

Oracle Exadata Database Machine X9M-8

The Oracle Exadata Database Machine is engineered to deliver dramatically better performance, cost effectiveness, and availability for Oracle databases. Exadata features a modern cloud-enabled architecture with scale-out high-performance database servers, scale-out intelligent storage servers with state-of-the-art PCIe flash, leading-edge storage cache using persistent memory, and cloud scale RDMA over Converged Ethernet (RoCE) internal fabric that connects all servers and storage. Unique algorithms and protocols in Exadata implement database intelligence in storage, compute, and networking to deliver higher performance and capacity at lower costs than other platforms, for all types of modern database workloads. The Exadata Database Machine X9M-8 excels at high-end online transaction processing (OLTP), multi-rack Data Warehouses, and In-Memory Database workloads. It is ideally suited for applications that desire large scale-up servers with large memory capacity.

Simple and fast to implement, the Exadata Database Machine X9M-8 powers and protects your most important databases. Exadata can be purchased and deployed on premises as the ideal foundation for a private database cloud.

Engineered For Fast And Reliable Deployment

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

All Exadata Database Machines are identically configured, therefore customers benefit from the experience of thousands of other customers' Exadata Database Machine deployments. Customer machines are also identical to the machines Oracle Support uses for problem identification and resolution, the machines Oracle Development uses for development and testing of the Oracle Database.

Exadata is the most thoroughly tested and tuned platform for running Oracle Database.



Oracle Exadata Database Machine X9M-8

Key Features

- Up to 576 CPU cores and 18 TB memory per rack for database processing
- Up to 448 CPU cores per rack dedicated to SQL processing in storage
- Up to 21 TB of Persistent Memory Acceleration per rack
- 100 Gb/sec RoCE Network
- Complete redundancy for high availability
- From 2 to 3 Database Servers per rack
- From 3 to 14 Storage Servers per rack
- Up to 716 TB of flash capacity (raw) per rack
- Up to 3.0 PB of disk capacity (raw) per rack

Any application that uses the Oracle Database today can be seamlessly migrated to the Exadata Database Machine, with no changes to the application. Likewise, any Oracle Database can also be easily migrated off Exadata, eliminating “lock-in” concerns.

Extreme System Scalability and Growth with Elastic Configurations

The Exadata Database Machine uses a scale-out architecture for both database servers and storage servers. As workloads grow, database CPUs, storage, and networking can be added to an Exadata Database Machine to scale without bottlenecks. The **architecture expands from small to extremely large configurations to accommodate workloads of any size.**

In Exadata X9M, a high-bandwidth low-latency active-active 100 Gb/sec RDMA over Converged Ethernet (RoCE) Network Fabric connects all the components inside an Exadata Database Machine. Specialized database networking protocols deliver much lower latency and higher bandwidth than is possible with generic communication protocols for **faster response time for OLTP operations and higher throughput for analytic workloads.** External connectivity to the Exadata Database Machine is via standard 10 Gb/sec or 25 Gb/sec Ethernet.

Exadata Database Machine is the most versatile database platform. The Exadata X9M-8 Database Machine uses powerful database servers, each with eight 24-core x86 processors and 3 TB of memory (expandable up to 6 TB). Exadata also uses scale-out, intelligent storage servers available in two configurations – High Capacity (HC) or Extreme Flash (EF):

- HC Storage Servers have four NVMe PCI Flash cards each with 6.4 TB (raw) Exadata Smart Flash Cache and twelve 18 TB 7,200 RPM disks.
- EF Storage Servers have an all-flash configuration with eight NVMe PCI Flash drives, each with 6.4 TB (raw) storage capacity.
- Exadata X9M HC and EF Storage include persistent memory, further enhancing performance. Each server is populated with twelve 128 GB Intel® Optane™ Persistent Memory modules as a new caching tier between DRAM and flash.

The minimal configuration of an Exadata Database Machine consists of two database servers and three storage servers, which **can be elastically expanded by adding one more database and/or storage servers within the same rack.** Elastic configurations provide a flexible and efficient mechanism to meet any size business need.

In addition to expanding within a rack, **multiple RoCE based Exadata can be interconnected using the integrated RoCE network fabric** to form even larger configurations. These racks can be of either Exadata X8M or Exadata X9M generations. For example, a system composed of four racks of Exadata X9M 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. Scaling out is easy, as Oracle Real Application Clusters (RAC) can

Key Benefits

- Pre-configured, pre-tested system optimized for all database applications
- Uncompressed I/O bandwidth of up to 1 TB/sec per full rack from SQL
- Ability to perform up to 15M 8K database read I/O operations, or 6.75M 8K Flash write I/O operations per second in a single rack
- Easily add compute or storage servers to meet the needs of any size application
- Scale by connecting multiple Exadata Database Machine X9M-8 racks or Exadata X9M Storage Expansion Racks. Up to 12 racks can be connected by simply adding RoCE cables and internal switches. Larger configurations can be built with external RoCE switches

dynamically add more processing power, and Automatic Storage Management (ASM) can dynamically add more storage capacity.




Figure 1 Elastic Scale-out to Multi-rack Exadata

When large storage capacity is required, the **Oracle Exadata X9M Storage Expansion Rack** is available. The Exadata Storage Expansion Rack enables customers to grow the storage capacity, OLTP IOPS, and scan throughput of any Exadata Database Machine. It is designed for database deployments with very large amounts of data, including historical or archive data, backups, documents, images, XML, JSON, and LOBs. The storage expansion rack connects to the Exadata Database Machine using the integrated RoCE network fabric and is configured with a few simple commands, as there are no LUNs or mount points. The starting configuration of the Oracle Exadata Storage Expansion Rack consists of four storage servers and can be expanded with additional storage servers.

Groundbreaking RDMA Based Network Fabric

The Exadata X9M release uses the same ultra-fast cloud scale networking fabric that was introduced in Exadata X8M, RDMA over Converged Ethernet (RoCE). RDMA (Remote Direct Memory Access) allows one computer to directly access data from another without Operating System or CPU involvement, for high bandwidth and low latency. The network card directly reads/writes memory with no extra copying or buffering and with very low latency. RDMA is an integral part of the Exadata high-performance architecture, and has been tuned and enhanced over the past decade, underpinning several Exadata-only technologies such as Exafusion Direct-to-Wire Protocol and Smart Fusion Block Transfer. As the RoCE API infrastructure is identical to InfiniBand's, **all existing Exadata performance features are available on RoCE.**

Shared Persistent Memory Acceleration

Exadata X8M introduced persistent memory (PMEM) data and commit accelerators in front of flash cache, enabling orders of magnitude lower latency

Related Products

- Oracle Database Exadata Cloud Service
- Oracle Database Exadata Cloud at Customer
- Oracle Exadata Database Machine X9M-2
- Oracle Exadata Storage Expansion Rack X9M
- Oracle Exadata Storage Server X9M-2
- Oracle Exadata Database Server X8M-8
- Oracle Database Enterprise Edition 11g, 12c, 18c, 19c and 21c
- Real Application Clusters
- Partitioning
- Multitenant
- Database In-Memory
- Advanced Compression
- Advanced Security
- Active Data Guard
- GoldenGate
- Real Application Testing
- OLAP
- Enterprise Manager
- Oracle Linux

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

accessing remotely stored data. Persistent memory is a modern silicon technology, adding a distinct storage tier of performance, capacity, and price between DRAM and Flash. As the persistent memory is physically present on the memory bus of the storage server, reads perform at memory speed, much faster than flash. Writes are persistent, surviving power cycles, unlike DRAM. By utilizing RDMA to access persistent memory remotely, **Exadata Persistent Memory Data and Commit Accelerators bypass the network and I/O stack, eliminating expensive CPU interrupts and context switches, reducing latency by 10x**, from 200 μ s to less than 19 μ s. Smart Exadata System Software also ensures data is mirrored across storage servers, which provides additional fault-tolerance. Exadata's unique end-to-end integration between Oracle Database and Exadata Storage Servers automatically caches the hottest data blocks efficiently between the database buffer cache, persistent memory, and flash cache. Adding persistent memory to the storage tier means the aggregate performance of this new cache tier can be dynamically used by any database on any server. This is a significant advantage over general-purpose storage architectures, which preclude sharing storage resources across database instances. Another smart Exadata System Software feature boosts log write performance. Log write latency is critical for OLTP performance, a faster log write means faster commit times. Inversely, any slowdown of log writes can cause the database to stall. With the RoCE based Exadata, **Exadata Persistent Commit Accelerator** automatically enables the database to issue a one-way RDMA log write to persistent memory. RDMA and persistent memory technologies allow the log write to occur without acknowledgement, and smart software distributes the write across multiple servers for resilience. This leads to a performance increase for log write operations.

Security and management of this tier are also automated. Persistent memory is configured automatically at installation time, with no user interaction required. Hardware monitoring is pre-configured. Persistent memory is only accessible to databases using database access controls, ensuring end to end security of data. Deploying persistent memory in Exadata X9M is so simple, it's transparent.




Figure 2. Intel® Optane™ Persistent Memory modules

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

Exadata **Extreme Flash (EF) Storage Server** is the foundation of a database-optimized all-flash Exadata Database Machine. Each EF Storage Server contains eight 6.4 TB Flash Accelerator F640v3 NVMe PCI Flash drives, offering 51.2 TB raw flash capacity per EF Storage Server, with an expected endurance of 8 years or more for typical database workloads. Exadata delivers ultra-high performance by placing these flash devices directly on the high speed PCIe4 interface rather than behind slow disk controllers. Exadata X9M adds the shared persistent

“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

memory acceleration tier, twelve 128 GB Intel® Optane™ Persistent Memory modules in front of flash to boost performance even more.

Exadata X9M uses a combination of scale-out storage, RDMA over Converged Ethernet networking, database offload, persistent memory accelerator and PCIe Flash to deliver extremely high performance from memory and flash. A single rack configuration of Exadata Database Machine X9M-8¹ can achieve up to **15.0 Million random 8K database read and 6.75 Million random 8K flash write I/O operations per second (IOPS)**.

For data warehouse environments that require the highest performance, Exadata X9M Extreme Flash storage servers are capable of **scanning up to 75GB/s per server for a total of 1TB/s²** on a traditional full rack configuration³.

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



Figure 3. Flash Accelerator PCIe Card

High Capacity Storage Server: Tiered Disk, Flash, And Persistent Memory Deliver Cost Of Disk With Shared Memory Performance

The second Exadata storage option is the **High Capacity (HC) Storage Server**. This server includes twelve 18 TB SAS disk drives with 216 TB total raw disk capacity. It also has four 6.4 TB flash accelerator cards for a total raw capacity of 25.6 TB of flash memory. Exadata X9M adds the shared persistent memory acceleration tier, twelve 128 GB Intel® Optane™ Persistent Memory modules in front of flash to boost performance even more. Deployed using smart software, **Exadata Persistent Memory Data Accelerator**, only the hottest database blocks are automatically cached in this new tier. Accessible over RDMA directly from the database delivers the highest I/O rates at an extremely low latency.

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 Persistent Memory Data Accelerator to deliver the best performance. Exadata Smart Flash Cache is used in conjunction with the Exadata Persistent Memory Data Accelerator to automatically cache frequently accessed data while keeping infrequently accessed data on disk, delivering the

“Oracle Exadata is a fantastic solution for us since it provides the high performance and availability that we need for mission-critical applications and fast backup and recovery. By far, Oracle has the best database engine on the planet.”

Drew Diamond
Technology Manager
Tractor Supply Co.

¹Elastic configuration with 3x Exadata X8M-8 Database Servers and 11x Exadata X9M-2 Extreme Flash Storage Servers

²Smart Scan rate measured at 1050 GB/sec

³Traditional full rack configuration with 2x Exadata X8M-8 Database Servers and 14x Exadata X9M-2 Extreme Flash Storage Servers

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 will negatively affect overall performance**. For example, if large write I/Os caused by backups or large table scans are likely to disrupt higher priority OLTP or scan operations, those large I/Os will bypass the flash cache and go straight to disk. Otherwise, Exadata System Software will utilize additional spare flash capacity and I/O bandwidth to optimize performance by caching these I/Os. Administrators can also manually (optional) 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. For example, a traditional full rack⁴ often has an effective flash capacity close to the usable disk capacity of 900 TB.

Exadata Smart Flash Cache also caches database block writes using Exadata Write Back Flash Cache technology. Write caching eliminates disk bottlenecks in large scale OLTP and batch workloads. The flash write capacity of a single full rack Exadata Database Machine X9M-8⁴ is **6 Million 8K flash write I/O operations per second (IOPS)**. 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 RAM, persistent 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 arrays. These flash-only arrays deliver better performance but cannot match the cost advantages of Exadata's smart tiering of data between disk and flash, as the overall size of data that can benefit from flash is limited to the size of expensive flash. And these flash arrays are unable to benefit from Exadata's unique database-aware storage optimization technologies. Generic data deduplication provided by some flash arrays is effective for Virtual Desktop Infrastructure environments, but not for databases.

Exadata not only delivers much more capacity than generic all-flash arrays, it also delivers better performance. Flash-only storage arrays cannot match the throughput of Exadata's integrated and optimized architecture with 100 Gb/sec active-active RDMA over Converged Ethernet based scale-out network, fast PCIe Flash, offload of data intensive operations to storage, and algorithms. All of which have been specifically optimized for Oracle Database.

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

A third Exadata storage option is the **Extended (XT) Storage Server**. Each Exadata XT Storage Server includes twelve 18 TB SAS disk drives with 216 TB total raw disk capacity.

“We take a certain amount of pride in the fact that we’re running the largest SAP ECC database on Oracle in the world. Our system that does the bulk of the delivery processing for our systems every day was doing about 1 million orders per day. And today, we’re averaging more around 3 million. And we’ve done all that with Exadata”

Mike White
Technology Manager
AmerisourceBergen Corporation

⁴ Traditional full rack with 2x Exadata X8M-8 Database Servers and 14x Exadata X9M-2 High Capacity Storage Servers

This storage option extends the operational and management benefits of Exadata to rarely accessed data that must be kept online. Exadata's Extended (XT) Storage Server is:

- Efficient – The XT server offers the same high capacity as the HC Storage server, including Hybrid Columnar Compression
- Simple – The XT server adds capacity to Exadata while remaining transparent to applications, transparent to SQL, and retains the same operational model
- Secure – The XT server enables customers to extend to low-use data the same Exadata security model and encryption used for online data
- Fast and Scalable – Unlike other low-use data storage solutions, the XT server is integrated to the Exadata fabric, for fast access and easy scale-out
- Compatible – The XT server is just another flavor of Exadata Storage server – you can just add XT servers to any Exadata rack
- Inexpensive – Does not require Exadata Storage Software licenses.

With Exadata Extended (XT) Storage Server, enterprises can meet their long-term data retention compliance requirements with the same trusted and continually validated Exadata solution, avoiding the operational risks and costs of managing information lifecycle across multiple platforms.

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

Accelerating Database Processing with Smart System Software

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

Exadata System Software enables Exadata's unparalleled performance without any of the bottlenecks of traditional storage, by implementing a unique highly efficient database-optimized storage infrastructure on the Exadata Storage Server. Each High Capacity and Extreme Flash Exadata Storage Server has two 16-core x86 processors that are used to offload database processing. The CPUs in the storage servers do not replace database CPUs. Instead 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**

“The highly available design of the Exadata infrastructure provides a future-proof, performant platform for DBaaS that enables high-density consolidation of Oracle Databases, with a high degree of embedded automation to reduce operational costs.”

Martin Lehner

Head of Unix, Middleware and Database Services
T-Systems Alpine

columns that are directly relevant to a query are sent to the database servers.

For example, if a query is executed to identify the customers who placed sales orders over \$1000 in the month of March, an Exadata system will offload the scanning of the table to the Exadata Storage Servers where filters extract only the relevant customer information for March with a minimum \$1000 spend and returns this reduced quantity of data to the database. This 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. Maintained automatically in-memory on the storage server, storage indexes 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 if rows with the specified column value might exist in a region of disk on the storage server. If the column value doesn't exist in the Storage Index, then scan I/O in that region for that query is skipped. Storage Indexes make many SQL operations run dramatically faster because large numbers of I/O operations are automatically replaced by a few in-memory lookups.

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

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

To further accelerate OLTP workloads Exadata X8M and X9M High Capacity and Extreme Flash Storage Servers include the **Exadata Persistent Memory Commit Accelerator**. Configured automatically, it eliminates interaction with the operating system and the overhead of transferring data via the standard I/O path, enabling the database to perform log writes via RDMA direct to the persistent memory buffer of multiple storage servers in parallel.

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

“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

continuously learns and tunes the database as the underlying data model or usage patterns change.

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

Optimizing Storage Use and I/O Through Compression

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

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

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

On OLTP systems, Hybrid Columnar Compression can be used to compress older, less active data while newer, more active and update-intensive data can be compressed using Advanced Row Compression. Oracle Database Release 18c and above provides the ability to change the type of compression used by individual table partitions online (even if there are global indexes on the table), to ensure seamless tiering across different compression types as data ages and becomes less active.

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

“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

accelerates reporting and analytic queries while maintaining excellent performance for OLTP-style single row lookups.

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

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

Exadata uniquely implements **In-Memory columnar formats in Flash Cache.** This feature extends the Exadata Columnar Flash Cache by automatically transforming data into In-Memory columnar formats as it is loaded into flash cache. Smart Scans also process multiple column values with a single instruction by leveraging ultra-fast Single Instruction Multiple Data (SIMD) Vector instructions. Smart Scan results are passed back to the database server in Oracle Database In-Memory formats, further reducing the load on database server CPUs. The effect is to seamlessly extend the In-Memory columnar store size from DRAM capacity in the database server to flash capacity in storage servers. An Exadata X9M-8 Full Rack HC⁵ has 360 TB of Flash Cache, capable of servicing some of the largest in-memory workloads. Databases not using Oracle Database In-Memory still benefit from Exadata Columnar Flash Cache without the vector processing optimizations.

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

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

Enterprise-Class Security with Extreme Performance

Exadata Database Machine is the world's most secure database machine. Building on the high security capabilities in the Oracle Database such as

“As a long-term Exadata customer, we are convinced by this platform and therefore decided to refresh our systems with the new Exadata X8M technology, which we expect us to be set for further data growth.”

Kaan Konak
Head of IT Governance, System & Infrastructure and Next Generation Technologies
Aksigorta

⁵ Traditional full rack with 2x Exadata X8M-8 Database Servers and 14x Exadata X9M-2 High Capacity Storage Servers

Transparent Data Encryption (TDE), Exadata **uniquely moves decryption processing from database server software into the Exadata Storage Server hardware**. Exadata storage leverages hardware decryption and compression together to provide the highest performance secure databases. **Encryption occurs after the data is compressed so that the cost of decryption is decreased by the degree of compression**. By leveraging both technologies, Exadata is able to query fully encrypted and compressed databases with minimal overhead at a rate of hundreds of gigabytes of (original) user data per second. 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 the task of ensuring the security of each individual software and hardware component, and ensuring that security is maintained across the full product stack. **Oracle delivers full stack security in the Exadata Database Machine.**

Exadata systems are designed, manufactured, and delivered to customers using a defense-in-depth approach, increasing the security posture of the system. Exadata systems are built using Oracle designed servers for both the database and storage tiers. The in-house design and development of the servers not only enables the implementation of features that are unique to Exadata but also enables tight control over the security of the design. This focus on security extends to the global supply chain of Oracle. Exadata security begins at power-up time with **Secure Boot**, which ensures that the system UEFI firmware only allows the execution of cryptographically signed boot loaders that the system recognizes as trustworthy. These signatures are verified with every reboot of the server **preventing malware from hiding embedded code** in the boot chain. The operating system that is 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 that are required to run the Oracle Database, and eliminates unnecessary packages. This serves to **minimize the attack surface** and results in increased security hardening of the system. Exadata leverages the ksplice capability of Oracle Linux to **apply security updates while the OS stays online**.

The disk and flash technologies used in Exadata X9M enable Stored Data Encryption. In Stored Data Encryption, disk and flash storage devices encrypt all user data as it enters the devices. Exadata's Secure Erase feature leverages this capability when an Exadata is re-purposed or decommissioned to instantly erase all user data present on storage devices by changing the encryption keys used to encrypt the user data. With Secure Erase, because the previous encryption key is deleted, there is no need to worry about latent data left on storage devices due to over-provisioning or sector sparing.

Exadata security has been probed and evaluated by hundreds of leading banks, telecoms, and government organizations worldwide. The security findings of all these evaluations have been incorporated into the Exadata standard

“As part of infrastructure modernization, FedPat chose Oracle Exadata X8M as its new core system for critical database applications. We experienced over 50% performance improvement in data warehousing and nearly 70% reduction in backup times”

Diego Calderon
CIO
Federación Patronal

configuration. Therefore, Exadata benefits from scrutiny both by Oracle Security experts and by hundreds of industry security experts around the world.

Mission Critical High Availability

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

Exadata in an MAA configuration is recognized by the analyst firm IDC as a system that **delivers at least 5-nines availability** and is categorized in the IDC AL4 fault-tolerant market segment⁶.

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 implements a unique RDMA-based **sub-second node death detection**, leading to virtual elimination of application brownout conditions.

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

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

Martin McGeough
Database Technical Architect
Vocalink

⁶ Worldwide AL4 Server Market Shares, 2019: Fault-Tolerant Systems Become Digital Transformation Platforms, IDC, Paul Maguranis Peter Rutten, July 2020

subsystem, Exadata redirects the I/O issued by the database server to another storage server.

Exadata automates monitoring of CPU, memory, input/output subsystems, file system, and network. This automation combines machine learning techniques with the deep lessons learned from thousands of mission critical real-world deployments. For example, Exadata can detect that anomalous use of system resources that negatively impact database performance and automatically identifies the offending process and issues an alert – all without any manual intervention.

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

Ideal Platform for Database as a Service

The Exadata Database Machine can host many databases, enabling database consolidation or a sophisticated Database as a Service private cloud. Multi-database environments inherently have diverse, complex, and unpredictable workloads mixing OLTP, analytics, and batch operations with sequential and random-access patterns. Exadata's ability to **run mixed database workloads with industry leading scalability and performance** makes it an ideal consolidation platform.

Multi-database environments create an inherent risk that one database will consume too many resources and therefore impact the quality of service of other databases. The Exadata Database Machine **uniquely** provides **end-to-end prioritization** of 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 can also be applied when multiple databases and multiple clusters are deployed on one Exadata rack, as is typical in a consolidated private cloud.

In X9M, 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.

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

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

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

Fast Deployment of Development and Test Databases with Exadata Snapshots

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

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

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

Comprehensive System Management

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

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

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

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

alerts to both administrators and Oracle support when faulty components are detected.

Highest Level of Service

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

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

IT Agility

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

Dramatically Lower Costs

Due to the extreme performance, high storage capacity, and unique compression capabilities delivered by the Exadata Database Machine, workloads that would require very large traditional hardware systems can be run on much smaller Exadata systems. In sizing exercises it is typical to see a 2-4x reduction in Exadata system size compared to a traditional system.

Exadata provides a huge RAM, flash, and disk footprint for large data sets. Raw disk storage on an Exadata full rack can exceed 3 PB (Petabytes) while raw flash storage can be up to 716 TB. Hybrid Columnar Compression also increases the effective storage and memory capacity by a factor of 10. By intelligently moving active data across disk, flash, and memory tiers, Exadata simultaneously delivers the highest performance and the lowest cost.

Exadata has the **unique** ability to consolidate many databases supporting multiple workloads in a single cloud platform. High-end OLTP, analytics, batch, reporting, and backups can all run simultaneously within and across databases

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

Roland Schiller
Architect
Swiss Re

with extreme performance. **The extreme performance and capacity of Exadata enables very large numbers of databases and workloads to be consolidated on Exadata.** Consolidating databases on Exadata reduces system hardware cost, software cost, and greatly reduces ongoing operations cost.

The uniformity of Exadata Database Machine configurations results in large cost savings. **Exadata standardizes not just technologies, but also integration, testing, security, hardening, tuning, and support.** Customers deploy Exadata systems much faster and with a lot less labor than traditional systems. Low level tuning, integration, and maintenance is reduced or eliminated. All Exadata users run a configuration that is identical to thousands of other users, and is identical to Oracle's internal configurations, it is far less likely that issues will be encountered, and issue resolution is quicker and simpler reducing both operations cost and downtime cost. When issues do occur, customers have to deal with one supplier – Oracle, as the entire system – hardware, firmware, operating system, 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 further increasing benefits.

Capacity-on-Demand Software Licensing

An Exadata X8M-8 Database Server has a substantial amount of compute capacity with eight 24-core x86 processors (192 cores in total). The Capacity-on-Demand feature allows a number of cores per database server to be turned off during the hardware installation, leaving at least 56 cores enabled. As your workload grows and more cores are needed, Capacity-on-Demand can be used to re-enable cores and license software 8 cores at a time. This pay-as-you-grow approach to software licensing is another way in which Exadata helps to align costs with business growth.

Exadata Business Benefits

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

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

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

“Oracle's continued engineering focus on Exadata performance has benefited our business over the last 5 years. We eagerly look forward to the next chapter of our Exadata journey.”

Kaan Konak
Head of IT Governance, System & Infrastructure and Next Generation Technologies
Aksigorta

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

Exadata enables:

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

Conclusion

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

“With Exadata, we were able to achieve 27[x] faster performance on our digital banking throughput for financial transactions. And in fact the results were amazing and exceeded our expectations.”

Ghassan Sarsak
CTO & CIO
ICS Financial Systems

Exadata Server Hardware ^{1,2}

SERVER TYPE	CPU	MEMORY	DISK	FLASH	NETWORK
Database Server	8 x 24-core Intel® Xeon® 8268 processors (2.9GHz)	3 TB (default) to 6 TB (max)	None	2 x 6.4 TB NVMe Flash cards	<ul style="list-style-type: none"> 8 x 10/25 Gb copper Ethernet ports (client) 8 x 1/10 Gb copper Ethernet ports (1 used for host ADMIN) 8 x 100 Gb QSFP28 RoCE Fabric ports 1 x ILOM Ethernet port
Storage Server High Capacity (HC)	2 x 16-core Intel® Xeon® 8352Y processors (2.2 GHz)	256 GB 1.5 TB Persistent Memory	12 x 18 TB 7,200 RPM disks	4 x 6.4 TB NVMe PCIe4.0 Flash cards	<ul style="list-style-type: none"> 2 x 100 Gb QSFP28 RoCE Fabric ports 1 x 1 Gb copper Ethernet port (mgmt) 1 x ILOM Ethernet port
Storage Server Extreme Flash (EF)	2 x 16-core Intel® Xeon® 8352Y processors (2.2 GHz)	256 GB 1.5 TB Persistent Memory	None	8 x 6.4 TB NVMe PCIe4.0 Flash cards	
Storage Server Extended (XT)	1 x 16-core Intel® Xeon® 8352Y processor (2.2 GHz)	96 GB	12 x 18 TB 7,200 RPM disks	None	

¹ All servers include redundant hot swappable fans and power supplies

² Table includes individually purchasable servers only

Exadata Rack Configurations ^{1, 2}

RACK SIZE	DATABASE SERVERS AND CORES	STORAGE SERVER AND CORES	HIGH CAPACITY STORAGE SERVER CAPACITY (RAW)		EXTREME FLASH STORAGE SERVER CAPACITY (RAW)
Half Rack ³	2 x servers, 384 cores	3 x servers, 96 cores for SQL offload	648 TB disk, 76.8 TB Flash, 4.5 TB Persistent Memory	or	153.6 TB Flash, 4.5 TB Persistent Memory
Full Rack ⁴	2 x servers, 384 cores	14 x servers, 448 cores for SQL offload	3024 TB disk, 358.4 TB Flash, 21 TB Persistent Memory		716.8 TB Flash, 21 TB Persistent Memory
+Database Servers	Up to 3x servers ³ , 576 cores max per rack	n/a	n/a		n/a
+Storage Servers	n/a	Up to 14x servers ³ , 448 cores max per rack	3024 TB disk, 358.4 TB Flash, 21 TB Persistent Memory		716.8 TB Flash, 21 TB Persistent Memory maximum per rack

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. Included replacement parts kit contains:

- 1 x 6.4 TB NVMe PCI Flash card and 1 x 18 TB High Capacity disk, or
- 1 x 6.4 TB NVMe PCI Flash card

² Elastic configurations allow adding database or storage servers to a half rack to achieve the exact ratio of compute to storage that the application needs. A full rack elastic configuration cannot exceed 16 servers and 42 RU (Rack Units). Database Servers = 5 RU, Storage Servers = 2 RU

³ Maximum number of database servers allowed in an elastic configuration is 3. Maximum number of storage servers allowed in an elastic configuration is 14.

⁴ Full Rack configurations added as examples of elastic configurations.

Other Elastic Expansion Options

MULTI-RACK CONNECTION	CONNECT ANY COMBINATION OF UP TO 12 EXADATA DATABASE MACHINE RACKS OR EXADATA STORAGE EXPANSION RACKS VIA THE ROCE NETWORK FABRIC. 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 READ IOPS ^{1,3}	MAXIMUM SQL WRITE IOPS ⁴	PERSISTENT MEMORY CAPACITY (RAW) ⁵	PCI FLASH CAPACITY (RAW) ⁵	DISK DATA CAPACITY (RAW)
Database Server	n/a	5,000,000	3,000,000	n/a	n/a	n/a
Storage Server High Capacity (HC) ¹	45 GB/s	2,300,000	614,000	1.5 TB	25.6 TB	216 TB
Storage Server Extreme Flash (EF) ¹	75 GB/s	2,300,000	614,000	1.5 TB	51.2 TB	n/a
Storage Server Extended (XT) ¹	n/a	n/a	n/a	n/a	n/a	216 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 database compression is used.

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 Typical Rack Configurations: Flash Capacity and Performance Metrics (HC & EF)

FLASH METRICS		MAXIMUM SQL FLASH BANDWIDTH ²	MAXIMUM SQL PMEM READ IOPS ^{1,3}	MAXIMUM SQL FLASH WRITE IOPS ⁴	PCI FLASH CAPACITY (RAW) ⁵
Full Rack ⁶	HC ¹	630 GB/s	10,000,000	6,000,000	358.4 TB
	EF ¹	1050 GB/s	10,000,000	6,000,000	716.8 TB
Half Rack ⁶	HC ¹	135 GB/s	6,900,000	1,842,000	76.8 TB
	EF ¹	225 GB/s	6,900,000	1,842,000	153.6 TB

¹ EF = Extreme Flash; HC = High Capacity; PMEM = Persistent Memory

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

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

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

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

⁶ Half and Full rack configurations added as examples of elastic configurations. Half = 2x DB Servers, 3x Storage Servers; Full = 2x DB Servers, 14x Storage Servers

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

FLASH METRICS	MAXIMUM SQL DISK BANDWIDTH ¹	MAXIMUM SQL DISK IOPS ²	DATA CAPACITY (RAW) ³
Full Rack⁴	25 GB/s	36,000	3024 TB
Half Rack	5.4 GB/s	7,800	648 TB

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

² 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.

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

⁴ Full rack configuration added as examples of elastic configurations. Full = 2x DB Servers, 14x Storage Servers.

Exadata Typical Rack Configurations: Combined Metrics (HC & EF)

COMBINED METRICS	DATA CAPACITY (USABLE) NORMAL REDUNDANCY ¹	DATA CAPACITY (USABLE) NORMAL REDUNDANCY 1	DATA CAPACITY (USABLE) - HIGH REDUNDANCY ¹	MAXIMUM DATA LOAD RATE ²
Full Rack³	HC ¹	1225.8 TB	898.0 TB	35.0 TB/hour
	EF ¹	281.8 TB	206.5 TB	35.0 TB/hour
Half Rack	HC ¹	245.4 TB	192.4 TB	7.5 TB/hour
	EF ¹	56.4 TB	44.2 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 Full rack configuration added as example of elastic configuration. Full = 2x DB Servers, 14x Storage Servers.

Exadata Database Machine Component Environmental Specifications

METRIC	EXADATA DATABASE SERVER X8M-8	EXADATA STORAGE SERVER X9M-2 HIGH CAPACITY (HC)	EXADATA STORAGE SERVER X9M-2 EXTREME FLASH	EXADATA STORAGE SERVER X9M-2 EXTENDED (XT)
Height	8.63 in (219.25 mm)		3.42 in (86.9 mm)	
Width	17.5 in (445 mm)		17.52 in (445.0 mm)	
Depth	32.8 in (833.0 mm)		29.88 in (759.0 mm)	
Acoustic Noise (operating)	8.7 B	8.0 B	8.0 B	8.0 B
Weight	198.8 lb (90.2 kg)	76.7 lb (34.8 kg)	60.6 lb (27.5 kg)	66.7 lb (30.2 kg)
Maximum Power Usage	3.5 kW (3.6 kVA)	0.8 kW (0.9 kVA)	0.9 kW (0.9 kVA)	0.5 kW (0.5 kVA)
Typical Power Usage ¹	2.5 kW (2.5 kVA)	0.6 kW (0.6 kVA)	0.6 kW (0.6 kVA)	0.4 kW (0.4 kVA)
Cooling at Maximum Usage	11,953 BTU/hour	2,883 BTU/hour	2,907 BTU/hour	1,720 BTU/hour
	12,610 kJ/hour	3,042 kJ/hour	3,067 kJ/hour	1,814 kJ/hour
Cooling at Typical Usage	8,367 BTU/hour	2,018 BTU/hour	2,035 BTU/hour	1,204 BTU/hour
	8,827 kJ/hour	2,129 kJ/hour	2,147 kJ/hour	1,270 kJ/hour
Airflow at Maximum Usage ²	553 CFM	133 CFM	135 CFM	80 CFM
Airflow at Typical Usage ²	387 CFM	93 CFM	94 CFM	56 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	FULL RACK ⁴	HALF RACK
Height	78.74 in (2000 mm)	
Width	23.62 in (600 mm)	
Depth	47.12 in (1197 mm)	
Acoustic Noise (operating)	9.5 B	9.3 B
Environmentals With High Capacity Disks		
Weight	2057.5 lb (933.3 kg)	1251.8 lb (567.8 kg)
Maximum Power Usage	19.7 kW (20.1 kVA)	10.4 kW (10.6 kVA)
Typical Power Usage ¹	13.8 kW (14.1 kVA)	7.3 kW (7.4 kVA)
Cooling at Maximum Usage	67,178 BTU/hour	35,462 BTU/hour
	70,873 kJ/hour	37,413 kJ/hour
Cooling at Typical Usage	47,025 BTU/hour	24,824 BTU/hour
	49,611 kJ/hour	26,189 kJ/hour
Airflow at Maximum Usage ²	3110 CFM	1642 CFM
Airflow at Typical Usage ²	2177 CFM	1149 CFM
Environmentals With Extreme Flash Drives		
Weight	1832.1 lb (831.0 kg)	1203.5 lb (545.9 kg)
Maximum Power Usage	19.8 kW (20.2 kVA)	10.4 kW (10.6 kVA)
Typical Power Usage ¹	13.9 kW (14.1 kVA)	7.3 kW (7.4 kVA)
Cooling at Maximum Usage	67,513 BTU/hour	35,534 BTU/hour
	71,226 kJ/hour	37,488 kJ/hour
Cooling at Typical Usage	47,259 BTU/hour	24,874 BTU/hour
	49,858 kJ/hour	26,242 kJ/hour
Airflow at Maximum Usage ²	3126 CFM	1645 CFM
Airflow at Typical Usage ²	2188 CFM	1152 CFM
<p>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 Airflow must be front-to-back. ⁴ Full rack configuration added as example of elastic configuration. Full = 2x DB Servers, 14x Storage Servers.</p>		

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, KN32, EN61000-3-11, EN61000-3-12
	Immunity:	EN55024, KN35
Certifications ^{2,3}	North America (NRTL), CE (European Union), International CB Scheme, HSE Exemption (India), BSMI (Taiwan), CCC (PRC), EAC (EAEU including Russia), 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)

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 Persistent Memory Data Accelerator
- Unique Persistent Memory Commit Accelerator
- Unique Database Aware PCI Flash
- Unique Exadata Smart Flash Caching
- Unique Exadata Smart Flash Logging
- Unique Smart Write-back Flash Cache
- Unique I/O Prioritization by cluster, workload, DB or user to ensure QOS
- Unique Exafusion Direct-to-Wire Protocol
- Unique Database Intelligent Network Resource Management
- Unique Exachk full-stack validation
- Unique Full-stack security scanning
- Unique Database scoped security
- Unique Cell-to-Cell Rebalance preserving Flash Cache and Storage Index
- Unique Full-Stack Secure Erase
- Unique Instant Data File Creation
- Unique Smart Fusion Block Transfer
- Unique Control of Flash Cache Size per Database
- Unique In-Memory OLTP Acceleration
- Unique Undo-Block Remote RDMA Read
- Unique Support for 4000 Pluggable Databases per Container Database with Multitenant Option

Exadata and Database Software Features – High Availability

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

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
- IPv6 Support for Ethernet Connections
- Capacity on Demand
- Cell software transparent restart
- Flash and disk life cycle management alert
- Automatic Disk Scrub and Repair
- 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 11g Release 2 Enterprise Edition, Oracle Database 12c Enterprise Edition Release 1 and 2, Oracle Database 18c Enterprise Edition, Oracle Database 19c Enterprise Edition, and Oracle Database 21c Enterprise Edition. Oracle Database Options such as Oracle Real Application Clusters, Oracle Partitioning, Oracle Multitenant, Oracle Active Data Guard. See the release specific documentation for feature support.
- **For storage servers:** Oracle Exadata System Software. Licenses are transferable from one system to another, or to a new system.

Oracle Software (included):

- **For database servers:** Oracle Linux 7 Update 9 with the Unbreakable Enterprise Kernel 5. 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 © 2021, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0921