



# Oracle Exadata Database Machine X10M



## Oracle Exadata Database Machine X10M

The Oracle Exadata Database Machine (Exadata) 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 PCIe flash, unique storage caching using RDMA accessible memory, and cloud-scale RDMA over Converged Ethernet (RoCE) internal fabric that connects all servers and storage. Unique algorithms and protocols in Exadata implement database intelligence in storage, compute, and networking to deliver higher performance and capacity at lower costs than other database platforms. Exadata is ideal for all types of modern database workloads, including Online Transaction Processing (OLTP), Analytics and Data Warehousing (DW), In-Memory Analytics, Artificial Intelligence (AI), Internet of Things (IoT), financial services, gaming, and compliance data management, as well as efficient consolidation of mixed database workloads.

Simple and fast to implement, the Exadata Database Machine X10M 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 acquired using a subscription model and deployed in the Oracle Public Cloud or Cloud@Customer with all infrastructure management performed by Oracle. The Oracle Autonomous Database is available exclusively on Exadata, either in the Oracle Public Cloud or Cloud@Customer.

### Engineered For Fast and Reliable Deployment

The Exadata Database Machine is the most cost-efficient and highest-performance platform for running Oracle Databases. Exadata is easy to deploy even for the most mission-critical systems, as the database servers, storage servers, and network are pre-configured, pre-tuned, and pre-tested by Oracle experts. Extensive end-to-end testing and validation ensure all components, including database software, operating system, hypervisor, drivers, and firmware, work seamlessly together and that there are no performance bottlenecks or single points of failure.

### Key Features

- Up to 2,880 CPU cores per rack for database processing
- Up to 33 TB memory per rack for database processing
- Up to 1,088 CPU cores per rack dedicated to SQL processing in storage
- Up to 21.25 TB of Exadata RDMA Memory per rack
- 100 Gb/sec RoCE Network
- Complete redundancy for high availability
- From 2 to 15 database servers per rack
- From 3 to 17 storage servers per rack
- Up to 462.4 TB of performance-optimized flash capacity (raw) per rack
- Up to 2 PB of capacity-optimized flash capacity (raw) per rack
- Up to 4.2 PB of disk capacity (raw) per rack

All Exadata Database Machines are identically configured; therefore, customers benefit from the experience of thousands of other customers' Exadata Database Machine deployments. Customer machines are also identical to the machines Oracle Support uses for problem identification and resolution, the machines Oracle Development develop and test the Oracle Database, and the machines that run Oracle's own public cloud services. **Exadata is the most thoroughly tested and tuned platform for running Oracle Database.**

**Any Oracle Database on any supported platform can be seamlessly migrated to the Exadata Database Machine** with no changes to the application using that database. Likewise, any Oracle Database can also be easily migrated off Exadata, eliminating "lock-in" concerns.

## Extreme System Scalability and Growth with Elastic Configurations

The Exadata Database Machine uses a scale-out architecture for both database and storage servers. As workloads grow, database, storage, and networking resources can be added to an Exadata Database Machine to scale without bottlenecks. The **architecture scales from small to extremely large configurations to accommodate workloads of any size.** In Exadata X10M, a high-bandwidth, low-latency active-active 100 Gb/sec RDMA over Converged Ethernet (RoCE) Network Fabric connects all the components inside an Exadata Database Machine. Specialized database networking protocols deliver much lower latency and higher bandwidth than is possible with generic communication protocols for **faster response time for OLTP operations and higher throughput for analytic workloads.** External data center connectivity to the Exadata Database Machine is via standard 10 Gb/sec, 25 Gb/sec, or 100 Gb/sec Ethernet.

**Exadata Database Machine is the most versatile database platform with extreme scalability inherent at every layer in its architecture.** The Exadata X10M Database Machines powerful database servers are equipped with two 96-core x86 processors and 512 GB of DDR5 memory (expandable up to 3 TB).

The scalability within each database server enables superior database and VM consolidation with higher OLTP transaction throughput and significantly more parallelized analytic workloads to coexist and consume fewer data center resources.

Exadata also uses scale-out, intelligent storage servers for database I/O processing, which are available in three configurations – High Capacity (HC), Extreme Flash (EF), or Extended Storage (XT):

- HC Storage Servers have four 6.8 TB performance-optimized Flash Accelerator F680 NVMe PCIe Flash cards for Exadata Smart Flash Cache and twelve 22 TB 7,200 RPM SAS disks for a total of 264 TB of raw storage.
- EF Storage Servers have an all-flash configuration. Each has four 6.8 TB performance-optimized Flash Accelerator F680 NVMe PCIe Flash drives for Exadata Smart Flash Cache and four 30.72 TB capacity-optimized NVMe PCIe Flash drives totaling 122.88 TB of raw storage.

### Key Benefits

- Pre-configured, pre-tested system optimized for all database applications
- Uncompressed I/O bandwidth of up to 1 TB/sec per rack from SQL
- Ability to perform up to 25.2M 8K database read I/O operations, or 10.9M 8K Flash write I/O operations per second in a single rack
- Easily add database or storage servers to meet the needs of any size application
- Extreme Scalability by connecting multiple Exadata Database Machine X10M racks or Exadata X10M Storage Expansion Racks. Up to 14 racks can be connected by simply adding RoCE cables and internal switches. Larger configurations can be built with external RoCE switches

- Exadata X10M HC and EF Storage Servers include Exadata RDMA Memory (XRMEM), further enhancing performance by delivering up to 2.8 million 8K OLTP Read IOPS – a 21% increase over the previous generation – and 45GB/s and 60GB/s SQL Scan throughput respectively. Each server is populated with 1.5 TB of DDR5 memory, 1.25 TB of which is used as a caching tier between the database buffer cache and Flash Cache, and the remaining 256 GB is used for Exadata System Software. Two 32-core x86 processors are included for Exadata System Software operations in each HC and EF storage server.
- XT Storage Servers are ideal for in-database data archival using Exadata Hybrid Columnar Compression. Each XT Storage Server is populated with twelve 7,200 RPM SAS disks with 22TB storage capacity, totaling 264 TB of raw storage per server and 128 GB of DDR5 memory. One 32-core x86 processor is included for Exadata System Software operations in each XT storage server. Oracle Exadata System Software may be optionally licensed on XT Storage Servers and enables features such as Exadata Smart Scan. Hybrid Columnar Compression is included with XT storage.

The minimal configuration of an Exadata Database Machine consists of two database servers and three HC storage servers or EF storage servers. This configuration **can elastically expand by adding more database and/or storage servers within the same rack**. Elastic configurations provide a flexible and efficient mechanism to meet any size business needs.

In addition to expanding within a rack, **multiple RoCE-based Exadata racks can be interconnected using the integrated RoCE network fabric** to form even larger configurations. These racks can be Exadata X8M, X9M, or X10M generations. For example, a system composed of four racks of Exadata X10M is simply four times as powerful as a single rack: it provides four times the I/O throughput, four times the storage capacity, and four times the processing power. It can be configured as a single system or logically partitioned for multiple databases and clusters. 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.




Figure 1 Elastic Scale-out to Multi-rack Exadata

**Related Products**

- Exadata Cloud Infrastructure
- Exadata Cloud@Customer
- Oracle Exadata Storage Expansion Rack X10M
- Oracle Exadata Storage Server X10M
- Oracle Exadata Database Server X10M
- Oracle Database Enterprise Edition 19c, 21c and 23ai
- Exadata Database Service
- Oracle Autonomous Database
- Real Application Clusters
- Partitioning
- Multitenant
- Database In-Memory
- Advanced Compression
- Advanced Security
- Active Data Guard
- GoldenGate
- Real Application Testing
- OLAP
- Enterprise Manager
- Oracle Linux
- Oracle Linux Virtualization

**Related Services**

The following services are available from Oracle:

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

For databases and workloads that require extreme storage capacity, the **Oracle Exadata X10M Storage Expansion Rack** is used to expand the storage tier of Exadata Database Machine. The Storage Expansion Rack expands the storage capacity, Flash Cache capacity, OLTP IOPS, and scan throughput 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 RoCE network fabric and is configured 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 HC or EF storage servers and can be expanded with additional storage servers.

*“We have implemented nearly 300 Exadata systems for our customers in manufacturing, financial services, construction and engineering, and public and private sector services.”*

**Dr. WP Hong**  
CIO  
Samsung SDS

## Groundbreaking RDMA-Based Network Fabric

The Exadata X10M release uses the same ultra-fast cloud-scale networking fabric that was introduced in Exadata X8M, RDMA over Converged Ethernet (RoCE). RDMA (Remote Direct Memory Access) allows one computer to directly access data from another without Operating System or CPU involvement for high bandwidth and low latency. The network card directly reads/writes memory with no extra copying or buffering and with very low latency. RDMA is an integral part of the Exadata high-performance architecture and has been tuned and enhanced with each new generation of Exadata, underpinning several Exadata-only technologies such as Exafusion Direct-to-Wire Protocol and Smart Fusion Block Transfer.

The Exadata X10M release implements a dual port PCIe Gen 5 network interface card capable of 2x 100 Gb/sec active-active RoCE network for a total throughput of 200 Gb/sec. By leveraging the RoCE network, **Oracle Database on Exadata can perform read I/O directly from memory in the shared storage servers.**

## Shared Exadata RDMA Memory Acceleration

Exadata RDMA Memory (XRMEM) in the Exadata Storage Servers is leveraged as a shared read accelerator. The XRMEM Data Accelerator is a memory cache tier in front of Flash Cache, enabling orders of magnitude lower latency accessing remotely stored data. By utilizing RDMA to access memory remotely, **XRMEM Data Accelerator bypasses the network and I/O stack, eliminating expensive CPU interrupts and context switches, and reducing latency by more than 10x**, from 200  $\mu$ s to less than 17  $\mu$ s. Smart Exadata System Software also ensures data is mirrored across storage servers, which provides additional fault tolerance. Exadata’s unique end-to-end integration between Oracle Database and Exadata Storage Servers automatically caches the hottest data blocks efficiently between the buffer cache in database servers and XRMEM and Flash Cache in storage servers. XRMEM is a shared storage tier across all of the storage nodes, which means the aggregate performance of this cache can be dynamically used by any database instance on any database server. This is a significant advantage over general-purpose storage architectures, which preclude sharing storage resources across database instances.

Real-world database workloads running on Exadata X10M, utilizing the shared XRMEM Data Accelerator, can achieve up to **25.2 million OLTP Read IOPS (8K IOs)<sup>1</sup>** in a single rack. This represents a 21% improvement over the same nine database and nine storage server configuration of the previous Exadata generation at 20.7 million<sup>2</sup>. This performance scales as additional racks are deployed.

Security and management of XRMEM are fully automated. XRMEM is configured automatically, with no user interaction required, and automatically managed thereafter. Hardware monitoring is pre-configured. Exadata RDMA

<sup>1</sup> Elastic configuration with 9x Exadata X10M Database Servers and 9x Exadata X10M Extreme Flash or 9x Exadata X10M High Capacity Storage Servers

<sup>2</sup> Elastic configuration with 9x Exadata X9M Database Servers and 9x Exadata X9M Extreme Flash Storage Servers or 9x High Capacity Storage Servers

4 Oracle Exadata Database Machine X10M / Version 1.2

Memory is only accessible to databases using database access controls, ensuring end-to-end security of data. XRMEM is entirely transparent to all applications.

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

Exadata **Extreme Flash (EF) Storage Server** is the foundation of a database-optimized all-flash Exadata Database Machine. Each EF Storage Server contains four capacity-optimized 30.72 TB flash drives for an aggregate, raw storage capacity of 122.88 TB. With the introduction of capacity-optimized flash, usable storage capacity<sup>3</sup> is now 2.4x<sup>4</sup> larger than prior generations.

In addition, each EF Storage Server includes four 6.8 TB performance-optimized flash drives, offering 27.2 TB of Exadata Smart Flash Cache. The **size of the Smart Flash Cache increases by 11.5x<sup>5</sup>** and is used to satisfy read and writes requests. Exadata delivers ultra-high performance by placing all the flash devices directly on the high-speed PCIe interface rather than behind slow disk controllers. Exadata EF Storage Servers include 1.25 TB of DDR5 Exadata RDMA Memory as a data acceleration tier in front of Flash Cache.

For data warehouse environments that require the highest performance, Exadata X10M Extreme Flash storage servers can achieve **up to 1 TB/s scan throughput<sup>6</sup>** from a single rack configuration.

These represent real-world, end-to-end performance metrics measured running SQL workloads with standard 8K database I/O sizes inside a single rack Exadata system. Exadata's performance on real Oracle Database workloads is orders of magnitude faster than traditional storage array architectures and much faster than current all-flash storage arrays.

## High Capacity Storage Server: Tiered Disk, Flash, and Exadata RDMA Memory Deliver Lower Cost of Disk Storage with Shared Memory Performance

The second Exadata storage option is the **High Capacity (HC) Storage Server**. This server includes twelve 22 TB disk drives with 264 TB total raw disk capacity. It also has four 6.8 TB performance-optimized flash drives for a total raw capacity of 27.2 TB of Exadata Smart Flash Cache and 1.25 TB of DDR5 XRMEM in front of flash to boost performance further.

Flash in the HC 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 and behind the Exadata RDMA Memory. Exadata Smart Flash Cache is used with the XRMEM Data Accelerator to automatically cache 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. Exadata uniquely understands database workloads and **knows when to avoid caching data that negatively affect overall performance**. For example, if large write I/Os caused by backups or large table scans are likely to disrupt higher-priority OLTP or scan operations, those large I/Os will bypass the flash cache and go straight to disk. Otherwise, Exadata System Software will utilize additional spare flash capacity and I/O bandwidth to optimize performance by caching these I/Os. Administrators can also manually provide SQL directives to ensure that specific tables, indexes, or partitions are preferentially retained in the flash cache.

It is common for hit rates in the Exadata Smart Flash Cache to be over 95%, or even 99% in real-world database workloads, yielding an effective flash capacity many times larger than the physical flash.

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 Exadata Database Machine X10M rack reaches **10.9 million 8K flash write I/O operations per second**

<sup>3</sup> Exadata X10M Extreme Flash Storage Servers with ASM High Redundancy mirroring and rebalance reservations applied.

<sup>4</sup> Exadata X7-X9M Extreme Flash Storage Servers allocated 48.9 TB raw capacity to data storage.

<sup>5</sup> Exadata X7 - X9M Extreme Flash Storage Servers allocated 2.32 TB to Flash Cache of the total 51.2 TB raw flash capacity per server.

<sup>6</sup> Elastic configuration with 2x Exadata X10M Database Servers and 17x Exadata X10M Extreme Flash Storage Servers

5 Oracle Exadata Database Machine X10M / Version 1.2



(IOPS)<sup>7</sup>. The Exadata Write-Back Flash Cache is transparent, persistent, and fully redundant, with performance comparable to dozens of enterprise disk arrays with thousands of disk drives.

The automatic data tiering between memory, flash, and disk in Exadata provides tremendous advantages over other flash-based solutions. Many storage vendors have developed flash-only arrays to achieve higher performance than traditional disk-based arrays. However, they cannot match the cost advantages of Exadata's smart data tiering between disk and flash. Traditional flash arrays lack Exadata's unique database-aware storage optimization. In addition, generic data deduplication provided by some flash arrays may be effective for workloads such as Virtual Desktop Infrastructure environments but are not for databases. In addition to utilizing its integrated and optimized hardware architecture, Exadata delivers superior performance by offloading data-intensive processing to unique algorithms in storage which have been specifically optimized for Oracle Database.

*“Our infrastructure needed to keep pace with exponential growth, quality and availability. The customer is the center of our strategy and one of our objectives is to always offer the best experience. With Oracle Exadata, we were able to improve our digital banking applications with 70% better performance and speed to enhance customer service.”*

**Everton Sims de Queiroz**  
Executive Infrastructure Manager  
Banco Original

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

The third Exadata storage option is the **Extended (XT) Storage Server**. Each Exadata XT Storage Server includes twelve 22 TB disk drives with 264 TB total raw disk capacity.

This storage option extends Exadata's operational and management benefits to rarely accessed data that must be kept online. Leveraging the same scale-out architecture as the HC and EF storage servers, XT storage simply and transparently expands capacity and integrates with Oracle Database security and data access controls.

With Exadata Extended (XT) Storage Server, enterprises 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 the information lifecycle across multiple platforms. XT Storage Servers include Hybrid Columnar Compression, while Exadata System Software may be optionally licensed to enable access to other smart features.

Combining Extreme Flash, High Capacity, and Extended Storage within an Exadata configuration allows customers to define a true Information Lifecycle Management policy. As data ages, it can be moved between the three storage tiers to ensure data is on the correct medium for its usage and retention requirements. Coupled with Automatic Data Optimization, part of Oracle Advanced Compression, customers can define policies to automate this movement between compression levels.

## Accelerating Database Processing with Smart System Software

As data volumes continue to grow, 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 tens to hundreds of gigabytes a second, far faster than conventional storage arrays can deliver data through their storage controllers and the storage network.

**Exadata System Software** enables Exadata's unparalleled performance by implementing a unique, highly efficient, database-optimized storage infrastructure on the Exadata Storage Server. Each storage server has CPUs used to

<sup>7</sup> Elastic configuration with 6x Exadata X10M Database Servers and 12x Exadata X10M Extreme Flash Storage Servers

offload database processing. These CPUs in the storage servers do not replace database CPUs. They 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 directly relevant to a query are sent to the database servers.**

For example, suppose a query is executed to identify the customers who placed sales orders over \$1000 in the month of March. Exadata will offload the scanning of the table to the Exadata Storage Servers, where filters extract only the relevant customer information for March with a minimum \$1000 spend and return this reduced quantity of data to the database. Offloading reduces the amount of data transferred to the database servers by orders of magnitude. Smart Scan greatly accelerates query execution, eliminates bottlenecks, and significantly reduces the CPU usage of the database servers.

**Storage Indexes** are another powerful and **unique** capability of Oracle Exadata that helps avoid unnecessary I/O operations and improves overall performance. Storage Indexes are maintained automatically in the storage server's memory and track minimum and maximum column values for table data contained in a storage region on that storage server. When a query specifies a WHERE clause, Exadata System Software examines the storage index to determine where rows with the specified column value exist in a disk region on the storage server. Rather than reading all the rows to satisfy the query and discarding the rows that do not match the WHERE clause, only the regions of the disk containing the rows matching the WHERE clause are read, avoiding I/Os for rows that would otherwise have been discarded. Storage Indexes make many SQL operations run dramatically faster because a few in-memory lookups automatically replace large numbers of I/O operations. Storage Indexes are automatically persisted to disk, avoiding the need to rebuild them and the associated consumption of extra I/O while ensuring consistent performance after planned or unplanned downtime.

The time it takes to commit user transactions or perform critical updates is sensitive to the latency of log writes. To accelerate OLTP workloads, the Exadata Smart Flash Cache implements unique algorithms to ensure consistent low latency of database log writes. **Exadata Smart Flash Log Write-Back** and **Exadata Smart Flash Log**, eliminate the storage disks as a potential log write throughput bottleneck, prevent log write latency outliers, and automatically and transparently store Oracle Database Redo Logs in the Smart Flash Cache. Smart Flash Log Write-Back **increases log write throughput by up to 2.5x**. Exadata uniquely prioritizes latency-sensitive I/O, such as log writes, over other I/O requests in the RoCE network and within the Exadata Storage Servers to further ensure other workloads do not impact mission-critical OLTP workloads.

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 ultra-fast RoCE network using **Remote Direct Memory Access (RDMA)**, bypassing the OS, kernel, and networking software overhead. Using RDMA 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.

Exadata **uniquely** uses Machine Learning to implement **Automatic Indexing with Oracle Database 19c** and later releases. Automatic Indexing continually analyzes SQL execution plans and creates new indexes to accelerate performance. 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.

After implementing the Oracle Exadata system, our client services saw performance improvements of 300% in data processing and 200% in data warehousing, while achieving zero downtime and zero data loss.

**Chang Rea Han**  
Vice President and CIO  
KCB

## 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 range from 5x to 20x. Average storage savings is an industry-leading 10x. Exadata Database Machine can offload decompression operations to processors in Exadata storage. As a result, there is reduced I/O because of the high compression achieved. Most analytic workloads, therefore, run faster using Hybrid Columnar Compression than without it.

Two modes of Hybrid Columnar Compression are available. **Warehouse compression** mode suits read-intensive workloads and provides large storage savings and enhanced analytic performance. **Archive compression** mode offers the highest degree of compression and targets data that is seldom accessed but must remain online. This data can be seamlessly stored on the XT storage server for further cost reduction.

OLTP systems can use Hybrid Columnar Compression to compress older, less active data while newer, more active and update-intensive data can be compressed using Advanced Row Compression. Oracle Database 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.

Exadata 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 Cache 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 I/O 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 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 storage tiers, 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 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.

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 the in-memory pool in the SGA to flash cache capacity in storage servers. An Exadata X10M Database Machine<sup>8</sup> has up to 462 TB of Flash Cache, capable of servicing some of the largest in-memory workloads. Exadata X10M utilizes a unique new algorithm to increase the compression of data stored in the **In-Memory columnar format by up to 1.25x, further increasing the effective Flash Cache capacity.** 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 non-Exadata cluster configuration, when a database node fails, the in-memory data on that node is lost. 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 may 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.

Finally, Oracle Database 19c and later enable the use of Database In-Memory Caching in Storage Servers without allocating memory to the Database In-Memory Column Store on database servers. By setting the `inmemory_force` parameter to 'CELLMEMORY\_LEVEL', databases can continue leveraging the optimizations and vector processing benefits of Database In-Memory and reaping the processing benefits of the shared storage tier and freeing up database server memory.

## Database Consolidation on Exadata

The Exadata Database Machine can host many databases, enabling 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 mixed database workloads with industry-leading scalability and performance** makes it an ideal consolidation platform.

Consolidated environments running on Exadata X10M may also use KVM-based Virtual Machines (Guests) and **Secure RDMA Fabric Isolation** for strong isolation between workloads. Isolation is critical in hosted, shared, service provider, and test/dev environments. When using virtualization, Exadata can safely deploy multiple RAC clusters running the same or different Exadata software, grid infrastructure, or database versions on the same set of database servers.

**Exadata Database Machine is the world's fastest virtualized Oracle database platform.** Exadata virtual machines use high-speed 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 significantly 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

---

<sup>8</sup> Elastic rack with 2x Exadata X10M Database Servers and 17x Exadata X10M Extreme Flash Storage Servers

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.

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** of an application workload's use of database CPU, memory, network, and storage. Workload 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 are applied when multiple databases and virtual clusters are deployed on one Exadata rack, as is typical in a consolidated private cloud.

In X10M, Exadata utilizes RDMA over Converged Ethernet protocols to ensure network-intensive workloads such as reporting, batch, and backups don't stall latency-sensitive interactive workloads. Latency-sensitive network operations, such as RAC Cache Fusion communication and log file writes, travel across high-priority network channels within the converged ethernet fabric. Non-latency-sensitive traffic travels on other channels with their own network switch buffers.

Due to Exadata's unique database consolidation and Database-as-a-Service capabilities, Exadata is the **only** platform supporting up to 4000 Pluggable Databases within a single Oracle Multitenant Container Database.

*“Our investment in Oracle Exadata has exceeded our expectations. Not only did we address our biggest challenge of increasing performance for the Siebel CRM platform by an average of 40% and 4-5X on large queries, we also reaped the benefits of consolidating all of our databases onto Exadata, reduced our 287 databases by half, lowered administration, improved uptime, and saved 50% of our legacy costs.”*

**Greg Ogle**  
Vice President  
Global IT Infrastructure  
Equinix

## **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 data is compressed, so the cost of decryption is decreased by the degree of compression**. By leveraging both technologies, Exadata can query fully encrypted and compressed databases with minimal overhead at hundreds of gigabytes of (original) user data per second. Oracle Transparent Data Encryption provides a complete key management solution to keep all data encrypted and secure.

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 ensuring the security of each software and hardware component and ensuring that security is maintained across the entire stack. **Oracle delivers full stack security in the Exadata Database Machine.**

Exadata virtual machines provide an added isolation layer at the operating system level. Additionally, in environments that leverage virtualization on Exadata, Exadata **Secure RDMA Fabric Isolation** ensures VM Guests in one cluster cannot communicate directly with other clusters on the same Exadata while still providing access to shared Exadata

storage. Such isolation is beneficial in consolidated environments where, for example, different organizational divisions share infrastructure and have various data security requirements.

Exadata systems are designed, manufactured, and delivered to customers using a defense-in-depth approach, increasing the system's security posture. Exadata systems are built using Oracle-designed database and storage servers. The in-house design and development of the servers enable the implementation of features unique to Exadata and enable tight control over the security of the design. This focus on security extends to the global supply chain of Oracle. 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. These signatures are verified every server reboot, **preventing the execution of malware hidden in embedded code** in the boot chain. The operating system installed on Exadata systems is a pared-down version of the standard Oracle Linux distribution, with Unbreakable Enterprise Kernel, unique to Exadata systems. This nano-kernel only includes packages required to run the Oracle Database and eliminates unnecessary packages, **minimizing the attack surface** and hardening the system's security. **Exadata Live Update** on the database servers leverages familiar Linux technologies – Ksplice and RPM – of Oracle Linux to **apply system software and security updates while the OS stays online**.

In addition, and complimentary to database encryption provided by Transparent Data Encryption, the disk and flash technologies used in Exadata X10M enable Stored Data Encryption to eliminate data leakage risk during proactive drive replacement or machine redeployment. In Stored Data Encryption, each disk and flash storage device encrypts 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. Because the previous drive encryption key is deleted with Secure Erase, 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 by Oracle Security experts and hundreds of industry security experts worldwide.

## 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 components to complex site failures and human errors. Each Exadata Database Machine has **completely redundant hardware**, including redundant 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.

Exadata in an MAA configuration is recognized by the analyst firm IDC as a system that **delivers at least 5-nines availability (99.999%)** and is categorized in the IDC AL4 fault-tolerant market segment<sup>9</sup>.

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 Failure Detection**. On non-Exadata platforms, detecting a server failure requires waiting for a long timeout, leading to

<sup>9</sup> Worldwide AL4 Server Market Shares, 2019: Fault-Tolerant Systems Become Digital Transformation Platforms, IDC, Paul Maguranis Peter Rutten, July 2020

extended application brownouts. RoCE-based Exadata Database Machines implement a unique RDMA-based **sub-second node failure detection**, leading to the virtual elimination of application brownout conditions.

Disk and flash devices occasionally exhibit high latency I/O operations due to internal recovery of failed sectors, internal firmware reboots, or wear leveling. These extended 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, I/O Latency Capping automatically redirects high latency write I/O operations to a healthy flash device, eliminating outliers during write operations. Exadata System Software uses machine learning techniques to predict components susceptible to failure and takes proactive action to take such components out of service gracefully. If disks fail, ASM performs a rebalance operation for the data resident on the disk, while applications continue to access the database with no interruption. Exadata allows hot-swapping of disks, fans, power supplies, and PCIe flash cards to avoid downtime. Exadata System software takes rebalance 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 the monitoring of CPU, memory, input/output subsystems, file system, and network. This automation combines machine learning techniques with the lessons learned from thousands of mission-critical real-world deployments. For example, Exadata can detect anomalous use of system resources that negatively impacts database performance and automatically identifies the process responsible, and issues an alert – all without any manual intervention.

As a result 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, web-based retailing, and many more. Exadata's mission-critical availability capabilities are not restricted to OLTP workloads; they also apply to data warehousing and analytics.

## **Fast Deployment of Development and Test Databases with Exadata Snapshots**

Administrators can quickly create space-efficient database snapshots directly on Exadata for test and development purposes. Exadata database snapshots are integrated with Oracle Multitenant to provide a straightforward interface for creating new pluggable database (PDB) snapshots.

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

Exadata Snapshots may also be created using a read-only Container Database (CDB) and all its PDBs, or a non-CDB architecture database. Oracle Data Guard enables the creation of point-in-time, read-only parents, from a primary database using redo to minimize data transfer.

All Exadata-specific features such as Smart Scan, Exadata RDMA Memory Data Accelerator, resource management, and Smart Flash Cache work seamlessly on database instances created via Exadata snapshots, providing an exact test and development environment while using a fraction of valuable storage resources. RMAN backups of snapshots on Exadata are also space efficient, with only the changed blocks included.

## 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, and network 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 the root cause of any performance bottlenecks.
- Lights-out monitoring within Enterprise Manager is optimized for Exadata with predefined metrics and thresholds so administrators receive timely notifications when issues arise and manage those exceptions. Hardware incidents are automatically detected, and service requests are logged to reduce problem resolution time.
- The Exachk tool, 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 find deviations in configuration across the racks or among the database servers of a rack. Exachk is a component of the Autonomous Health Framework (AHF). AHF issues early warnings or automatically solves operational runtime issues faced by Database and System administrators in availability and performance.
- Exadata's built-in Management Server (MS) processes constantly monitor the health of hardware and software components and send alerts to administrators and Oracle support when faulty components are detected.
- Exadata Real-time Insights streams fine-grained performance data directly from the Management Server (MS) processes on all Exadata servers to power real-time performance dashboards, enabling DBAs to monitor performance at a fleet level with per-second level accuracy.

## Highest Level of Service

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

**Oracle Platinum Services** 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 covers all software and hardware within an Engineered System, including the Oracle Database – the highest level of support ever for a full-stack software/hardware platform. 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 networking. Management of a traditional database system is typically spread across the teams of each component, such as the database team, the storage team, and the system administrators. In contrast, an **Exadata system is typically managed by a unified Database Machine Administration (DMA) team**. Database Machine Administrators have complete control of all resources in the Exadata Database Machine, including storage resources. Database Machine Administrators can implement new database deployments and configuration changes 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 low-level configuration issues.



## Dramatically Lower Costs

Due to 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. In sizing exercises, it is typical to see a 2x-4x reduction in Exadata system size compared to a traditional system.

Exadata provides a huge memory, flash, and disk footprint for large data sets. Raw disk storage on an Exadata system<sup>10</sup> can reach 4.2 PB (Petabytes), while raw flash storage can be up to 2 PB. Hybrid Columnar Compression may also increase the effective storage and memory capacity by an average factor of 10. By intelligently moving active data across disk, flash, and memory tiers, Exadata simultaneously delivers the highest performance and the lowest cost.

Exadata can **uniquely consolidate many databases supporting multiple workloads** in a single cloud platform. High-end OLTP, analytics, batch, reporting, and backups can run simultaneously within and across databases with extreme performance. **Exadata's extreme performance and capacity enable many databases and workloads to be consolidated on Exadata.** Consolidating databases on Exadata reduces system hardware and software costs and ongoing operations costs.

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 less labor than traditional systems. Low-level tuning, integration, and maintenance is reduced or eliminated. All Exadata users run a configuration that is identical to thousands of other users and are identical to Oracle's internal configurations, making it far less likely that issues will be encountered. When issues occur, the resolution is simpler as customers work with one supplier – Oracle, as the entire system – hardware, firmware, operating system, hypervisor, and database layers are all owned and supported by Oracle. The “one-hand-to-shake” support model enables faster problem resolution times and reduces downtime and associated costs, further increasing economic benefits.

## Capacity-on-Demand Software Licensing

An Exadata X10M Database Server has a substantial compute capacity with two 96-core x86 processors (192 cores per database server). The Capacity-on-Demand feature allows a subset of cores (minimum of 14) in each database server to be enabled during the hardware installation. As your workload grows and more cores are needed, Capacity-on-Demand can be used to increase CPU resources in 2-core increments. Since software licenses are only required for the enabled cores, this pay-as-you-grow approach to software licensing is another way Exadata helps align costs with business growth.

## Exadata in Oracle Cloud

Customers can run Oracle databases on Exadata in Oracle Cloud, with the same extreme performance and availability experienced by thousands of organizations running Exadata on-premises. Exadata in Oracle Cloud combines the world's #1 database – Oracle, and the most powerful database platform – Exadata with the simplicity, automation, operations, and economics of the cloud, both in the public cloud (Oracle Cloud Infrastructure) and the hybrid cloud (Exadata Cloud@Customer).

With Exadata in Oracle Cloud:

- Compute resources can be online scaled up and down enabling customers to pay for only what they use
- Oracle Cloud operations manage all infrastructure, eliminating many administration activities previously performed by customer staff

---

<sup>10</sup> Exadata X10M elastic configuration with 2x X10M Database Servers and 16x X10M HC or 16c X10M EF Storage Servers

- Powerful cloud automation exposed through a browser-based UI and REST APIs simplifies common lifecycle management tasks
- All Exadata System software and hardware is included with the infrastructure subscription

Customers can choose between bringing their own on-premises Oracle Database licenses, or subscribe to an inclusive license for **all Oracle Database options and features** – such as Oracle Multitenant, In-Memory Database, Real Application Clusters (Oracle RAC), Active Data Guard, Partitioning, Advanced Compression, and Advanced Security. Exadata in Oracle Cloud also includes all Oracle Enterprise Manager packs. Oracle databases deployed on Exadata in Oracle Cloud are **100% compatible** with those deployed on-premises, ensuring a smooth transition to the cloud, and a seamless hybrid cloud strategy. With elastic Exadata configurations, pay-per-use Oracle Database licensing, and infrastructure managed by Oracle experts, Exadata in Oracle Cloud enables business agility and operational flexibility with **zero CapEx**. Exadata is available as cloud infrastructure in the Oracle public cloud with Exadata Cloud Infrastructure or in customer data centers with Exadata Cloud@Customer. A single Exadata Cloud Infrastructure or Exadata Cloud@Customer can run both – Oracle Autonomous Database and Oracle Exadata Database Service – the most advanced database services available.

Exadata in Oracle Cloud 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
- Reducing costs when running variable workloads whose resource requirements change over time.
- Consolidating a variety of workloads in the cloud using multiple Oracle databases or Oracle Multitenant
- Easy provisioning of Oracle standby or replica databases for disaster recovery 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, and proof-of-concept

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 IT footprint simply expands to include the elasticity and flexibility of Exadata in Oracle Cloud. They also do not have to invest in multiple database cloud platforms for multiple workloads since Exadata provides a unified platform for all workloads – analytics, data warehousing, OLTP, consolidation, in-memory, and mixed-workloads.

Customers who are not able to migrate their databases to Exadata in one of Oracle's public cloud regions can run Exadata in Oracle Cloud in their own data center with Exadata Cloud@Customer to satisfy security, compliance, and data residency requirements. Since the data is stored in customer data centers next to existing on-premises applications, customers can also easily maintain their current system dependencies.

**Uniquely** engineered for extreme performance for all workloads, along with fast deployment, simplified management, low operating costs and reduced risks, Exadata in Oracle Cloud is the best cloud database platform.

## **Exadata Business Benefits**

Beyond the operational benefits of extreme performance, availability, security, and deployment flexibility across on-premises and cloud, Exadata also directly benefits the business bottom line.

**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 risks of unexpected system issues after go-live are 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, large memory, and flash capacity enhance employee productivity and customer satisfaction by significantly improving user response times. **Users spend more time doing valuable 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

*“It is not an exaggeration to say that Oracle Exadata has been the living proof and the most important companion in Hyundai Home Shopping’s digital innovation journey. Sales, revenue, and operating profit margins have all grown significantly.”*

**Bae-hyun Kim**  
Team Leader, Security and Infrastructure  
Hyundai Home Information, Hyundai IT&E

## Conclusion

Exadata delivers a fully integrated database platform with the latest hardware technologies and **unique** software to deliver extreme performance, availability, and security. 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 – on-premises or in the cloud.

## Exadata Server Hardware 1,2

Server Type	CPU	Memory (DDR5)	Disk	Flash	Network
<b>Database Server</b>	2 x 96-core AMD EPYC™ 9J14 processors, 2.6 GHz (up to 3.7 GHz)	512 GB (factory option) 1,536 GB (factory option and field upgrade) 2,304 GB (factory option and field upgrade) 3,072 GB (factory option and field upgrade, max)	None	2 x 3.84 TB NVMe Flash SSD (hot swappable), (upgradeable to 4 x 3.84 TB)	<ul style="list-style-type: none"> <li>Client/backup adapter 1: 2 x 10/25 Gb Ethernet ports (SFP28)</li> <li>Client/backup adapter 2: 2 x 10/25 Gb Ethernet ports (SFP28)</li> <li>Client/backup adapter 3, 4, or 5:                             <ul style="list-style-type: none"> <li>4 x 10 Gb Ethernet ports (RJ45), or</li> <li>2 x 10/25 Gb Ethernet ports (SFP28), or</li> <li>2 x 100 Gb optical Ethernet ports (QSFP28)</li> </ul> </li> <li>1 x 1 Gb Ethernet port (RJ45, management)</li> <li>1 x 1 Gb Ethernet port (RJ45, ILOM)</li> <li>2 x 100 Gb QSFP28 RoCE Fabric ports</li> </ul>
<b>Database Server Eighth Rack<sup>2</sup></b>	1 x 32-core AMD EPYC™ 9334 processors 2.7 GHz (up to 3.9 GHz)	384 GB (factory option) 768 GB (factory option and field upgrade) 1,152 GB (factory option and field upgrade, max)	None	2 x 3.84 TB NVMe Flash SSD (hot swappable), (upgradeable to 4 x 3.84 TB)	<ul style="list-style-type: none"> <li>Client/backup adapter 1: 2 x 10/25 Gb Ethernet ports (SFP28)</li> <li>Client/backup adapter 2: 2 x 10/25 Gb Ethernet ports (SFP28)</li> <li>Client/backup adapter 3:                             <ul style="list-style-type: none"> <li>4 x 10 Gb Ethernet ports (RJ45), or</li> <li>2 x 10/25 Gb Ethernet ports (SFP28), or</li> <li>2 x 100 Gb optical Ethernet ports (QSFP28)</li> </ul> </li> <li>1 x 1 Gb Ethernet port (RJ45, management)</li> <li>1 x 1 Gb Ethernet port (RJ45, ILOM)</li> <li>2 x 100 Gb QSFP28 RoCE Fabric ports</li> </ul>
<b>Storage Server High Capacity (HC)</b>	2 x 32-core AMD EPYC™ 9334 processors 2.7 GHz (up to 3.9 GHz)	256 GB 1.25 TB XRMEM	12 x 22 TB 7,200 RPM disks	4 x 6.8 TB NVMe PCIe4.0 performance-optimized Flash cards	<ul style="list-style-type: none"> <li>2 x 100 Gb QSFP28 RoCE Fabric ports</li> <li>1 x 1 Gb Ethernet port (RJ45, management)</li> <li>1 x 1 Gb Ethernet port (RJ45, ILOM)</li> </ul>
<b>Storage Server Extreme Flash (EF)</b>	2 x 32-core AMD EPYC™ 9334 processors	256 GB 1.25 TB XRMEM	None	4 x 6.8 TB NVMe PCIe4.0 performance-optimized Flash cards, and	<ul style="list-style-type: none"> <li>2 x 100 Gb QSFP28 RoCE Fabric ports</li> <li>1 x 1 Gb Ethernet port (RJ45, management)</li> <li>1 x 1 Gb Ethernet port (RJ45, ILOM)</li> </ul>

	2.7 GHz (up to 3.9 GHz)			4 x 30.72 TB NVMe PCIe4.0 capacity- optimized Flash cards	
<b>Storage Server Extended (XT)</b>	1 x 32-core AMD EPYC™ 9334 processors 2.7 GHz (up to 3.9 GHz)	128 GB	12 x 22 TB 7,200 RPM disks	None	<ul style="list-style-type: none"> <li>• 2 x 100 Gb QSFP28 RoCE Fabric ports</li> <li>• 1 x 1 Gb Ethernet port (RJ45, management)</li> <li>• 1 x 1 Gb Ethernet port (RJ45, ILOM)</li> </ul>
<b>Storage Server Eighth Rack High Capacity (HC)</b>	1 x 32-core AMD EPYC™ 9334 processor 2.7 GHz (up to 3.9 GHz)	192 GB 576 GB XRMEM	6 x 22 TB 7,200 RPM disks	2 x 6.8 TB NVMe PCIe4.0 Flash cards	<ul style="list-style-type: none"> <li>• 2 x 100 Gb QSFP28 RoCE Fabric ports</li> <li>• 1 x 1 Gb Ethernet port (RJ45, management)</li> <li>• 1 x 1 Gb Ethernet port (RJ45, ILOM)</li> </ul>

<sup>1</sup> All servers include redundant hot swappable fans and power supplies

<sup>2</sup> Table includes servers available for individual purchase only. Eighth Rack Database Servers are available only in the Eighth Rack configuration shown below and not individually purchasable.

## Exadata Rack Configurations <sup>1, 2</sup>

Rack Size	Database Servers and Cores	Storage Servers and Cores	High Capacity Storage Server Capacity (Raw)	Extreme Flash Storage Server Capacity (Raw)
<b>Eighth Rack<sup>3</sup></b>	2 x servers, 64 cores	3 x servers, 96 cores for SQL offload	<ul style="list-style-type: none"> <li>• 396 TB disk</li> <li>• 40.8 TB performance-optimized Flash,</li> <li>• 1.69 TB Exadata RDMA Memory</li> </ul>	<ul style="list-style-type: none"> <li>• N/A</li> </ul>
<b>Quarter Rack</b>	2 x servers, 384 cores	3 x servers, 192 cores for SQL offload	<ul style="list-style-type: none"> <li>• 792 TB disk</li> <li>• 81.6 TB performance-optimized Flash</li> <li>• 3.75 TB Exadata RDMA Memory</li> </ul>	or <ul style="list-style-type: none"> <li>• 368.6 TB capacity-optimized Flash</li> <li>• 81.6 TB performance-optimized Flash</li> <li>• 3.75 TB Exadata RDMA Memory</li> </ul>
<b>Elastic Configuration (Half rack)<sup>5</sup></b>	4 x servers, 768 cores	7 x servers, 448 cores for SQL offload	<ul style="list-style-type: none"> <li>• 1,848 TB disk</li> <li>• 190.4 TB performance-optimized Flash</li> <li>• 8.75 TB Exadata RDMA Memory</li> </ul>	or <ul style="list-style-type: none"> <li>• 860.2 TB capacity-optimized Flash</li> <li>• 190.4 TB performance-optimized Flash</li> <li>• 8.75 TB Exadata RDMA Memory</li> </ul>
<b>Elastic Configuration 1 (Example)<sup>3</sup></b>	9 x servers, 1,728 cores	9 x servers, 576 cores for SQL offload	<ul style="list-style-type: none"> <li>• 2,376 TB disk</li> <li>• 244.8 TB performance-optimized Flash</li> <li>• 11.25 TB Exadata RDMA Memory</li> </ul>	or <ul style="list-style-type: none"> <li>• 1,105.9 TB capacity-optimized Flash</li> <li>• 244.8 TB performance-optimized Flash</li> <li>• 11.25 TB Exadata RDMA Memory</li> </ul>
<b>Elastic Configuration 2(Example)<sup>5</sup></b>	2 x servers, 384 cores	16 x HC servers, 1,024 cores for SQL offload,	<ul style="list-style-type: none"> <li>• 4,224 TB disk</li> <li>• 435.2 TB performance-optimized Flash</li> <li>• 20 TB Exadata RDMA Memory</li> </ul>	or <ul style="list-style-type: none"> <li>• 2,089 TB capacity-optimized Flash</li> </ul>



		or 17 x EF servers, 1,088 cores for SQL offload			<ul style="list-style-type: none"> <li>• 462.4 TB performance-optimized Flash</li> <li>• 21.25 TB Exadata RDMA Memory</li> </ul>
<b>+Database Servers</b>	Up to 15x servers <sup>4</sup> , 2,880 cores max per rack	N/A	N/A		N/A
<b>+Storage Servers</b>	N/A	Maximum per rack:  Up to 16x HC servers <sup>4</sup> , 1,024 cores,  or  Up to 17x EF servers <sup>4</sup> , 1,088 cores	Maximum per rack:  <ul style="list-style-type: none"> <li>• 4,224 TB disk</li> <li>• 435.2 TB performance-optimized Flash</li> <li>• 20 TB Exadata RDMA Memory</li> </ul>	or	Maximum per rack:  <ul style="list-style-type: none"> <li>• 2,089 TB capacity-optimized Flash</li> <li>• 462.4 TB performance-optimized Flash</li> <li>• 21.5 TB Exadata RDMA Memory maximum per rack</li> </ul>

<sup>1</sup> Each rack is 42 RU (Rack Units) in height, has 2x redundant Power Distribution Units (PDUs), 2x 36-port 100 Gb/s RoCE switches and 1x 48-port Management Ethernet switch for administration.

<sup>2</sup> 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. An elastic configuration cannot exceed 19 servers and 38 RU (Rack Units) per rack. Database Servers = 2 RU, Storage Servers = 2 RU

<sup>3</sup> Elastic Configuration (Half Rack), Elastic Configuration 1 and Elastic Configuration 2 configurations added as examples of elastic configurations.

<sup>4</sup> Maximum number of database servers allowed in an elastic configuration is 15. Maximum number of storage servers allowed in an elastic configuration is 17.

## Other Elastic Expansion Options

<b>Multi-Rack Connection</b>	<p>Connect any combination of up to 14 Exadata Database Machine racks or Exadata Storage Expansion Racks via the RoCE Network Fabric.</p> <ul style="list-style-type: none"> <li>• Larger configurations can be built with external RoCE switches.</li> <li>• Connected racks must contain Exadata RoCE hardware.</li> </ul>
<b>Eighth Rack Expansion Options</b>	<p>Expand just compute or just storage or both, described as follows:</p> <ul style="list-style-type: none"> <li>• Eighth Rack Database Servers can be expanded by replacing the 32 core CPU with two 96 core CPU and installing an additional 128 GB Memory per server.</li> </ul> <p>Storage can be expanded by adding additional Eighth Rack High Capacity, High Capacity (HC), Extreme Flash (EF) and/or Extended (XT) Storage Servers.</p>

## Exadata Capacity and Performance Metrics: Individual Servers

Server Type	Maximum SQL Flash Bandwidth <sup>2</sup>	Maximum SQL Read IOPS <sup>1,3</sup>	Maximum SQL Write IOPS <sup>4</sup>	Exadata RDMA Memory Capacity (Raw) <sup>5</sup>	Performance-Optimized PCI Flash Capacity (Raw) <sup>5</sup>	Capacity-Optimized PCI Flash Capacity (Raw) <sup>5</sup>	Disk Data Capacity (Raw) <sup>5</sup>
<b>Database Server</b>	N/A	2,800,000	2,000,000	N/A	N/A	N/A	7.68 TB
<b>Storage Server High Capacity (HC)<sup>1</sup></b>	45 GB/s	2,800,000	916,000	1.25 TB	27.2 TB	N/A	264 TB
<b>Storage Server Extreme Flash (EF)<sup>1</sup></b>	60 GB/s	2,800,000	916,000	1.25 TB	27.2 TB	122.88 TB	N/A
<b>Storage Server Extended (XT)<sup>1</sup></b>	N/A	N/A	N/A	N/A	N/A	N/A	264 TB
<b>Storage Server Eighth Rack High Capacity (HC)<sup>1</sup></b>	22.5 GB/s	1,400,000	458,000	576 GB	13.6 TB	N/A	132 TB

<sup>1</sup> Actual system performance varies by application.

<sup>2</sup> Bandwidth is peak physical scan bandwidth achieved running SQL, assuming no database compression. Effective user data bandwidth is higher when using database compression.

<sup>3</sup> Based 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.

<sup>4</sup> Based 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.

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

## Exadata Elastic Rack Configurations: Flash Capacity and Performance Metrics (HC and EF)

Flash Metrics		Maximum SQL Flash Bandwidth <sup>2</sup>	Maximum SQL XRMEM Read IOPS <sup>1,3</sup>	Maximum SQL Flash Write IOPS <sup>4</sup>	Performance-Optimized PCI Flash Capacity (Raw) <sup>5</sup>
<b>Elastic Configuration 1 (Example)<sup>6</sup></b>	HC <sup>1</sup>	405 GB/s	25,200,000	8,244,000	244.8 TB
	EF <sup>1</sup>	540 GB/s	25,200,000	8,244,000	244.8 TB
<b>Elastic Configuration 2 (Example)<sup>6</sup></b>	HC <sup>1</sup>	720 GB/s	5,600,000	4,000,000	435.2 TB
	EF <sup>1</sup>	1,020 GB/s	5,600,000	4,000,000	462.4 TB
<b>Elastic Configuration (Half Rack)<sup>6</sup></b>	HC <sup>1</sup>	315 GB/s	11,200,000	6,412,000	190.4 TB
	EF <sup>1</sup>	420 GB/s	11,200,000	6,412,000	190.4 TB
<b>Quarter Rack</b>	HC <sup>1</sup>	135 GB/s	5,600,000	2,748,000	81.6 TB
	EF <sup>1</sup>	180 GB/s	5,600,000	2,748,000	81.6 TB
<b>Eighth Rack</b>	HC <sup>1</sup>	67.5 GB/s	2,800,000	1,374,000	40.8 TB

<sup>1</sup> EF = Extreme Flash; HC = High Capacity

<sup>2</sup> Bandwidth is peak physical scan bandwidth achieved running SQL, assuming no database compression. Effective user data bandwidth is higher when using database compression.

<sup>3</sup> 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.

<sup>4</sup> 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.

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

<sup>6</sup> Elastic Configuration (Half Rack), Elastic Configuration 1 and Elastic Configuration 2 configurations added as examples of elastic configurations. Half = 4x DB Servers, 7x Storage Servers; Elastic Configuration 1 = 9x DB and 9x Storage Servers; Elastic Configuration 2 = 2x DB and 16x HC or 17x EF Storage Servers

## Exadata Elastic Rack Configurations: Disk Capacity and Performance Metrics (HC)

Disk Metrics	Maximum SQL Disk Bandwidth <sup>1</sup>	Maximum SQL Disk IOPS <sup>2</sup>	Data Capacity (Raw) <sup>3</sup>
<b>Elastic Configuration 1 (Example)<sup>4</sup></b>	16 GB/s	23,000	2,376 TB
<b>Elastic Configuration 2 (Example)<sup>4</sup></b>	28.8 GB/s	41,600	4,224 TB
<b>Elastic Configuration (Half Rack)<sup>4</sup></b>	12.5 GB/s	18,000	1,848 TB
<b>Quarter Rack</b>	5.4 GB/s	7,800	792 TB
<b>Eighth Rack</b>	2.7 GB/s	3,900	396 TB

<sup>1</sup> Bandwidth is peak physical scan bandwidth achieved running SQL, assuming no database compression. Effective user data bandwidth is higher when database compression is used.

<sup>2</sup> 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.

<sup>3</sup> Raw capacity is measured in standard disk drive terminology with 1 GB = 1 billion bytes.

<sup>4</sup> Elastic Configuration (Half Rack), Elastic Configuration 1 and Elastic Configuration 2 configurations added as examples of elastic configurations. Half = 4x DB Servers, 7x Storage Servers; Elastic Configuration 1 = 9x DB and 9x Storage Servers; Elastic Configuration 2 = 2x DB and 16x HC

## Exadata Elastic Rack Configurations: Combined Metrics (HC and EF)

Combined Metrics		Data Capacity (Usable) Normal Redundancy <sup>1</sup>	Data Capacity (Usable) High Redundancy <sup>1</sup>	Maximum Data Load Rate <sup>2</sup>
<b>Elastic Configuration 1 (Example)<sup>3</sup></b>	HC <sup>1</sup>	983.3 TB	720.4 TB	22.5 TB/hour
	EF <sup>1</sup>	457.6 TB	335.2 TB	22.5 TB/hour
<b>Elastic Configuration 2 (Example)<sup>3</sup></b>	HC <sup>1</sup>	1,748.1 TB	1,280.6 TB	8.8 TB/hour
	EF <sup>1</sup>	864.3 TB	633.2 TB	8.8 TB/hour
<b>Elastic Configuration (Half Rack)<sup>3</sup></b>	HC <sup>1</sup>	764.7 TB	560.2 TB	17.5 TB/hour
	EF <sup>1</sup>	355.9 TB	260.7 TB	17.5 TB/hour
<b>Quarter Rack</b>	HC <sup>1</sup>	306.1 TB	240.1 TB	7.5 TB/hour
	EF <sup>1</sup>	142.5 TB	111.7 TB	7.5 TB/hour
<b>Eighth Rack</b>	HC <sup>1</sup>	153.1 TB	120 TB	3.8 TB/hour

<sup>1</sup> 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. Normal redundancy calculations reflect the use of Grid Infrastructure version 12.2.0.1 or later.

<sup>2</sup> Load rates are typically limited by database server CPU, not I/O. Rates vary based on load method, indexes, data types, compression, and partitioning.

<sup>3</sup> Elastic Configuration (Half Rack), Elastic Configuration 1 and Elastic Configuration 2 configurations added as examples of elastic configurations. Half = 4x DB Servers, 7x Storage Servers; Elastic Configuration 1 = 9x DB and 9x Storage Servers; Elastic Configuration 2 = 2x DB and 16x HC or 17x EF Storage Servers

## Exadata Database Machine Component Environmental Specifications

Metric	Exadata Database Server X10M	Exadata Database Server X10M (3TB Memory)	Exadata Storage Server X10M High Capacity (HC)	Exadata Storage Server X10M Extreme Flash (EF)	Exadata Storage Server X10M Extended (XT)	Exadata Eighth Rack Storage Server X10M High Capacity (HC)
Height	3.42" ( 86.9 mm )					
Width	17.52 " ( 445.0 mm )					
Depth	30.51" ( 775.0 mm )					
Acoustic Noise (operating)	8.4 B	8.6 B	8.4 B	8.2 B	8.1 B	8.4 B
Weight	52 lb. ( 23.6 kg )	53 lb. ( 24 kg )	74 lb. ( 33.6 kg )	60 lb. ( 27.2 kg )	68 lb. ( 30.8 kg )	63 lb. ( 28.6 kg )
Maximum Power Usage	1.3 kW ( 1.3 kVA )	1.5 kW ( 1.5 kVA )	1 kW ( 1.1 kVA )	1 kW ( 1 kVA )	0.4 kW ( 0.5 kVA )	0.8 kW ( 0.8 kVA )
Typical Power Usage <sup>1</sup>	0.9 kW ( 0.9 kVA )	1 kW ( 1.1 kVA )	0.7 kW ( 0.7 kVA )	0.7 kW ( 0.7 kVA )	0.3 kW ( 0.3 kVA )	0.5 kW ( 0.6 kVA )
Cooling at Typical Usage	4,299 BTU/hour	5,050 BTU/hour	3,559 BTU/hour	3,337 BTU/hour	1,512 BTU/hour	2,634 BTU/hour
	4,536 kJ/hour	5,328 kJ/hour	3,755 kJ/hour	3,521 kJ/hour	1,595 kJ/hour	2,778 kJ/hour
Airflow at Maximum Usage <sup>2</sup>	3,010 BTU/hour	3,535 BTU/hour	2,491 BTU/hour	2,336 BTU/hour	1,058 BTU/hour	1,843 BTU/hour
	3,175 kJ/hour	3,729 kJ/hour	2,628 kJ/hour	2,464 kJ/hour	1,116 kJ/hour	1,945 kJ/hour
Airflow at Typical Usage <sup>2</sup>	199 CFM	234 CFM	165 CFM	154 CFM	70 CFM	122 CFM
	139 CFM	164 CFM	115 CFM	108 CFM	49 CFM	85 CFM
Operating temperature/humidity: 5 °C to 32 °C (41 °F to 89.6 °F), 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						
<sup>1</sup> Typical power usage varies by application load <sup>2</sup> Airflow must be front-to-back						



## Exadata Database Machine Environmental Specifications

Metric	Elastic Configuration 1 (Example) <sup>3</sup>	Elastic Configuration 2 (Example) <sup>3</sup>	Elastic Configuration (Half Rack) <sup>3</sup>	Quarter Rack	Eighth Rack
<b>Height</b>	78.74 in ( 2000 mm )				
<b>Width</b>	23.62 in ( 600 mm )				
<b>Depth</b>	47.12 in ( 1197 mm )				
<b>Acoustic Noise (operating)</b>	9.6 B	9.7 B	9.5 B	9.3 B	9.3 B
<b>Environmental With High Capacity Disks</b>					
<b>Weight</b>	1,904.7 lb. ( 864.0 kg )	2,058.7 lb. ( 933.8 kg )	1,411.1 lb. ( 640.1 kg )	937.8 lb. ( 425.4 kg )	904.8 lb. ( 410.4 kg )
<b>Maximum Power Usage</b>	21.6 kW ( 22 kVA )	20.1 kW ( 20.5 kVA )	13.2 kW ( 13.5 kVA )	6.5 kW ( 6.6 kVA )	5.2 kW ( 5.3 kVA )
<b>Typical Power Usage<sup>1</sup></b>	15.1 kW ( 15.4 kVA )	14 kW ( 14.3 kVA )	9.2 kW ( 9.4 kVA )	4.6 kW ( 4.6 kVA )	3.6 kW ( 3.7 kVA )
<b>Cooling at Maximum Usage</b>	73,631 BTU/hour	68,448 BTU/hour	45,016 BTU/hour	22,182 BTU/hour	17,687 BTU/hour
	77,680 kJ/hour	72,212 kJ/hour	47,492 kJ/hour	23,402 kJ/hour	18,659 kJ/hour
<b>Cooling at Typical Usage</b>	51,541 BTU/hour	47,913 BTU/hour	31,511 BTU/hour	15,528 BTU/hour	12,381 BTU/hour
	54,376 kJ/hour	50,549 kJ/hour	33,245 kJ/hour	16,382 kJ/hour	13,062 kJ/hour
<b>Airflow at Maximum Usage<sup>2</sup></b>	3409 CFM	3169 CFM	2084 CFM	1027 CFM	819 CFM
<b>Airflow at Typical Usage<sup>2</sup></b>	2386 CFM	2218 CFM	1459 CFM	719 CFM	573 CFM
<b>Environmentals With Extreme Flash Disks</b>					
<b>Weight</b>	1,778.7 lb. ( 806.8 kg )	1,907.0 lb. ( 865.0 kg )	1,313.1 lb. ( 595.6 kg )	895.8 lb. ( 406.3 kg )	N/A
<b>Maximum Power Usage</b>	21 kW ( 21.4 kVA )	20 kW ( 20.4 kVA )	12.7 kW ( 13 kVA )	6.3 kW ( 6.4 kVA )	N/A
<b>Typical Power Usage<sup>1</sup></b>	14.7 kW ( 15 kVA )	14 kW ( 14.3 kVA )	8.9 kW ( 9.1 kVA )	4.4 kW ( 4.5 kVA )	N/A
<b>Cooling at Maximum Usage</b>	71,635 BTU/hour	68,236 BTU/hour	43,464 BTU/hour	21,517 BTU/hour	N/A
	75,574 kJ/hour	71,989 kJ/hour	45,854 kJ/hour	22,700 kJ/hour	N/A
<b>Cooling at Typical Usage</b>	50,144 BTU/hour	47,765 BTU/hour	30,425 BTU/hour	15,062 BTU/hour	N/A
	52,902 kJ/hour	50,392 kJ/hour	32,098 kJ/hour	15,890 kJ/hour	N/A

<b>Airflow at Maximum Usage<sup>2</sup></b>	3316 CFM	3159 CFM	2012 CFM	996 CFM	N/A
<b>Airflow at Typical Usage<sup>2</sup></b>	2321 CFM	2211 CFM	1409 CFM	697 CFM	N/A
<b>Operating temperature/humidity: 5 °C to 32 °C (41 °F to 89.6 °F), 10% to 90% relative humidity, non-condensing</b> <b>Altitude Operating: Up to 3,048 m, max. ambient temperature is de-rated by 1° C per 300 m above 900 m</b>					
<sup>1</sup> Typical power usage varies by application load <sup>2</sup> Airflow must be front-to-back. <sup>3</sup> Elastic Configuration (Half Rack), Elastic Configuration 1 and Elastic Configuration 2 configurations added as examples of elastic configurations. Half = 4x DB Servers, 7x Storage Servers; Elastic Configuration 1 = 9x DB and 9x Storage Servers; Elastic Configuration 2 = 2x DB and 16x HC or 17x EF Storage Servers					

## Exadata Database Machine Regulations and Certifications

<b>Regulations</b> <sup>1,2,3</sup>	<b>Product Safety:</b>	UL/CSA 60950-1, EN 60950-1, IEC 60950-1 CB Scheme with all country differences
		UL/CSA 62368-1, EN 62368-1, IEC 62368-1 CB Scheme with all country differences
	<b>EMC</b>	
	<b>Emissions:</b>	FCC CFR 47 Part 15, ICES-003, EN55032, KS C 9832, EN61000-3-11, EN61000-3-12
	<b>Immunity:</b>	EN55024, KS C 9835
<b>Certifications</b> <sup>2,3</sup>	North America (NRTL), CE (European Union), International CB Scheme, HSE Exemption (India), BSMI (Taiwan), KC (Korea), RCM (Australia), VCCI (Japan), UKCA (United Kingdom)	
<b>European Union Directives</b> <sup>3</sup>	2014/35/EU Low Voltage Directive, 2014/30/EU EMC Directive, 2011/65/EU RoHS Directive, 2012/19/EU WEEE Directive	
<sup>1</sup> All standards and certifications referenced are to the latest official version. For additional detail, please contact your sales representative. <sup>2</sup> Other country regulations/certifications may apply. <sup>3</sup> 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 X10M

Each Exadata Database Machine X10M 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 and 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 Reverse Offload to DB servers if Storage CPUs are Busy
- Unique Automatic Data Columnarization
- Unique Automatic Conversion of Data to In-Memory Formats when Loading into Flash Cache

### Exadata and Database Software Features – OLTP

- Unique Exadata RDMA Memory Data Accelerator
- Unique Exadata RDMA Memory Commit Accelerator (X8M and X9M only)
- 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 cluster, workload, DB or user 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 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 4000 Pluggable Databases per Container Database 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
- Unique Storage Index persistence to avoid rebuild on storage server restart
- Unique Database In-Memory Columnar Cache persistence to avoid rebuild on storage server restart

### Manageability Features

- Oracle Embedded Integrated Lights Out Manager (ILOM) with upgrade pre-staging optimizations
- Oracle Enterprise Manager Exadata Plug-in
- Unique Active AWR includes storage stats for end-to-end monitoring

- Real-time Insights server metric streaming
- 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 Linux Virtualization
- Automated VLAN Creation
- Oracle Exadata Deployment Assistant
- Separate Management Switch and Connectivity
- Exaccli command line management from remote servers
- Cellcli command line management of Storage Servers
- DCLI distributed command line automation tool
- Automatic Service Request and Patch Manager (patchmgr) support for:
  - database servers,
  - storage servers
  - power distribution units, and
  - Cisco RoCE and management switches

**Oracle Database Software (available separately):**

- **For database servers:** Oracle Database 19c Enterprise Edition, Oracle Database 21c Enterprise Edition, and Oracle Database 23ai 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. Oracle Grid Infrastructure 19c or higher is required.
- **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 8 with the Unbreakable Enterprise Kernel 6. Zero-loss Zero-copy Datagram Protocol (ZDP) RoCEv2 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)

## Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

 [blogs.oracle.com](https://blogs.oracle.com)

 [facebook.com/oracle](https://facebook.com/oracle)

 [twitter.com/oracle](https://twitter.com/oracle)

Copyright © 2024, Oracle and/or its affiliates. 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, Java, and MySQL are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.