ORACLE EXADATA DATABASE MACHINE X7-8

The Oracle Exadata Database Machine is engineered to deliver dramatically better performance, cost effectiveness and availability for Oracle databases. Exadata features a modern cloud-based architecture with scale-out high-performance database servers, scale-out intelligent storage servers, state-of-the-art PCI flash storage servers and an ultra-fast InfiniBand internal fabric that connects all servers and storage. Unique software algorithms 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 deployed on premises as the ideal foundation for a private database cloud, or can be used on the Oracle Public Cloud for same-day deployment with no infrastructure management or Capex.

The Exadata Database Machine X7-8 is a high-end database machine that has the same extreme performance, storage, and InfiniBand as the X7-2, but uses large-scale 8-socket SMP servers instead of the 2-socket servers in X7-2. Each of the 8-socket servers in the X7-8 has a total of 192 processor cores and 3 terabytes to 6 terabytes of DRAM. The Exadata Database Machine X7-8 excels at high-end online transaction processing (OLTP), multi-rack Data Warehouses, and In-Memory Database workloads. It is ideally suited for applications that desire large scale-up servers with large memory capacity.

Engineered System for Fast Deployment of All Your Databases

The Exadata Database Machine is an easy to deploy system that includes 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 a mission critical high performance system. Extensive end-to-end testing ensures all components including database software, OS, firmware, drivers, etc. work seamlessly together and 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 including most of the top Banks, Telcos and Retailers in the world. Customer machines are also identical to the machines Oracle Support uses for problem identification and resolution, the machines Oracle Engineering uses for development and testing of Oracle Database, and the machines Oracle uses to implement its own SaaS and PaaS public cloud. Hence, **Exadata is the most thoroughly tested and tuned platform for running Oracle Database and is also the most supportable platform.**

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 of Exadata eliminating any fear of "lock-in".

Customers thinking of deploying databases on the Public Cloud, now or in the future, can be confident that **Exadata provides 100% compatibility between private and public clouds** enabling easy migration to the public cloud and simple 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. The Exadata configuration carefully balances CPU, I/O and network throughput to avoid bottlenecks. As an Exadata Database Machine grows, database CPUs, storage, and networking are added in a balanced fashion ensuring scalability without bottlenecks.

The scale-out architecture accommodates any size workload and allows seamless expansion from small to extremely large configurations while avoiding performance bottlenecks and single points of failure.

Within the Exadata family, Exadata X7-8 is a high-end database machine that is particularly well suited to large OLTP databases that benefit from running on small numbers of large compute nodes, in-memory database workloads requiring very large memory capacity, large multi-rack configurations that benefit from reduced numbers of database nodes, and databases with high I/O requirements that can benefit from very large buffer caches.

Each Exadata X7-8 machine comes with very powerful 8-socket database servers and 2-socket storage servers. Each database server has eight 24-core x86 processors with 3TB of memory (expandable up to 6TB). Each storage server has two 10-core x86 processors for SQL processing. For High Capacity (HC) storage configurations, each storage server comes with four 6.4TB NVMe PCI Flash cards and twelve 10TB 7,200 RPM HC disks. For Extreme Flash (EF) storage configurations, each storage server comes with eight 6.4TB NVMe PCI Flash drives.

In addition to the inherent scalability and power of Exadata, elastic configurations provide an extremely flexible and efficient mechanism to expand computing power

KEY FEATURES

- Up to 576 CPU cores and up to 18 TB of memory for database processing per rack
- Up to 280 CPU cores dedicated to SQL processing in storage
- From 2 to 3 8-socket database servers per rack
- From 3 to 14 Oracle Exadata Storage Servers
- 40 Gb/second InfiniBand
- Up to 720TB of Flash Storage per rack
- Hybrid Columnar Compression often delivers 10X-15X compression ratios
- Complete redundancy for high availability
- Fault-tolerant In-Memory Duplication for Oracle Database In-Memory
- Oracle Linux database servers

KEY BENEFITS

- Uncompressed I/O bandwidth of up to 350 GB/second per rack from SQL
- Ability to perform up to 5.48 million database 8K read I/O operations per second
- Scale by connecting multiple Exadata Database Machine X7-8 racks or Exadata Storage Expansion Racks.
- Up to 18 racks can be connected by simply connecting via InfiniBand cables and using internal switches. Larger configurations can be built with external InfiniBand switches
- Pre-configured system optimized for all database applications

and/or storage capacity of any given Exadata system. With Elastic Configurations, the system can be custom configured to meet any business use case.

The starting configuration of an Exadata X7-8, which is called the Half Rack configuration for Exadata X7-8, consists of 2 database servers and 3 storage servers. This can be further elastically expanded by adding more database or storage severs as requirements grow. The Exadata X7-8 can have 2 to 3 database servers and 3 to 14 storage servers. A full rack Exadata Database Machine X7-8 consists of 2 database servers and 14 storage servers.

A high-bandwidth low-latency 40 Gb/second InfiniBand network connects all the components inside an Exadata Database Machine. Specialized database networking protocols run over the InfiniBand network and provide much lower latency and higher bandwidth communication than is possible using generic communication protocols. This enables both faster response time for OLTP operations, and higher throughput for Analytic workloads. External connectivity to the Exadata Database Machine is provided using standard 1,10 or 25 Gigabit Ethernet.

Exadata supports multiple racks to be connected using the integrated InfiniBand fabric to form even larger configurations. For example, a system composed of four Full Racks is simply four times as powerful as a single rack system – providing quadruple the I/O throughput, quadruple the storage capacity, and quadruple the processors. It can be configured as a large single system or logically partitioned for consolidation of multiple databases. Scaling out is easy with Exadata Database Machine. Oracle Real Application Clusters (RAC) can dynamically add more processing power, and Oracle Automatic Storage Management (ASM) can dynamically add more storage.

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 you to grow the Exadata storage capacity and I/O bandwidth of any Exadata Database Machine. It is designed for database deployments that require very large amounts of data including:

RELATED PRODUCTS

- Exadata Database Machine X7-2
- Exadata Storage Server X7-2
- Exadata Storage Expansion X7-2
- Oracle SuperCluster
- Oracle Database 11g,12c, 18c and 19c
- Real Application Clusters
- Database In-Memory
- Partitioning
- Multitenant
- Advanced Compression
- Advanced Security
- Active Data Guard
- GoldenGate
- Real Application Testing
- OLAP
- Advanced Analytics
- Business Intelligence
- Enterprise Manager
- Oracle Linux

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

historical or archive data, backups, documents, images, XML, JSON and, LOBs. It is extremely simple to configure as there are no LUNs or mount points to set up. It connects to the Exadata Database Machine using the integrated InfiniBand fabric. Storage is configured and added to a database online with a few simple commands.

Exadata Database Machines protect your investment by allowing newer generation servers and storage to be deployed seamlessly into existing Exadata Database Machines. Similarly, new software releases are compatible with previous generation Exadata Database Machines. All currently supported Exadata platforms can be combined in a single configuration and can run the latest Exadata software.

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. Exadata X7 enhances previous generation EF Storage Servers, by doubling the flash capacity. Each EF Storage Server contains eight 6.4 TB state-of-the-art Flash Accelerator F640 PCI flash drives, offering 51.2 TB raw flash capacity per EF Storage Server. In addition, Exadata delivers ultra-high performance by placing the flash drives directly on the high speed PCI bus rather than behind slow disk controllers and directors. Finally, Exadata flash uses the latest **NVMe** (Non-Volatile Memory Express) flash to achieve extremely low latency and CPU overhead
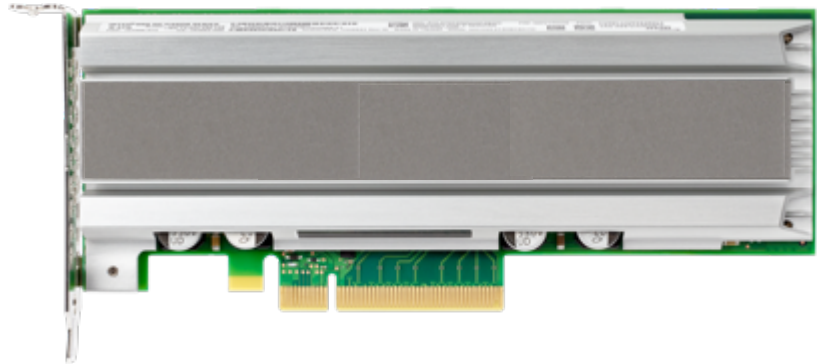

Fig. 2: Flash Accelerator F640 PCIe Card

Flash performance is often limited and bottlenecked by 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 rates from flash. A single rack Exadata Database Machine elastic configuration X7-8, with 3 database servers and 11 Extreme Flash storage servers can achieve up to 5.48 Million random 8K read I/O operations per second (IOPS) when running database workloads. A full rack Exadata Database Machine X7-8, with 2 database servers and 14 Extreme Flash storage servers, can achieve up to **350 GB per second of analytic scan bandwidth from SQL**, **at 4.38 Million Flash IOPS** when running database workloads. This performance is orders of magnitude faster than traditional storage array architectures, and is also much faster than current all-flash storage arrays. It is important to note that these are real-world end-to-end performance figures measured running SQL workloads

with realistic I/O sizes inside a single rack Exadata system. They are not component-level measurements based on low-level I/O tools.

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

The second Exadata storage option in X7-8 is the Exadata X7-2 **High Capacity Storage Server**. This server includes twelve 10 TB SAS disk drives with 120 TB total raw disk capacity. It also has four Flash Accelerator F640 NVMe PCIe cards, and similar to the EF Storage Server, Exadata X7 doubles the flash capacity for HC storage servers with a total capacity of 25.6 TB of flash. 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 since intelligent caching provides flash level performance for much more data than fits directly into flash.

Exadata Smart Flash Cache automatically caches frequently accessed data while keeping infrequently accessed data on disk. 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 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, 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 always retained in the flash cache. Tables can be retained in flash cache without the need to move the table to different tablespaces, files or LUNs as is often required with traditional storage. A single full rack Exadata Database Machine X7-8, with 2 database servers and 14 High Capacity storage servers can achieve up to **350 GB per second of data scan bandwidth, and up to 4.38 Million random 8K read I/O operations per second (IOPS)** when running database workloads.

Exadata's Smart Flash Cache is designed to deliver flash-level I/O rates and response times for data that is many times larger than the physical flash capacity in the machine by moving active data into flash, while leaving cold data on disk. It is common for hit rates in the Exadata Smart Flash Cache to be over 90%, or even 98% in real-world database workloads even though flash capacity is more than 4 times smaller than raw disk capacity. 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 cache. For example, a full rack Exadata Database Machine X7-8 with 2 database servers and 14 High Capacity Storage Servers often has an effective flash capacity equal to the usable triple-mirrored disk capacity of 500 TB.

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 full rack Exadata Database Machine X7-8 with 2 database servers and 14 High Capacity Storage Servers exceeds **4.38 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 all-flash arrays. These all-flash arrays deliver higher performance than traditional arrays but give up the cost advantages of smart tiering of data between disk and flash. Therefore the overall size of data that benefits from flash is limited to the size of expensive flash. Third party flash arrays will also not benefit from Exadata Hybrid Columnar Compression. Data deduplication provided by some flash arrays is very effective for VDI (Virtual Desktop Infrastructure) environments but is ineffective for databases.

Exadata not only delivers much more capacity than all-flash arrays, it also delivers better performance. All-flash 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 that are specifically optimized for databases.

Accelerating Database Processing with Smart System Software

As data volumes grow exponentially, 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.

The technology that 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 10-core x86 processors that are used to offload database processing. A rack of Exadata Database Machine can have a total of up to 280 processor cores in the storage servers that can be used to offload the database servers. The CPUs in the storage servers do not replace database CPUs. Instead they accelerate data 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 names. The result is that the data transferred to the database servers is reduced by orders of magnitude. This greatly accelerates query execution, eliminates bottlenecks, and significantly reduces the CPU usage of the database servers.

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, 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 avoided. 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 the 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 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. Based on internal tests, it has been observed that Smart Block Transfer increases throughput (about 40% higher) and decreases response times (about 33% less) for communication intensive workloads.

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 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. The Exadata Smart Flash Logging algorithms are unique to Exadata.

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. As the name implies, this technology utilizes 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 tremendous 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 by consuming CPU for decompression. Because the Exadata Database Machine is able to offload decompression overhead into large numbers of processors in Exadata storage, and in addition there is reduced I/O because of the high compression achieved, most analytic workloads run faster using Hybrid Columnar Compression than they do without it. Hybrid Columnar Compression delivers the compression and analytic performance benefits of columnar storage while avoiding the dramatic slowdown that columnar-only data stores experience for drilldown operations that often involve single row access.

Two modes of Hybrid Columnar Compression are available. **Warehouse compression** mode is suitable for read intensive workloads such as Data Warehouses and provides large storage savings while providing 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.

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 18c 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, which benefits from pure columnar access, 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.

Elastic configurations enable users to configure systems with very large system memory. Each 8-socket SMP server on X7-8 can directly access terabytes of memory using its ultra-fast memory interconnect. This often enables in-memory queries to avoid the communication overhead of distributing processing across multiple nodes of a cluster. When in-memory queries are distributed across nodes, the 40Gb per second InfiniBand network provides extremely high throughput and very low latencies.

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's loaded into flash cache. Smart Scans then leverage ultra fast Single Instruction Multiple Data (SIMD) Vector instructions, thus processing multiple column values with a single instruction. Smart Scan results are passed back to the database server in Oracle Database In-Memory formats, thus further reducing the CPU utilization at the database servers. This has the net effect of seamlessly extending the In-Memory columnar store size from DRAM capacity in the database server to flash capacity in storage servers. An Exadata X7-8 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**. Memory Fault Tolerance is unique to Oracle Engineered Systems. 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 of data on a surviving database server and processing will 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 DRAM memory installed in Exadata Storage Servers as an extension of the memory cache (buffer cache) on database servers. Specialized algorithms transfer data between the buffer cache on database servers and DRAM cache on storage servers. This reduces the IO latency to 100 microseconds for all IOs served from in-memory cache. Exadata **uniquely** keeps only one in-memory copy of data across database and storage servers, avoiding memory wastage 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.

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, 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 hundreds of gigabytes of user data per second.

The Exadata system is designed and delivered as an integrated whole, and not a collection of components. In traditional database deployments, the customer takes on all the integration tasks for the system – 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 Linux distributions to ensure that just the RPMs needed to run Oracle Database are installed and enabled. With this approach, system security is stronger than default installations and common security vulnerabilities are avoided.

Exadata security has been probed and evaluated by hundreds of leading banks, telcoms, and government organizations worldwide. The security findings of all these evaluations have been incorporated into the Exadata standard configuration, making it a highly secure database platform.

Mission Critical High Availability

The Exadata Database Machine is engineered to provide the highest levels of availability. **All types of failures are protected against** including simple failures such as disk, server, or network, as well as 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 configured in a Maximum Availability Architecture (MAA) configuration to transparently maintain a real-time copy of the database at a remote site to provide full protection against primary database failures and site disasters.

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 Systems[1].

---

[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

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 very quickly determine that the suspect server is not reachable through any network path and can 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 IO operations due to internal recovery of failed sectors, internal firmware reboots, or wear leveling. These long IO 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. If disks do fail, ASM performs a rebalance operation for the data that was resident on the disk. 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.

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 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 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 provides unique **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 implements unique **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 deployment.

Exadata also implements unique **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 more than 252 Pluggable Databases within a single Oracle Multitenant Container Database.

Fast Deployment of Development and Test Databases with Exadata Snapshots

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 can be cleansed of any sensitive information. 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, multiple test and development environments can share space while maintaining independent databases for each task.

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.

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 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 1.68 Petabytes while raw all-flash storage can be up to 716.8 TB. In addition, Hybrid Columnar Compression often expands storage and memory capacity 10X. By intelligently moving active data across storage and memory tiers, Exadata simultaneously delivers the highest performance at 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, 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 X7-8 database server has a substantial amount of compute capacity with eight 24-core processors (192 cores in total).

The Capacity-on-Demand feature allows up to nearly 70% of the cores per server to be turned off during the hardware installation, leaving at least 56 cores enabled for the server. As your workload grows and more cores are needed, Capacity-on-Demand can be used to re-enable cores and license software 8 cores at a time. This pay-as-you-grow approach to software licensing is yet another way in which Exadata helps to align costs with business growth.

Exadata Business Benefits

Beyond the operational benefits of extreme performance, availability, and security at low cost, 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 Database.

EXADATA DATABASE MACHINE X7-8 HARDWARE

Components

2 to 3 Database Servers, each with:

- *8 x Twenty-Four-Core Xeon® 8168 Processors (2.7 GHz)*
- *3 TB Memory (expandable up to 6 TB)*
- *Disk Controller HBA with 2GB Cache*
- *2 x 6.4 TB PCI NVMe Flash Cards*
- *8 x InfiniBand QDR (40Gb/s) Ports*
- *8 x 10/25 Gb Ethernet Ports based*
- *8 x 1/10 Gb Ethernet Ports (One of which is used for Host Admin)*
- *1 x ILOM Ethernet Port*
- *4 x Redundant Hot-Swappable Power Supplies*

3 to 14 Exadata Storage Servers X7-2 HC or EF each with:

HC	EF
• *2 x Ten-core processor for SQL processing* • *4 PCI Flash cards each with 6.4 TB (raw) Exadata Smart Flash Cache* • *12 x 10 TB 7,200 RPM High Capacity disks*	• *2 x Ten-core processor for SQL processing* • *8 x 6.4 TB NVMe PCI Flash Drives*

2 x 36 port QDR (40 Gb/sec) InfiniBand Switches

Additional Hardware Components:

- *42U Rack*
- *Ethernet switch for administrative connectivity to servers in the Database Machine*
- *2 x Redundant Power Distributions Units (PDUs)*

Included Spare Parts Kit Contains:

- *1 x 6.4 TB NVMe PCI Flash Card and 1 x 10TB High Capacity disk, or*
- *1 x 6.4 TB NVMe PCI Flash drive*

EXADATA X7-8 TYPICAL RACK CONFIGURATIONS: KEY CAPACITY AND PERFORMANCE METRICS

Metric	Full Rack HC[1]	Full Rack EF[1]	Half Rack HC	Half Rack EF
Number of Servers				
Database Servers	2	2	2	2
Storage Servers	14	14	3	3
Flash Metrics				
Maximum SQL Flash Bandwidth [2]	350 GB/sec	350 GB/sec	75 GB/sec	75 GB/sec
Maximum SQL Flash Read IOPS[3]	4,382,000	4,382,000	1,425,000	1,494,000
Maximum SQL Flash Write IOPS[4]	4,382,000	4,382,000	1,260,000	1,350,000
Flash Data Capacity (raw) [5]	358.4 TB	716.8 TB	76.8 TB	153.6 TB
Disk Metrics				
Maximum SQL Disk bandwidth [2]	25 GB/sec	N/A	5.4 GB/s	N/A
Maximum SQL Disk IOPS [3]	36,000	N/A	7,800	N/A
Disk Data Capacity (raw) [5]	1680 TB	N/A	360 TB	N/A
Combined Metrics				
Data Capacity (usable) – Normal Redundancy [6]	681 TB	282TB	136 TB	56 TB
Data Capacity (usable) – High Redundancy [6]	499 TB	206 TB	107 TB	44 TB
Maximum Data Load Rate [7]	35 TB/hr	35 TB/hr	8 TB/hr	8 TB/hr

Actual system performance varies by application.

[1] EF = Extreme Flash; HC = High Capacity

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

[3] Based on 8K IO requests running SQL. Note that the IO size greatly affects Flash IOPS. Others quote IOPS based on smaller IOs and are not relevant for databases.

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

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

[6] Usable capacity is measured using normal powers of 2 space terminology with 1 TB = 1024 * 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, DBFS disk group, and OS images and binaries. Normal redundancy calculations reflect the use of Grid Infrastructure version 12.2.0.1 or later.

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

EXADATA DATABASE MACHINE X7-8 EXPANSION AND UPGRADES

Expansion & Upgrades
Expansion: 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 or X7 generation hardware.
Upgrade Support Services: • *Hardware Installation and Software configuration*

EXADATA DATABASE MACHINE X7-8 SUPPORT SERVICES

Components

- *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)*

EXADATA DATABASE MACHINE X7-8 ENVIRONMENTAL SPECIFICATIONS

Metric	Full Rack	Half Rack
Height Width Depth	• *78.74" - 2000 mm* • *23.66" – 601 mm* • *47.13" – 1197 mm*	
Acoustic noise (operating)	9.5 B	9.2 B
Environmentals With High Capacity Disks		
Weight	2044.2 lbs (927,2 kg)	1247.3 lbs (565.8 kg)
Maximum power usage Typical power usage [1]	16.7 kW (17.0 kVA) 11.7 kW (11.9 kVA)	9.8 kW (10.0 kVA) 6.8 kW (7.0 kVA)
Cooling at maximum usage Cooling at typical usage	56,932 BTU/hour 60,063 kJ/hour 39,852 BTU/hour 42,044 kJ/hour	33,323 BTU/hour 35,156 kJ/hour 23,326 BTU/hour 24,609 kJ/hour
Airflow at maximum usage [2] Airflow at typical usage [2]	2636 CFM 1845 CFM	1543 CFM 1080 CFM
Environmentals With Extreme Flash Drives		
Weight	1830 lbs (830.1 kg)	1201.4 lbs (544.9 kg)
Maximum power usage Typical power usage [1]	16.8 kW (17.1 kVA) 11.7 kW (12.0 kVA)	9.8 kW (10.0 kVA) 6.8 kW (7.0 kVA)
Cooling at maximum usage Cooling at typical usage	57,218 BTU/hour 60,365 kJ/hour 40,053 BTU/hour 42,256 kJ/hour	33,384 BTU/hour 35,221 kJ/hour 23,369 BTU/hour 24,654 kJ/hour
Airflow at maximum usage [2] Airflow at typical usage [2]	2649 CFM 1854 CFM	1546 CFM 1082 CFM

Operating temperature/humidity: 5 ºC to 32 ºC (41 ºF to 89.6 ºF), 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

Regulations [3]

- *Safety: UL/CSA 60950-1, EN 60950-1, IEC 60950-1 CB Scheme with all country differences*
- *RFI/EMI: EN55022, EN61000-3-11, EN61000-3-12*
- *Immunity: EN 55024*
- *Emissions and Immunity: EN300 386*

Certifications [3]

- *North America (NRTL), European Union (EU), International CB Scheme, BSMI (Taiwan), C-Tick (Australia), CCC (PRC), MSIP (Korea), CU EAC (Customs Union), VCCI (Japan)*

European Union Directives [3]

- *2014/30/EU Low Voltage Directive, 2014/30/EU EMC Directive, 2011/65/EU RoHS Directive, 2012/19/EU WEEE Directive*

[1] Typical power usage varies by application load.

[2] Airflow must be front-to-back

[3] All standards and certifications referenced are to the latest official version at the time the data sheet was written. Other country regulations/certifications may apply. In some cases, as applicable, regulatory and certification compliance were obtained at the component level.

KEY FEATURES & FUNCTIONALITY

Exadata and Database Software Features - Analytics

- *Automatically Parallelize and Offload Data Scans to storage*
- *Filter Rows in Storage based on 'where' clause*
- *Filter Rows in Storage based on columns selected*
- *JSON and XML Offload*
- *Filter rows in Storage based on Join with other Table*
- *Hybrid Columnar Compression*
- *Storage Index data skipping*
- *I/O Resource Management by User, Query, Service, DB, etc.*
- *Automatic Transformation to Columnar Format in Flash Cache*
- *Smart Flash Caching for Table Scans*
- *Offload Index Fast Full Scans*
- *Offloads Scans on Encrypted Data, with FIPS compliance*
- *Storage offload for LOBs and CLOBs*
- *Storage offload for min/max operations*
- *Data Mining Offload*
- *All Ports Active InfiniBand Messaging*
- *Reverse Offload to DB servers if Storage CPUs are Busy*
- *Automatic Data Compression in Flash Cache*
- *Offload JSON and XML analytic queries*

Exadata and Database Software Features - OLTP

- *Database Aware PCI Flash*
- *Exadata Smart Flash Caching*
- *Exadata Smart Flash Logging*
- *Write-back Flash Cache*
- *I/O Prioritization by DB, User, or workload to ensure QOS*
- *Exafusion Direct-to-Wire Protocol*
- *Network Resource Management*
- *Exachk full-stack validation*
- *Full-stack security scanning*
- *NVMe flash interface for lowest latency IO*
- *Database scoped security*
- *Cell-to-Cell Rebalance preserving Flash Cache*
- *Secure disk and flash erase*
- *InfiniBand Partitioning*
- *Instant data file creation*
- *Active Bonding of InfiniBand*
- *Smart Fusion Block Transfer*
- *Automatic VLAN creation*
- *Set Minimum or Maximum Flash Cache Size per Database*

Exadata and Database Software Features - High Availability

- *Instant Detection of Node or Cell Failure*
- *In-Memory Fault Tolerance*
- *Sub-second Failover of I/O on stuck disk or flash*
- *Offload backups to storage servers*
- *Exadata Data Validation (H.A.R.D.)*
- *Prioritize rebalance of critical files*
- *Automatic Hard Disk scrub and repair*
- *Power cycle failed drives to Eliminate false drive failures*
- *Avoid reading Predictive failed disks*
- *Cell software transparent restart*
- *Flash and disk life cycle management alert*
- *Confinement of temporarily poor performing drives*
- *Prevent shutdown if mirror server is down*
- *Detection and disabling of unreliable network links*
- *Preserve Storage Index on Rebalance*
- *Automatic disk scrub and repair*

Manageability Features

- *Oracle Embedded Integrated Lights Out Manager (ILOM)*
- *Oracle Enterprise Manager Exadata Plug-in*
- *Active AWR includes storage stats for end to end monitoring*
- *IPv6 Support for Ethernet Connections*
- *Capacity on Demand*
- *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, Oracle Database 18c Enterprise Edition and Oracle Database 19c Enterprise Edition. 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)