

Oracle Exadata Database Machine X11M


The Oracle Exadata Database Machine (Exadata) is engineered to deliver exceptional performance, availability, cost-effectiveness, and efficiency for Oracle databases. Exadata uses a modern, cloud-enabled scale-out architecture with high-performance database servers and intelligent storage servers. State-of-the-art flash, RDMA-accessible memory for unique storage caching, and a cloud-scale RDMA over Converged Ethernet (RoCE) fabric connect all servers and storage.

Exadata's unique algorithms and protocols enable database intelligence across storage, compute, and networking layers. It delivers greater performance, capacity, and cost efficiency than other platforms.

Exadata is the ideal platform for a broad range of modern database workloads, including Artificial Intelligence (AI), Analytics, Data Warehousing (DW), In-Memory Analytics, Online Transaction Processing (OLTP), and Internet of Things (IoT).

Exadata X11M is the latest generation of Exadata and implements Exadata Exascale, the next-generation, data-intelligent software architecture that enables revolutionary innovations, including database-aware intelligent clones and snapshots.

Exadata X11M is simple to implement and powers and protects your most important databases. Exadata X11M is designed to run everywhere. It can be deployed on-premises as the ideal foundation for a private database cloud or in the Oracle Cloud Infrastructure, Cloud@Customer in a hybrid cloud model, or Multicloud (Azure, Google Cloud, and AWS) using a subscription model with all infrastructure management performed by Oracle. The Oracle Autonomous Database is available exclusively on Exadata, either in the Oracle Cloud Infrastructure, Multicloud, or Cloud@Customer.



**Oracle Exadata Database
Machine X11M**

Key Features

- Up to 2,880 CPU cores per rack for database processing
- Up to 42 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
- 2X100 Gb/sec active-active 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 4.4 PB of disk capacity (raw) per rack

Engineered For Fast and Reliable Deployment

Exadata 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 there are no performance bottlenecks or single points of failure.

All Exadata systems are identically configured; therefore, customers benefit from the experience of thousands of other customers' Exadata deployments. Customer systems are also identical to the systems that Oracle Support uses for problem identification and resolution, the systems on which Oracle Development develops and tests the Oracle Database, and the systems that run Oracle's SaaS applications, such as Oracle Fusion Applications and NetSuite. **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 Exadata 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

Exadata 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 deployment to scale without bottlenecks. The **architecture scales from small to extremely large configurations to accommodate workloads of any size**. In Exadata X11M, a high-bandwidth, low-latency active-active 100 Gb/sec RDMA over Converged Ethernet (RoCE) Network Fabric connects all the components. Specialized database networking protocols deliver much lower latency and higher bandwidth than is possible with generic communication protocols, enabling **faster response time for OLTP operations and higher throughput for analytic workloads**. External data center connectivity to Exadata is via standard 10 Gb/sec, 25 Gb/sec, or 100 Gb/sec Ethernet.

Exadata is the most versatile database platform, with extreme scalability inherent at every layer in its architecture. Exadata X11M extends this innovation further with powerful database servers equipped with two of the latest 96-core x86 processors and 512 GB of DDR5 memory (expandable up to 3 TB).

Exadata X11M also introduces the new Database Server-Z, equipped with a single 32-core x86 processor and 768 GB of DDR5 memory (expandable up to 1,152 GB). This provides customers with the choice of database servers to match their workload requirements, such as those requiring fewer CPU cores or consolidating a few virtual machines or databases.

Key Benefits

- Pre-configured, pre-tested system optimized for all database applications
- Uncompressed I/O bandwidth of up to 8.5 TB/sec per rack from SQL
- Ability to perform up to 25.2M 8K database read I/O operations, or 13M 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 X11M racks or Exadata X11M 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

The CPU scalability within each database server enables superior performance and utilization efficiency by enabling greater database and VM consolidation, 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 four configurations in Exadata X11M: High Capacity (HC), Extreme Flash (EF), High Capacity-Z (HC-Z), and Extended (XT)

- Exadata X11M HC Storage Servers have four 6.8 TB performance-optimized Flash Accelerator F680 v2 NVMe PCIe 5.0 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.
- Exadata X11M EF Storage Servers have an all-flash configuration. Each server has four 6.8 TB performance-optimized Flash Accelerator F680 v2 NVMe PCIe 5.0 flash cards for Exadata Smart Flash Cache and four 30.72 TB capacity-optimized NVMe PCIe Flash drives totaling 122.88 TB of raw storage.
- Exadata X11M HC and EF Storage Servers include Exadata RDMA Memory (XRMEM), further enhancing performance by delivering up to 2.8 million 8K OLTP Read IOPS. Each server is populated with 1.5 TB of DDR5 memory, 1.25 TB of which is configured as Exadata RDMA Memory and used as a caching tier between the database buffer cache and Flash Cache, and the remaining 256 GB is used for Exadata System Software. SQL Scan throughput from flash is 2.2X faster than the prior generation, reaching 100 GB/s per storage server. In addition, Smart Scan also reads columnar data in XRMEM accelerating analytics query throughput up to 500 GB/s. Two 32-core x86 processors are included for Exadata System Software operations in each HC and EF storage server.
- Exadata X11M HC-Z Storage Servers are ideal for running small to medium workloads with two 6.8 TB performance-optimized Flash Accelerator F680 v2 NVMe PCIe 5.0 flash cards for Exadata Smart Flash Cache and six 22 TB 7,200 RPM SAS disks. Each server is populated with 768 GB of DDR5 memory, 576 GB of which is configured as Exadata RDMA Memory and used as a caching tier between the database buffer cache and Flash Cache, and the remaining 192 GB is used for Exadata System Software. One 32-core x86 processor is included for Exadata System Software operations in each HC-Z storage server.
- Exadata X11M XT Storage Servers are ideal for storing less frequently used data, and include Hybrid Columnar Compression. Each XT Storage Server provides 264 TB of raw storage using twelve 7,200 RPM SAS disks, each with a 22 TB capacity. Each server includes two 6.8 TB performance-optimized Flash Accelerator F680 v2 NVMe PCIe 5.0 flash cards for Exadata Smart Flash Cache and can be deployed with Exascale as a separate storage pool or with ASM as a separate disk group. A 32-core x86 processor and 256 GB of DDR5 memory run Exadata System Software. An optional license for Oracle Exadata System Software enables advanced features such as Exadata Smart Scan.

Related Products

- Exadata Cloud Infrastructure
- Exadata Cloud@Customer
- Oracle Exadata Storage Expansion Rack X11M
- Oracle Exadata Storage Server X11M
- Oracle Exadata Database Server X11M
- Oracle Database Enterprise Edition 19c, 21c, and 23ai
- Exadata Exascale
- 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

The minimal configuration of an Exadata consists of two database servers and three HC, EF, or HC-Z 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. Customers can mix both X11M database server and any of the storage server configurations within the same rack, enabling even greater control of resources to meet varying performance, high availability, security, and consolidation requirements. Capacity for less-frequently used data can be added using XT storage servers.

In addition to expanding within a rack, **multiple Exadata racks can be interconnected using the integrated RoCE network fabric** to form even larger configurations. These racks can be Exadata X8M, X9M, X10M, or X11M generations. A system composed of four racks of Exadata X11M 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 Exadata Exascale and Automatic Storage Management (ASM) can dynamically add more storage capacity.




Figure 1 Elastic Scale-out to Multi-rack Exadata

For databases and workloads that require extreme storage capacity, the **Oracle Exadata X11M Storage Expansion Rack** is used to expand the storage tier of Exadata. The Storage Expansion Rack expands the storage capacity, Flash Cache capacity, OLTP IOPS, and scan throughput of any Exadata. It is designed for database deployments with very large amounts of data, including historical or archive data, backups, vectors, documents, images, XML, JSON, and LOBs. The Storage Expansion Rack connects to Exadata 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 X11M uses an ultra-fast cloud-scale RDMA over Converged Ethernet (RoCE) networking fabric. RDMA (Remote Direct Memory Access) allows one computer to directly access data from another without Operating System or CPU involvement for high bandwidth and very low latency. The network card directly reads/writes memory with no extra copying or buffering. RDMA is an integral part of the Exadata high-performance architecture and is tuned and enhanced with each new generation of Exadata, underpinning several Exadata-only technologies, including

Exafusion Direct-to-Wire Protocol for faster inter-instance OLTP communication, or to **read undo blocks** from other Real Application Clusters (RAC) instances to increase OLTP performance.

Exadata X11M implements a dual-port PCIe 5.0 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 17X**, from 250 microseconds¹ to as low as **14 microseconds**. 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 the storage servers, 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 X11M, utilizing the shared XRMEM Data Accelerator, can achieve up to **25.2 million OLTP Read IOPS (8K IOs)²** in a single rack. This represents a 21% improvement over the same nine database and nine storage server configurations of the Exadata X9M generation at 20.7 million³ and 430% higher than Exadata X7⁴. This performance scales as additional racks are deployed.

XRMEM's security and management are fully automated. It is configured automatically, with no user interaction required, and is automatically managed thereafter. Hardware monitoring is pre-configured. XRMEM is only accessible to databases using database access controls, ensuring end-to-end data security. XRMEM is entirely transparent to all applications.

Record-breaking Analytics I/O Performance

Each X11M High Capacity (HC) and Extreme Flash (EF) Storage Server includes four 6.8 TB performance-optimized flash drives, offering 27.2 TB of Exadata Smart Flash Cache. 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 HC and EF Storage Servers include 1.25 TB of DDR5 Exadata RDMA Memory as a data acceleration tier in front of Flash Cache.

For analytics environments that require the highest performance, Exadata X11M HC and EF storage servers are both capable of scanning data in flash at **up to 100 GB/s with Smart Scan, 2.2X faster** than the previous generation. In addition, Smart Scan can achieve **up to 8.5 TB/s scan throughput⁵** from a single rack configuration by reading columnar data cached in Exadata RDMA Memory.

This represents real-world, end-to-end performance 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.

¹ Exadata X7 8K OLTP Read Latency from flash in storage servers

² Elastic configuration with 9X Exadata X11M Database Servers and 9X Exadata X11M Extreme Flash or 9X Exadata X11M High Capacity Storage Servers

³ Elastic configuration with 9X Exadata X9M Database Servers and 9X Exadata X9M Extreme Flash Storage Servers or 9X Exadata X9M High Capacity Storage Servers

⁴ Exadata X7-2 Full Rack with 8X Exadata X7 Database Servers and 14X Exadata X7 Extreme Flash Storage Servers or 14X High Capacity Storage Servers

⁵ Elastic configuration with 2X Exadata X11M Database Servers and 17X Exadata X11M High Capacity or 17X Exadata X11M Extreme Flash Storage Servers

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

The most common 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 with a total raw capacity of 27.2 TB of Exadata Smart Flash Cache and 1.25 TB of DDR5 Exadata RDMA Memory (XRMEM) in front of the 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 affects 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 X11M rack reaches **13 million 8K flash write I/O operations per second (IOPS)**⁶. 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, due to the need for such arrays to move data to the database servers instead of processing data in the storage, the network infrastructure becomes a significant performance bottleneck as a single flash card can overwhelm the network. Therefore, all-flash storage arrays cannot match the cost advantages of Exadata's smart data tiering between disk and flash because they lack Exadata's unique database-aware storage optimizations. 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 that 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

⁶ Elastic configuration with 6X Exadata X11M Database Servers and 13X Exadata X11M Extreme Flash Storage Servers or 13X Exadata X11M High Capacity Storage Servers

Extreme Flash Storage Server: Lowest Latency, All Flash

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 and increasing the usable storage capacity by 2.4X⁷. Exadata EF storage servers deliver extreme high performance, flash capacity, and low latency for database workloads with the highest performance demands. In addition, X11M EF storage includes four 6.8 TB performance-optimized flash drives for a total raw capacity of 27.2 TB, increasing the **size of the Smart Flash Cache by 11.5X**⁸, and is used to satisfy read and write requests. 1.25 TB of DDR5 Exadata RDMA Memory (XRMEM) is available in front of flash to boost performance further.

High Capacity-Z Storage Server: High Performance Storage Accelerates Small to Medium Database Workloads

The third Exadata storage option is the **High Capacity-Z (HC-Z) Storage Server**. Each HC-Z Storage server includes six 22 TB disk drives with 132 TB total raw disk capacity. To boost performance further, the HC-Z Storage Server also includes two 6.8 TB NVMe PCIe Flash Cards for 13.6 TB of Exadata Smart Flash Cache and 576 GB of DDR5 XRMEM in front of the flash.

The High Capacity-Z storage server is ideal for smaller workloads or environments. It delivers Exadata's extraordinary performance, scalability, security, and manageability while making Exadata accessible to all organizations.

Exadata Extended (XT) Storage Server: Cost-Effective Storage for Less Frequently Accessed Data

The Exadata Extended (XT) Storage Server delivers cost-effective, scalable storage for less frequently accessed data that must remain online for retention or compliance. Each XT server features twelve 22 TB hard drives, providing up to 264 TB of raw capacity, along with two 6.8 TB flash cards to accelerate data caching. Deploy XT Storage to expand Exascale storage or to create or extend ASM XTND disk groups for large volumes of infrequently accessed data.

XT storage works seamlessly with Exascale and ASM and integrates built-in Oracle Database features, including security, access control, and Hybrid Columnar Compression. XT storage supports Oracle Database 23ai and Exascale, enabling instant database-aware snapshots and clones so you can manage test and development environments more easily.

Like other Exadata storage options, Extreme Flash (EF), High Capacity (HC), and High Capacity-Z (HC-Z), XT Storage uses a scale-out architecture for easy capacity expansion and unified management, extending Exadata's proven security, manageability, and reliability to less frequently used data, without the hassle of managing multiple platforms or adding operational risk.

XT storage servers keep your compliance and long-term retention data online and accessible. Optional licensing of Exadata System Software enables smart features, including Smart Scan and Storage Indexes, and delivers **SQL scan throughput up to 50 GB/s**.

By combining all Exadata storage options, Extreme Flash, High Capacity, High Capacity-Z, and Extended (XT) Storage, organizations can implement an efficient information lifecycle management strategy. As data ages, it can be automatically moved between storage tiers and compression levels using Oracle's Automatic Data Optimization, ensuring that storage resources align with data usage and retention requirements for maximum cost efficiency.

⁷ Compared to X9M Extreme Flash Storage Servers

⁸ 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.

Accelerate 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 offload database processing. These CPUs 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** eliminates the storage disks as a potential log write throughput bottleneck, prevents log write latency outliers, and automatically and transparently stores Oracle Database Redo Logs in the Smart Flash Cache. Smart Flash Log Write-Back **increases log write throughput by up to 2.5X**. In addition, **Pipelined Log Writes increase redo write throughput by up to 1.4X** on Exadata⁹ by automatically using the high-performance RoCE network to execute parallel redo log writes without impacting latency, minimizing the need to queue redo log writes, and improving OLTP performance. 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

⁹ Available on Exadata X10M and later generations

(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 Exadata, 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

AI Smart Scan Accelerates Artificial Intelligence Workloads

Oracle Database 23ai introduces significant new capabilities, known as AI Vector Search, to enable and accelerate Artificial Intelligence (AI) workloads, enabling applications to query data based on semantic similarity. With AI Vector Search, applications can search unstructured data, such as contracts, knowledge management documents, images, video, audio, etc., and then combine that unstructured data with data stored in relational or other data types already stored in Oracle Database. This continuation of Oracle's Converged Database strategy helps eliminate data fragmentation by avoiding the need for a specialized vector database and elevates vector processing into the enterprise-grade data management domain.

AI Smart Scan accelerates AI Vector Search queries by up to 30X with optimizations that deliver low latency parallelized scans across large volumes of vector data and enables large concurrent AI user communities. AI Smart Scan processes vector data at memory speed from ultra-fast Exadata RDMA Memory and Flash Cache in storage servers and performs vector distance computations and **Adaptive Top-K Filtering** where the data resides, avoiding unnecessary network data transfer and database work resulting in **up to 4.6X faster queries** and **up to 4.7X greater data filtering** respectively. Additionally, AI queries are **up to 32X faster when using BINARY vector dimension formats** in Oracle Database 23ai.

Exadata Exascale Intelligent Data Architecture for the Cloud Era

Exadata Exascale is the world's only intelligent data architecture for the cloud. Exascale is a revolutionary leap that combines the best of Exadata and the best of the cloud. This next-generation, hyper-elastic, multi-tenant software architecture is designed to handle mission-critical OLTP, analytics, AI, JSON, and mixed Oracle Database workloads.

Exascale reimagines how compute and storage resources are managed on Exadata platforms by decoupling and simplifying storage management, paving the way for innovative new capabilities. It ensures industry-leading database performance, availability, and security standards that organizations expect from Exadata.

Exascale features a **reimagined approach to database snapshots and clones on Exadata.** It enables **space-efficient thin clones from any read/write database or pluggable database**, significantly boosting developer

productivity. Exascale seamlessly integrates with development, test, and deployment pipeline requirements while providing native Exadata performance.

Databases on Exascale are automatically distributed across all available storage in the Exadata storage servers, providing **ultra-low latency RDMA for I/O and database-aware intelligent Smart Scan with up to thousands of cores available to all workloads**. Automatic data replication across multiple storage servers ensures fault tolerance and reliability.

Energy Efficiency Meets Extreme Performance

Customers are increasingly required to meet specific power consumption reduction targets to align with their corporate sustainability policies and responsibilities. Exadata X11M is designed to assist customers achieve environmental and sustainable goals while enabling extreme performance. Exadata X11M reduces energy consumption, space requirements, and cooling needs, helping customers reduce environmental impact while running high-performance workloads.

The first step in reducing power consumption is choosing the right database server for your specific workload. Exadata X11M offers two configurations for varying power and performance needs: a two-socket, 96-core CPU Database Server and a one-socket, 32-core CPU Database Server-Z. This flexibility allows customers to align power consumption with their workload and sustainability targets.

Exadata X11M also introduces innovative features to **optimize energy efficiency**, such as **setting a specific CPU power target or dynamically adjusting power consumption based on workload demands**—ideal for quieter periods like nights or weekends. With Exadata Capacity-on-Demand, customers can activate only the cores they need at deployment, reducing energy use and only paying for the cores required by their workloads.

Designed to run AI, OLTP, analytics, and mixed workloads on a single platform, Exadata X11M provides extreme scalability and performance, making it the ideal choice for workload consolidation. By consolidating multiple databases on Exadata, customers can increase utilization efficiency, reduce their data center footprint, and lower power consumption. Exadata is designed to be the ideal platform for Oracle Database. The X11M generation continues to execute this mission while introducing capabilities to help customers affect tangible power savings and increased data center efficiency, driving a more sustainable future.

Optimize Storage Consumption and I/O Using 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 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.

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 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.

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. Exadata X11M provides up to 462 TB of Flash Cache in a single rack¹⁰, designed to handle some of the largest in-memory workloads. It leverages a unique algorithm to increase the compression of data stored in the In-Memory columnar format by up to 1.25X¹¹, further expanding the effective capacity of the Flash Cache. Databases not using 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 time 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.

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, reaping the processing benefits of the shared storage tier, and freeing up database server memory.

¹⁰ Elastic rack with 2X Exadata X11M Database Servers and 17X Exadata X11M Extreme Flash Storage Servers or 17X Exadata X11M High Capacity Storage Servers

¹¹ Compared to X9M

11 Oracle Exadata Database Machine X11M / Version 1.2

Increase Utilization Efficiency on Exadata Through Consolidation

Exadata 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 AI, OLTP, analytics, and batch operations with sequential and random-access patterns. Exadata's ability to **run more database workloads with industry-leading scalability and performance** makes it an ideal consolidation platform.

Consolidated environments running on Exadata X11M 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 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 can utilize RDMA-enabled Exascale Volumes to increase performance and consolidation density with up to 50 VMs per database server.

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

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. Exadata **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 X11M, 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 4096 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 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 maintaining security across the entire stack. **Oracle delivers full-stack security in Exadata**.

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. Secure Boot can be used for bare metal and virtual machines. 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 X11M enable Instant Secure Erase to eliminate data leakage risk during proactive drive replacement or machine redeployment. Exadata's Secure Erase feature leverages this on-device capability when an Exadata is re-purposed or decommissioned to instantly erase all user data present on storage devices by cryptographically erasing the device, rather than require the device to be overwritten. Because the previous device

cryptographic 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 is 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

Exadata 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 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 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 segment¹².

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 extended application brownouts. RoCE-based Exadata 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 a 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, a rebalance operation is performed 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 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, automatically identify the process responsible, and issue an alert—all without any manual intervention.

¹² Worldwide AL4 Server Market Shares, 2019: Fault-Tolerant Systems Become Digital Transformation Platforms, IDC, Paul Maguranis Peter Rutten, July 2020

Exadata's industry-leading availability makes it the platform of choice for the world's top companies running their most critical applications—ranging from interbank fund transfers and online securities trading to real-time call tracking and web-based retailing. Its exceptional availability isn't restricted to OLTP workloads; it extends to AI and analytics as well, ensuring reliable performance across all mission-critical environments.

Ultra-Fast Deployment of Development and Test Databases with Exadata Exascale

Exadata Exascale features a **reimagined approach to database snapshots and clones on Exadata**. It enables **space-efficient thin clones from any read/write database or pluggable database**, significantly boosting developer productivity. Exascale seamlessly integrates with development, test, and deployment pipeline requirements while providing native Exadata performance.

Administrators and users can easily create read-only database snapshots and read-write thin-provisioned clones from any Z3ai source database, including Data Guard standby databases, for use in testing, development, and other purposes. Exascale introduces a new redirect-on-write architecture that allows snapshots and clones to share the same underlying storage, significantly reducing the space required for changes made in the clones. This space-efficient design enables development teams to provide each developer with their own production-like environment without the need for costly storage investments.

Database clones on Exascale can be created from any read/write or read-only database or Pluggable Database (PDB), including from Data Guard standby and existing clones, without requiring any alteration to the source database. Each clone is entirely independent of its source database. Source databases can continue to be used for their intended purpose or even dropped without impacting any clones created, increasing flexibility and operational efficiency.

Exascale is integrated with Oracle Multitenant to provide a simple interface for creating new pluggable database (PDB) snapshots and clones. Exascale enables flexible cloning workflows by allowing PDBs to be cloned within the same Container Database (CDB), or between CDBs in the same Exadata Exascale infrastructure. The **PDB Snapshot Carousel** capability of Oracle Database automatically creates PDB snapshots at regular intervals for subsequent use as a point-in-time PDB clone source. Data Guard standby databases can be utilized to create snapshots and thin-provisioned read-write copies of the source databases. Oracle REST Data Services (ORDS) enables REST API control of PDB snapshots, clones, and the PDB Snapshot Carousel for easy integration with CI/CD development pipelines and workflows.

Space-efficient clones of a CDB and all the PDBs it contains can be created using the gDBCLone utility. gDBCLone simplifies the process of cloning CDBs with an easy-to-use command-line interface.

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. This provides a production-like 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 and provides comprehensive capabilities from monitoring and reporting to active lifecycle management. It enables:

- **Unified Monitoring:** Oracle Enterprise Manager 24ai 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 Insight 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, 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 Exadata's extreme performance, high storage capacity, and unique compression capabilities, 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 system¹⁵ can reach 4.4 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 of 10X. By intelligently moving active data across disk, flash, and memory tiers, Exadata simultaneously delivers the highest performance and the lowest cost.

¹⁵ Exadata X11M elastic configuration with 2X X11M Database Servers and 17X X11M High Capacity Storage Servers or 17X X11M Extreme Flash Storage Servers

Exadata can **uniquely consolidate many databases supporting multiple workloads** in a single cloud platform. High-end OLTP, AI, 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 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 are reduced or eliminated. All Exadata users run a configuration that is identical to thousands of other users and is 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 X11M Database Server has 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 the Cloud (Exadata Cloud), with the same extreme performance and availability experienced by thousands of organizations running Exadata on-premises. Exadata 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. Exadata Cloud is available in the Oracle public cloud (Oracle Cloud Infrastructure - OCI), within partner public clouds (Azure, Google Cloud, and AWS), and as a hybrid cloud platform (Exadata Cloud@Customer) within customer data centers.

With Exadata Cloud:

- Compute resources can be online scaled up and down, enabling customers to pay only for what they use, starting with a highly affordable entry-level subscription on OCI
- Customers choose between fully managed Autonomous Database or, for total control, Exadata Database Service, or run both services on the same Exadata infrastructure
- All Exadata System Software and hardware is included with the infrastructure subscription
- Customers can bring their own on-premises Oracle Database licenses or subscribe to an all-inclusive license containing all Oracle Database options and features
- Oracle Cloud operations manage all infrastructure, eliminating many administration activities previously performed by customer staff
- Powerful cloud automation exposed through a browser-based UI and REST APIs simplifies common lifecycle management tasks
- Customers who are unable to migrate their databases to the public cloud because of security, compliance, data residency or dependencies on other on-premises systems can run Exadata Cloud in their own datacenter with Exadata Cloud@Customer.

Oracle databases deployed on Exadata Cloud are **100% compatible** with those deployed on-premises, ensuring a smooth transition to the cloud, and a seamless hybrid cloud strategy. For existing Oracle Database customers, applications and data models **do not need to change** to include the elasticity and flexibility of Exadata Cloud. They also do not have to invest in separate database cloud services for different data models or workloads since Exadata provides a unified platform for all applications – AI, OLTP, analytics, consolidation, and mixed workloads.

Exadata Cloud is an ideal fit for:

- Business-critical production databases at any scale without incurring the capital expenditure and complexity of maintaining the underlying IT infrastructure
- Minimizing costs for workloads whose resource requirements vary over time
- Consolidating a variety of systems, databases, and database services on powerful Exadata infrastructure, avoiding the cost and complexity of multiple database services
- 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 ad-hoc Oracle databases for development, functionality testing, application certification, and proof-of-concept

Uniquely engineered for extreme performance for all workloads, along with fast deployment, simplified management, low operating costs, and reduced risks, Exadata Cloud is simply 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 results 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

Server Type	CPU	Memory (DDR5)	Disk	Flash	Standard Network Adapters	Optional Network Adapters
Database Server²	2 x 96-core AMD EPYC 9J25 processors, 2.6 GHz (up to 4.5 GHz)	Choose: 512 GB ³ 1,536 GB ⁴ 2,304 GB ⁴ 3,072 GB ⁵		2 x 3.84 TB NVMe Flash SSD (hot swappable), (upgradeable to 4 x 3.84 TB)	<ul style="list-style-type: none"> Client/backup adapter 1: 2 x 10/25 Gb Ethernet ports (SFP28) Client/backup adapter 2: 2 x 10/25 Gb Ethernet ports (SFP28) 1 x 1 Gb Ethernet port (RJ45, management) 1 x 1 Gb Ethernet port (RJ45, ILOM) 2 x 100 Gb QSFP28 RoCE Fabric ports 	Client/backup adapter 3, 4, or 5: <ul style="list-style-type: none"> 4 x 10 Gb Ethernet ports (RJ45), or 2 x 10/25 Gb Ethernet ports (SFP28), or 2 x 100 Gb optical Ethernet ports (QSFP28)
Database Server-Z²	1 x 32-core AMD EPYC 9J15 processor 2.95 GHz (up to 4.4 GHz)	Choose: 768 GB ³ 1,152 GB ⁵		2 x 3.84 TB NVMe Flash SSD (hot swappable), (upgradeable to 4 x 3.84 TB)	<ul style="list-style-type: none"> Client/backup adapter 1: 2 x 10/25 Gb Ethernet ports (SFP28) Client/backup adapter 2: 2 x 10/25 Gb Ethernet ports (SFP28) 1 x 1 Gb Ethernet port (RJ45, management) 1 x 1 Gb Ethernet port (RJ45, ILOM) 2 x 100 Gb QSFP28 RoCE Fabric ports 	Client/backup adapter 3: <ul style="list-style-type: none"> 4 x 10 Gb Ethernet ports (RJ45), or 2 x 10/25 Gb Ethernet ports (SFP28), or 2 x 100 Gb optical Ethernet ports (QSFP28)
High Capacity Storage Server	2 x 32-core AMD EPYC 9J15 processors 2.95 GHz (up to 4.4 GHz)	256 GB and 1,280 GB Exadata RDMA Memory	12 x 22 TB 7,200 RPM disks	4 x 6.8 TB NVMe PCIe 5.0 performance-optimized Flash cards	<ul style="list-style-type: none"> 2 x 100 Gb QSFP28 RoCE Fabric ports 1 x 1 Gb Ethernet port (RJ45, management) 1 x 1 Gb Ethernet port (RJ45, ILOM) 	
Extreme Flash Storage Server	2 x 32-core AMD EPYC 9J15 processors 2.95 GHz (up to 4.4 GHz)	256 GB and 1,280 GB Exadata RDMA Memory		4 x 6.8 TB NVMe PCIe 5.0 performance-optimized Flash cards, and 4 x 30.72 TB NVMe PCIe 4.0 capacity-optimized Flash cards	<ul style="list-style-type: none"> 2 x 100 Gb QSFP28 RoCE Fabric ports 1 x 1 Gb Ethernet port (RJ45, management) 1 x 1 Gb Ethernet port (RJ45, ILOM) 	
High Capacity-Z² Storage Server	1 x 32-core AMD EPYC 9J15 processor 2.95 GHz (up to 4.4 GHz)	192 GB and 576 GB Exadata RDMA Memory	6 x 22 TB 7,200 RPM disks	2 x 6.8 TB NVMe PCIe 5.0 performance-optimized Flash cards	<ul style="list-style-type: none"> 2 x 100 Gb QSFP28 RoCE Fabric ports 1 x 1 Gb Ethernet port (RJ45, management) 1 x 1 Gb Ethernet port (RJ45, ILOM) 	
Extended Storage Server	1 x 32-core AMD EPYC 9J15 processor 2.95 GHz (up to 4.4 GHz)	256 GB	12 x 22 TB 7,200 RPM disks	2 x 6.8 TB NVMe PCIe 5.0 performance-optimized Flash cards	<ul style="list-style-type: none"> 2 x 100 Gb QSFP28 RoCE Fabric ports 1 x 1 Gb Ethernet port (RJ45, management) 1 x 1 Gb Ethernet port (RJ45, ILOM) 	

¹ All servers include redundant hot swappable fans and power supplies

² Database Server-Z and High Capacity-Z Storage Servers can be combined with standard servers in the same rack. Minimum server requirements apply for Normal and High Redundancy

³ Factory only option

⁴ Factory option and field upgrade

⁵ Maximum configuration, factory option and field upgrade

Exadata Rack Configurations ^{1, 2, 5}

Rack Size	Database Servers and Cores	Storage Servers and Cores	High Capacity Storage Server Capacity (Raw)	Extreme Flash Storage Server Capacity (Raw)
Quarter Rack	2 x servers, 384 cores	3 x servers, 192 cores for SQL offload	<ul style="list-style-type: none"> 792 TB disk 81.6 TB performance-optimized Flash 3.75 TB Exadata RDMA Memory 	or <ul style="list-style-type: none"> 368.6 TB capacity-optimized Flash 81.6 TB performance-optimized Flash 3.75 TB Exadata RDMA Memory
Quarter Rack (Database Server-Z)	2 x servers, 64 cores	3 x HC-Z servers, 96 cores for SQL offload or 3 x HC or EF servers, 192 cores for SQL offload	Using Storage Server High Capacity-Z <ul style="list-style-type: none"> 396 TB disk 40.8 TB performance-optimized Flash 1.69 TB Exadata RDMA Memory or Using Storage Server High Capacity <ul style="list-style-type: none"> 792 TB disk 81.6 TB performance-optimized Flash 3.75 TB Exadata RDMA Memory 	or <ul style="list-style-type: none"> 368.6 TB capacity-optimized Flash 81.6 TB performance-optimized Flash 3.75 TB Exadata RDMA Memory
Elastic Configuration 1 (Example)³	9 x servers, 1,728 cores	9 x servers, 576 cores for SQL offload	<ul style="list-style-type: none"> 2,376 TB disk 244.8 TB performance-optimized Flash 11.25 TB Exadata RDMA Memory 	or <ul style="list-style-type: none"> 1,105.9 TB capacity-optimized Flash 244.8 TB performance-optimized Flash 11.25 TB Exadata RDMA Memory
Elastic Configuration 2(Example)³	2 x servers, 384 cores	17 servers, 1,088 cores for SQL offload	<ul style="list-style-type: none"> 4,488 TB disk 462.4 TB performance-optimized Flash 21.25 TB Exadata RDMA Memory 	or <ul style="list-style-type: none"> 2,088 TB capacity-optimized Flash 462.4 TB performance-optimized Flash 21.25 TB Exadata RDMA Memory

+Database Servers	Up to 15X servers ⁴ , 2,880 cores max per rack			
+Storage Servers		Maximum per rack: Up to 17x HC or EF servers ⁴ , 1,088 cores or Up to 17x HC-Z servers ⁴ , 544 cores or Up to 14x XT servers ⁶ , 448 cores	Maximum per rack using HC storage: <ul style="list-style-type: none"> • 4,488 TB disk • 462.4 TB performance- optimized Flash • 21.25 TB Exadata RDMA Memory or Maximum per rack using HC- Z storage <ul style="list-style-type: none"> • 2,244 TB disk • 231.2 TB performance- optimized Flash • 9.6 TB Exadata RDMA Memory or Maximum per rack using XT storage <ul style="list-style-type: none"> • 3,696 TB disk • 190.4 TB performance- optimized Flash 	or Maximum per rack using EF storage: <ul style="list-style-type: none"> • 2,088 TB capacity- optimized Flash • 462.4 TB performance- optimized Flash • 21.25 TB Exadata RDMA Memory

¹ 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.

² 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

³ Elastic Configuration 1 and Elastic Configuration 2 configurations added as examples of elastic configurations.

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

⁵ Exadata X11M enables significant flexibility by allowing 1 socket (Database Server-Z and High Capacity-Z) and 2 socket database and storage servers to be mixed in almost any combination. For example, two Database Server-Z and three Extreme Flash (EF) Storage servers can be configured for workloads that require the extreme low latency of flash while having low database CPU and/or memory requirements. The same type of storage server must be used when expanding existing storage servers. E.g. Extreme Flash must be used to expand Extreme Flash servers including disk groups and storage pools. When adding a new storage server type, a minimum of two servers is required for Normal redundancy and three for High redundancy (recommended).

⁶ Exadata X11M XT storage servers can be used to provide an ASM disk group, Exascale storage pool, or both. For ASM, the minimum number of storage servers is two. For Exascale, the minimum number of storage servers is three. When ASM and Exascale are deployed on the same XT servers, the minimum number of storage servers is three. XT servers cannot be used without at least three HC, HC-Z or EF storage servers in the configuration.

Other Elastic Expansion Options

Multi-Rack Connection	Connect any combination of up to 14 Exadata Database Machine racks or Exadata Storage Expansion Racks via the RoCE Network Fabric. <ul style="list-style-type: none"> • Larger configurations can be built with external RoCE switches. • Connected racks must contain Exadata RoCE hardware
------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Exadata Capacity and Performance Metrics: Individual Servers

Server Type	Maximum SQL Flash Bandwidth ²	Maximum SQL Exadata RDMA Memory Bandwidth ²	Maximum SQL Read IOPS ^{1,3}	Maximum SQL Write IOPS ⁴	Exadata RDMA Memory Capacity (Raw) ⁵	Performance-Optimized PCI Flash Capacity (Raw) ⁵	Capacity-Optimized PCI Flash Capacity (Raw) ⁵	Disk Data Capacity (Raw) ⁵
Database Server			2,800,000	2,500,000				7.68 TB
Database Server-Z			1,400,000	1,250,000				7.68 TB
High Capacity (HC) ¹ Storage Server	100 GB/s	500 GB/s	2,800,000	1,000,000	1,280 GB	27.2 TB		264 TB
Extreme Flash (EF) ¹ Storage Server	100 GB/s	500 GB/s	2,800,000	1,000,000	1,280 GB	27.2 TB	122.88 TB	
High Capacity-Z (HC-Z) Storage Server	50 GB/s	250 GB/s	1,400,000	500,000	576 GB	13.6 TB		132 TB
Extended (XT) Storage Server	50 GB/s		600,000	500,000		13.6 TB		264 TB

¹ Actual system performance varies by application.

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

³ 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.

⁴ 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.

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

Exadata Elastic Rack Configurations: Capacity and Performance Metrics (HC, EF, and HC-Z)

Flash Metrics		Maximum SQL Flash and Exadata RDMA Memory Bandwidth ²	Maximum SQL XRMEM Read IOPS ^{1,3}	Maximum SQL Flash Write IOPS ⁴	Performance-Optimized PCI Flash Capacity (Raw) ⁵	Disk Capacity (Raw) ⁵	Capacity-Optimized Flash Capacity (Raw) ⁵
Elastic Configuration 1 (Example)⁶	HC ¹	4,500 GB/s	25,200,000	9,000,000	244.8 TB	2,376 TB	
	EF ¹	4,500GB/s	25,200,000	9,000,000	244.8 TB		1,105.9 TB
Elastic Configuration 2 (Example)⁶	HC ¹	8,500 GB/s	5,600,000	5,000,000	462.4 TB	4,488 TB	
	EF ¹	8,500 GB/s	5,600,000	5,000,000	462.4 TB		2,089 TB
Quarter Rack	HC ¹	1,500 GB/s	5,600,000	3,000,000	81.6 TB	792 TB	
	EF ¹	1,500 GB/s	5,600,000	3,000,000	81.6 TB		368.6 TB
Quarter Rack (Database Server-Z)	HC-Z ¹	750 GB/s	2,800,000	1,500,000	40.8 TB	396 TB	
	HC ¹	1,500 GB/s	2,800,000	2,500,000	81.6 TB	792 TB	
	EF ¹	1,500 GB/s	2,800,000	2,500,000	81.6 TB		368.6 TB

¹ EF = Extreme Flash; HC = High Capacity, HC-Z = High Capacity-Z
² Bandwidth is peak physical scan bandwidth achieved running SQL, assuming no database compression. Effective user data bandwidth is higher when using database compression.
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 Elastic Configuration 1 and Elastic Configuration 2 configurations added as examples of elastic configurations. Elastic Configuration 1 = 9X DB and 9X Storage Servers; Elastic Configuration 2 = 2X DB and 17x HC or 17x EF Storage Servers

Exadata Elastic Rack Configurations: Combined Metrics (HC, EF and HC-Z)

Combined Metrics		Data Capacity (Usable) Normal Redundancy ¹	Data Capacity (Usable) High Redundancy ¹	Maximum Data Load Rate ²
Elastic Configuration 1 (Example)³	HC ¹	983.2 TB	720.3 TB	22.5 TB/hour
	EF ¹	457.6 TB	335.2 TB	22.5 TB/hour
Elastic Configuration 2 (Example)³	HC ¹	1857.2 TB	1360.6 TB	8.8 TB/hour
	EF ¹	864.2 TB	633.1 TB	8.8 TB/hour
Quarter Rack	HC-Z ¹	153.1 TB	120 TB	3.8 TB/hour
	HC ¹	306.1 TB	240.1 TB	7.5 TB/hour
	EF ¹	142.5 TB	111.7 TB	7.5 TB/hour
Quarter Rack (Database Server-Z)	HC-Z ¹	153.1 TB	120 TB	3.8 TB/hour
	HC ¹	306.1 TB	240.1 TB	7.5 TB/hour
	EF ¹	142.5 TB	111.7 TB	7.5 TB/hour

¹ 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.

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

³ Elastic Configuration 1 and Elastic Configuration 2 configurations added as examples of elastic configurations. Elastic Configuration 1 = 9X DB and 9X Storage Servers; Elastic Configuration 2 = 2X DB and 17x HC or 17x EF Storage Servers

Exadata Database Machine Component Environmental Specifications

Metric	Database Server X11M	Database Server-Z X11M	Database Server X11M (1.5TB Memory)	Database Server X11M (3TB Memory)	X11M High Capacity (HC) Storage Server	X11M Extreme Flash (EF) Storage Server	X11M High Capacity-Z (HC-Z) Storage Server	X11M Extended (XT) Storage Server
Height	3.42" (86.9 mm)							
Width	17.52 " (445.0 mm)							
Depth	30.51" (775.0 mm)							
Acoustic Noise (operating)	8.6 B	8.5 B	8.5 B	9.0 B	8.6 B	8.0 B	8.6 B	7.5 B
Weight	53 lb (24 kg)	53 lb (24 kg)	53 lb (24 kg)	53 lb (24 kg)	74 lb (33.6 kg)	60 lb (27.2 kg)	63 lb (28.6 kg)	68 lb (30.8 kg)
Maximum Power Usage	1.2 kW (1.2 kVA)	1.1 kW (1.1 kVA)	1.4 kW (1.4 kVA)	1.4 kW (1.4 kVA)	1 kW (1 kVA)	0.9 kW (0.9 kVA)	0.8 kW (0.8 kVA)	0.5 kW (0.5 kVA)
Typical Power Usage¹	0.8 kW (0.8 kVA)	0.8 kW (0.8 kVA)	1 kW (1 kVA)	1 kW (1 kVA)	0.7 kW (0.7 kVA)	0.6 kW (0.6 kVA)	0.5 kW (0.5 kVA)	0.4 kW (0.4 kVA)
Cooling at Maximum Usage	4,054 BTU/hour	3,783 BTU/hour	4,726 BTU/hour	4,705 BTU/hour	3,474 BTU/hour	3,037 BTU/hour	2,570 BTU/hour	1,825 BTU/hour
	4,277 kJ/hour	3,991 kJ/hour	4,986 kJ/hour	4,964 kJ/hour	3,665 kJ/hour	3,204 kJ/hour	2,712 kJ/hour	1,926 kJ/hour
Cooling at Typical Usage	2,838 BTU/hour	2,648 BTU/hour	3,308 BTU/hour	3,294 BTU/hour	2,431 BTU/hour	2,126 BTU/hour	1,799 BTU/hour	1,278 BTU/hour
	2,994 kJ/hour	2,794 kJ/hour	3,490 kJ/hour	3,475 kJ/hour	2,565 kJ/hour	2,243 kJ/hour	1,898 kJ/hour	1,348 kJ/hour
Airflow at Maximum Usage²	188 CFM	175 CFM	219 CFM	218 CFM	161 CFM	141 CFM	119 CFM	85 CFM
Airflow at Typical Usage²	131 CFM	123 CFM	153 CFM	152 CFM	113 CFM	98 CFM	83 CFM	59 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 ¹ Typical power usage varies by application load ² Airflow must be front-to-back.								

Exadata Database Machine Environmental Specifications

Metric	Elastic Configuration 1 (Example) ^{3,4}	Elastic Configuration 2 (Example) ^{3,4}	Quarter Rack ⁴	Quarter Rack (Database Server-Z and High Capacity-Z)
Height	78.74 in (2000 mm)			
Width	23.62 in (600 mm)			
Depth	47.12 in (1197 mm)			
High Capacity Storage Servers				
Acoustic Noise (operating)	9.7 B	9.9 B	9.4 B	9.4 B
Weight	1904.7 lb (864 kg)	2145 lb (972.9 kg)	937.8 lb (425.4 kg)	904.8 lb. (410.4 kg)
Maximum Power Usage	22.5 kW (22.9 kVA)	20.9 kW (21.3 kVA)	6.7 kW (6.8 kVA)	5.3 kW (5.4 kVA)
Typical Power Usage ¹	15.7 kW (16.1 kVA)	14.6 kW (14.9 kVA)	4.7 kW (4.8 kVA)	3.7 kW (3.8 kVA)
Cooling at Maximum Usage	76,702 BTU/hour	71,409 BTU/hour	22,779 BTU/hour	18,185 BTU/hour
	80,920 kJ/hour	75,337 kJ/hour	24,032 kJ/hour	19,185 kJ/hour
Cooling at Typical Usage	53,691 BTU/hour	49,987 BTU/hour	15,946 BTU/hour	12,730 BTU/hour
	56,644 kJ/hour	52,736 kJ/hour	16,823 kJ/hour	13,430 kJ/hour
Airflow at Maximum Usage ²	3551 CFM	3306 CFM	1055 CFM	842 CFM
Airflow at Typical Usage ²	2486 CFM	2314 CFM	738 CFM	589 CFM
Extreme Flash Storage Servers				
Acoustic Noise (operating)	9.7 B	9.4 B	9.2 B	
Weight	1778.7 lb (806.8 kg)	1907 lb (865 kg)	895.8 lb (406.3 kg)	
Maximum Power Usage	21.3 kW (21.8 kVA)	18.8 kW (19.1 kVA)	6.3 kW (6.4 kVA)	
Typical Power Usage ¹	14.9 kW (15.2 kVA)	13.1 kW (13.4 kVA)	4.4 kW (4.5 kVA)	
Cooling at Maximum Usage	72,771 BTU/hour	63,984 BTU/hour	21,469 BTU/hour	
	76,773 kJ/hour	67,504 kJ/hour	22,650 kJ/hour	
Cooling at Typical Usage	50,940 BTU/hour	44,789 BTU/hour	15,028 BTU/hour	
	53,741 kJ/hour	47,253 kJ/hour	15,855 kJ/hour	
Airflow at Maximum Usage ²	3369 CFM	2962 CFM	994 CFM	
Airflow at Typical Usage ²	2358 CFM	2074 CFM	696 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

¹ Typical power usage varies by application load

² Airflow must be front-to-back.

³ Elastic Configuration 1 and Elastic Configuration 2 configurations added as examples of elastic configurations. Elastic Configuration 1 = 9X DB and 9X Storage Servers ; Elastic Configuration 2 = 2X DB and 17X HC or 17X EF Storage Servers

⁴ Environmental specifications based on Exadata X11M database server with 1.5 TB memory

Exadata Database Machine Regulations and Certifications

Regulations ^{1,2,3}	Product Safety:	UL/CSA 60950-1, EN 60950-1, IEC 60950-1 CB Scheme with all country differences
		UL/CSA 62368-1, EN 62368-1, IEC 62368-1 CB Scheme with all country differences
	EMC	
	Emissions:	FCC CFR 47 Part 15, ICES-003, EN55032, KS C 9832, EN61000-3-11, EN61000-3-12
	Immunity:	EN55035, KS C 9835
Certifications ^{2,3}	North America (NRTL), CE (European Union), International CB Scheme, HSE Exemption (India), BSMI (Taiwan), KC (Korea), RCM (Australia), VCCI (Japan), UKCA (United Kingdom)	
European Union Directives ³	2014/35/EU Low Voltage Directive, 2014/30/EU EMC Directive, 2011/65/EU RoHS Directive, 2012/19/EU WEEE Directive	

¹ All standards and certifications referenced are to the latest official version. For additional detail, please contact your sales representative.

² Other country regulations/certifications may apply.

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

Exadata Database Machine Support Services

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

Optional Customer Supplied Ethernet Switch Installation in Exadata Database Machine X11M

Each Exadata Database Machine X11M 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 Control of Flash Cache Size per Database
- Unique In-Memory OLTP Acceleration
- Unique Undo-Block Remote RDMA Read
- Unique Support for 4096 Pluggable Databases per Container Database with Multitenant Option

Exadata and Database Software Features – Artificial Intelligence

- Unique Automatic parallelization and offload of AI Vector Search to storage
- Unique Filtering of rows based on vector distance
- Unique Column Projection of AI Vector Search results including vector distance column from storage
- Unique Vector Distance Computation on Storage Servers
- Unique per-storage server Adaptive Top-K calculation

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 7. 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

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2025, 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.