

Altitude Operating: Up to 3,048 m, max. ambient temperature is de-rated by 1° C per 300 m above 900 m

¹ Typical power usage varies by application load

² Airflow must be front-to-back.

EXADATA DATABASE MACHINE ENVIRONMENTAL SPECIFICATIONS

Metric	Full Rack	Half Rack	Quarter Rack	Eighth Rack
Height	78.74" - 2000 mm 23.66" – 601 mm 47.13" – 1197 mm			
Width				
Depth				
Acoustic noise (operating)	9.4 B	9.2 B	8.8 B	8.8 B

Environmentals With High Capacity Disks

Weight	2019.1 lbs (915.8 kg)	1298.8 lbs (589.1 kg)	909.8 lbs (412.7 kg)	882.3 lbs (400.2 kg)
Maximum power usage	16.0 kW (16.3 kVA)	8.4 kW (8.6 kVA)	4.3 kW (4.4 kVA)	3.5 kW (3.5 kVA)
Typical power usage ¹	11.2 kW (11.4 kVA)	5.9 kW (6.0 kVA)	3.0 kW (3.0 kVA)	2.4 kW (2.5 kVA)
Cooling at maximum usage	54,499 BTU/hour 57,496 kJ/hour	28,645 BTU/hour 30,220 kJ/hour	14,560 BTU/hour 15,360 kJ/hour	11,789 BTU/hour 12,437 kJ/hour
Cooling at typical usage	38,149 BTU/hour 40,247 kJ/hour	20,051 BTU/hour 21,154 kJ/hour	10,192 BTU/hour 10,752 kJ/hour	8,252 BTU/hour 8,706 kJ/hour
Airflow at maximum usage ²	2523 CFM	1326 CFM	674 CFM	546 CFM
Airflow at typical usage ²	1766 CFM	928 CFM	472 CFM	382 CFM

Environmentals With Extreme Flash Drives

Weight	1793.7 lbs (813.6 kg)	1186.1 lbs (538.0 kg)	861.5 lbs (390.8 kg)	856.4 lbs (388.5 kg)
Maximum power usage	15.8 kW (16.1 kVA)	8.3 kW (8.5 kVA)	4.2 kW (4.3 kVA)	3.5 kW (3.5 kVA)
Typical power usage ¹	11.1 kW (11.3 kVA)	5.8 kW (5.9 kVA)	3.0 kW (3.0 kVA)	2.4 kW (2.5 kVA)
Cooling at maximum usage	54,021 BTU/hour 56,992 kJ/hour	28,406 BTU/hour 29,968 kJ/hour	14,457 BTU/hour 15,252 kJ/hour	11,782 BTU/hour 12,430 kJ/hour
Cooling at typical usage	37,815 BTU/hour 39,895 kJ/hour	19,884 BTU/hour 20,978 kJ/hour	10,120 BTU/hour 10,677 kJ/hour	8,247 BTU/hour 8,701 kJ/hour
Airflow at maximum usage ²	2501 CFM	1315 CFM	669 CFM	545 CFM
Airflow at typical usage ²	1751 CFM	921 CFM	469 CFM	382 CFM

Operating temperature/humidity: 5 °C to 32 °C (41 °F to 89.6 °F), as measured by an industry grade temperature measurement device directed at the front bezel of the servers 10% to 90% relative humidity, non-condensing

Altitude Operating: Up to 3,048 m, max. ambient temperature is de-rated by 1° C per 300 m above 900 m

¹ Typical power usage varies by application load.

² Airflow must be front-to-back.

EXADATA DATABASE MACHINE X8-2 REGULATIONS AND CERTIFICATIONS

Regulations ^{1,2,3}	Product Safety:	UL/CSA 60950-1, EN 60950-1, IEC 60950-1 CB Scheme with all country differences
	EMC	
	Emissions:	FCC CFR 47 Part 15, ICES-003, EN55032, EN61000-3-11, EN61000-3-12
	Immunity:	EN55024
Certifications ^{2,3}	North America (NRTL), European Union (EU), International CB Scheme, HSE Exemption (India), BSMI (Taiwan), CCC (PRC), EAC (EAEU including Russia), RCM (Australia), VCCI (Japan)	
European Union Directives ³	2014/35/EU Low Voltage Directive, 2014/30/EU EMC Directive, 2011/65/EU RoHS Directive, 2012/19/EU WEEE Directive	

¹All standards and certifications referenced are to the latest official version at the time the data sheet was written. For additional detail, please contact your sales representative.

² Other country regulations/certifications may apply.

³ In some cases, as applicable, regulatory and certification compliance were obtained for the shelf-level systems only.

EXADATA DATABASE MACHINE SUPPORT SERVICES

- *Hardware Warranty: 1 year with a 4 hr web/phone response during normal business hours (Mon-Fri 8AM-5PM), with 2 business day on-site response/Parts Exchange*
 - *Oracle Premier Support for Systems includes Oracle Linux support and 24x7 with 2 hour on-site hardware service response (subject to proximity to service center)*
 - *Oracle Premier Support for Operating Systems*
 - *Oracle Customer Data and Device Retention*
 - *System Installation Services*
 - *Software Configuration Services*
 - *Oracle Platinum Services*
 - *Business Critical Service for Systems*
 - *Oracle Exadata Start-Up Pack*
 - *System Upgrade Support Services including hardware installation and software configuration*
 - *Oracle Auto Service Request (ASR)*
-

OPTIONAL CUSTOMER SUPPLIED ETHERNET SWITCH INSTALLATION IN EXADATA DATABASE MACHINE X8-2

- *Each Exadata Database Machine X8-2 rack has 2U available at the top of the rack that can be used by customers to optionally install their own client network Ethernet switches in the Exadata rack instead of in a separate rack. Some space, power, and cooling restrictions apply.*
-

KEY FEATURES & FUNCTIONALITY

Exadata and Database Software Features - Analytics

- *Unique Automatic Parallelization and Offload of Data Scans to storage*
 - *Unique Filtering of Rows in Storage based on 'where' clause*
 - *Unique Filtering of Rows in Storage based on columns selected*
 - *Unique Storage Offload of JSON and XML Analytic Queries*
 - *Unique Filtering of rows in Storage based on Join with other Table*
 - *Unique Hybrid Columnar Compression*
 - *Unique Storage Index Data Skipping*
 - *Unique I/O Resource Management by User, Query, Service, DB, etc.*
 - *Unique Automatic Transformation to Columnar Format in Flash Cache*
 - *Unique Smart Flash Caching for Table Scans*
 - *Unique Storage Offload of Index Fast Full Scans*
 - *Unique Storage Offloads of Scans on Encrypted Data, with FIPS compliance*
 - *Unique Storage offload for LOBs and CLOBs*
 - *Unique Storage offload for min/max operations*
 - *Unique Data Mining Offload to Storage*
 - *Unique Intelligent Routing Across All InfiniBand Ports and Links to Minimize Congestion*
 - *Unique Reverse Offload to DB servers if Storage CPUs are Busy*
 - *Unique Automatic Data Columnarization in Flash Cache*
 - *Unique Automatic Conversion of Data to In-Memory Formats when Loading into Flash Cache*
-

Exadata and Database Software Features - OLTP

- *Unique Database Aware PCI Flash*
 - *Unique Exadata Smart Flash Caching*
 - *Unique Exadata Smart Flash Logging*
 - *Unique Smart Write-back Flash Cache*
 - *Unique I/O Prioritization by DB, User, or workload to ensure QOS*
 - *Unique Exafusion Direct-to-Wire Protocol*
 - *Unique Database Intelligent Network Resource Management*
 - *Unique Exachk full-stack validation*
 - *Unique Full-stack security scanning*
 - *Unique Database scoped security*
 - *Unique Cell-to-Cell Rebalance preserving Flash Cache and Storage Index*
 - *Unique Full-Stack Secure Erase*
 - *Unique Oracle VM with InfiniBand SRIOV*
 - *Unique Automated InfiniBand Partitioning for Workload Isolation*
 - *Unique Instant Data File Creation*
 - *Unique Smart Fusion Block Transfer*
 - *Unique Control of Flash Cache Size per Database*
 - *Unique In-Memory OLTP Acceleration*
 - *Unique Undo-Block Remote RDMA Read*
 - *Unique Support for More Than 252 Pluggable Databases with Multitenant Option*
-

Exadata and Database Software Features - High Availability

- *Unique Instant Detection of Node or Cell Failure*
 - *Unique In-Memory Fault Tolerance*
 - *Unique Sub-second Failover of I/O on stuck disk or Flash*
 - *Unique Offload backups to storage servers*
 - *Unique Exadata Data Validation (extended H.A.R.D.)*
 - *Unique Prioritize Recovery of Critical Database Files*
 - *Unique Automatic Repair of Corrupt Disk Data By Reading Other Storage Servers*
 - *Unique Avoidance of Read I/Os on Predictive failed disks*
 - *Unique Confinement and power cycle of temporarily poor performing drives*
 - *Unique Shutdown Prevention If Mirror Storage Server is Down*
 - *Unique Detection and Disabling of Unreliable Network Links*
 - *Unique Preservation of Storage Index on Rebalance*
-

Manageability Features

- *Oracle Embedded Integrated Lights Out Manager (ILOM)*
- *Oracle Enterprise Manager Exadata Plug-in*
- *Unique Active AWR includes storage stats for end to end monitoring*
- *IPv6 Support for Ethernet Connections*
- *Capacity on Demand*
- **Cell software transparent restart**
- **Flash and disk life cycle management alert**
- **Automatic Disk Scrub and Repair**
- *Trusted Partitions for Oracle Virtual Machine*
- *Automated VLAN Creation*
- *Oracle Exadata Deployment Assistant*
- *Separate Management Switch and Connectivity*
- *Exacli command line management from remote servers*
- *Cellcli command line management of Storage Servers*
- *DCLI distributed command line automation tool*

Oracle Database Software (available separately):

For database servers: Oracle Database 11g Release 2 Enterprise Edition, Oracle Database 12c Enterprise Edition Release 1 and 2, Oracle Database 18c Enterprise Edition Release 1, and Oracle Database 19c. Oracle Database Options such as Oracle Real Application Clusters, Oracle Partitioning, Oracle Multitenant, Oracle Active Data Guard. See the release specific documentation for feature support.

For storage servers: Oracle Exadata System Software. Licenses are transferable from one system to another, or to a new system

Oracle Software (included):


For database servers: Oracle Linux 7 Update 5 with the Unbreakable Enterprise Kernel 4. Zero-loss Zero-copy Datagram Protocol (ZDP) InfiniBand protocol used to communicate between the Exadata Storage Servers and the Oracle Database which is based on the Reliable Datagram Sockets (RDS) OpenFabrics Enterprise Distribution (OFED)



CONTACT US

For more information about Oracle Exadata, visit oracle.com/exadata or call +1.800.ORACLE1 to speak to an Oracle representative.

CONNECT WITH US



Integrated Cloud Applications & Platform Services

Copyright © 2017, 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.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.



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