

ORACLE EXADATA DATABASE MACHINE X5-8



The Oracle Exadata Database Machine is engineered to be the highest performing and most available platform for running the Oracle Database. Exadata is a modern architecture featuring scale-out industry-standard database servers, scale-out intelligent storage servers, state-of-the-art PCI flash storage servers and an extremely high speed InfiniBand internal fabric that connects all servers and storage. Unique software algorithms in Exadata implement database intelligence in storage, PCI based flash, and InfiniBand networking to deliver higher performance and capacity at lower costs than other platforms. Exadata runs all types of database workloads including Online Transaction Processing (OLTP), Data Warehousing (DW), In-Memory Analytics and consolidation of mixed workloads. Simple and fast to implement, the Exadata Database Machine is designed to power and protect your most important databases and an ideal foundation for a consolidated database cloud.

The Exadata Database Machine X5-8 is a high-end database machine that has the same extreme performance, storage, and InfiniBand as the X5-2, but uses large-scale 8-socket SMP servers instead of the 2-socket servers in X5-2. Each of the 8-socket servers in the X5-8 has total of 144 processor cores and 2 terabytes to 6 terabytes of DRAM. The X5-8 is especially well suited for high-end OLTP workloads, in-memory workloads, Database as a Service or database consolidation, and multi-rack data warehouses.



Engineered System for Fast Deployment of All Your Databases

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

Because all Exadata Database Machines are identically configured, customers benefit from the experience of thousands of other users that have deployed the Exadata

KEY FEATURES

- Up to 576 CPU cores and up to 24 TB of memory for database processing per rack
- Up to 224 CPU cores dedicated to SQL processing in storage
- From 2 to 4 8-socket database servers
- From 3 to 14 Oracle Exadata Storage Servers
- 40 Gb/second InfiniBand
- Up to 180 TB of Flash Storage per rack
- Hybrid Columnar Compression often delivers 10X-15X compression ratios
- Complete redundancy for high availability
- Fault-tolerant In-Memory Duplication for Oracle Database In-Memory
- Oracle Linux database servers

KEY BENEFITS

- Uncompressed I/O bandwidth of up to 263 GB/second per rack from SQL
- Ability to perform up to 4,140,000 database 8K read I/O operations per second
- Scale by connecting multiple Exadata Database Machine X5-8 racks or Exadata Storage Expansion Racks.
- Up to 18 racks can be connected by simply connecting via InfiniBand cables and using internal switches. Larger configurations can be built with external InfiniBand switches
- Pre-configured system optimized for all database applications

Database Machine for their mission critical applications. Customer machines are also identical to the machines Oracle Support uses for problem identification and resolution, and the machines Oracle Engineering uses for development and testing of the Oracle Database. Hence, Exadata is the most thoroughly tested and tuned platform for running the Oracle Database and is also the most supportable platform. The Oracle Exadata Database Machine runs the standard Oracle Database. Therefore, any application that uses the Oracle Database today can be seamlessly migrated to use the Exadata Database Machine with no changes to the application.

Unlike competing hardware platforms such as IBM® mainframes, Teradata®, or IBM® PureData™ System for Analytics (powered by Netezza), Exadata platform does not “lock in” customers. These competing platforms have extensive proprietary software interfaces that make it extremely complex and expensive to migrate applications to them. Because Exadata is based on the industry standard Oracle database, applications can be easily and quickly migrated on or off the Exadata platform.

Extreme System Scalability and Growth with Elastic Configurations

The Exadata Database Machine uses a scale-out architecture for both database servers and storage servers. The Exadata configuration carefully balances CPU, I/O and network throughput to avoid bottlenecks. As an Exadata Database Machine grows, database CPUs, storage, and networking are added in a balanced fashion ensuring scalability without bottlenecks.

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

Within the Exadata family, Exadata X5-8 is a high-end database machine that is particularly well suited to large OLTP databases that benefit from running on small numbers of large compute nodes, in-memory database workloads requiring very large memory capacity, high-scale consolidation that benefits from a large memory footprint, large multi-rack configurations that benefit from reduced numbers of database nodes, and databases with high I/O requirements that can benefit from very large buffer caches.

Each Exadata X5-8 machine comes with very powerful 8-socket database servers and 2-socket storage servers. Each database server has eight 18-core x86 processors with 2TB of memory (expandable up to 6TB). Each storage server has two 8-core x86 processors for SQL processing. For High Capacity (HC) configurations, each storage server comes with four 1.6TB PCI Flash cards and twelve 8TB 7,200 RPM HC disks. For Extreme Flash (EF) configurations, each storage server comes with eight 1.6TB NVMe PCI Flash drives. The combined hardware enhancements in Exadata X5-8 enable it to perform up to 25% faster than the previous generation Exadata X4-8.

In addition to the inherent scalability and power of Exadata, elastic configurations provide an extremely flexible and efficient mechanism to expand computing power and/or storage capacity of any given Exadata system. With Elastic Configurations, the system can be custom configured to meet any business use case.

The starting configuration of an Exadata X5-8, which is called the Half Rack configuration for Exadata X5-8, consists of 2 database servers and 3 storage servers.

RELATED PRODUCTS

- Exadata Database Machine X5-2
- Exadata Storage Server X5-2
- Exadata Storage Expansion X5-2
- Oracle Database 11g and 12c
- Real Application Clusters
- Database In-Memory
- Partitioning
- Advanced Compression
- Advanced Security
- Active Data Guard
- GoldenGate
- Real Application Testing
- OLAP
- Advanced Analytics
- Business Intelligence
- Enterprise Manager
- Oracle Linux

RELATED SERVICES

The following services are available from Oracle:

- Advanced Customer Services
- Oracle Premier Support for Systems
- Oracle Infrastructure as a Service On-Premise (IaaS)
- Oracle Platinum Services
- Oracle PlatinumPlus Services
- Consulting Services
- Oracle University courses

This can be further elastically expanded by adding more database or storage servers as requirements grow. The Exadata X5-8 can have between 2 to 4 database servers and 3 to 14 storage servers. A full rack Exadata Database Machine X5-8 consists of 2 database servers and 14 storage servers.

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

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



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

When even larger storage capacity is required the Oracle Exadata Storage Expansion Rack is available. The Exadata Storage Expansion Rack enables you to grow the Exadata storage capacity and I/O bandwidth of any Exadata Database Machine. It is designed for database deployments that require very large amounts of data including: historical or archive data, backups, documents, images, XML and, LOBs. The expansion rack is extremely simple to configure as there are no LUNs or mount points to set up. It connects to the Exadata Database Machine using the integrated InfiniBand fabric. Storage is configured and added to a database online with a few simple commands.

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

Extreme Performance by Offloading Data Intensive Processing

As data volumes grow exponentially, conventional storage arrays struggle to quickly transfer data from disk and flash to database servers at a rate that keeps the CPUs busy. Modern many-core servers can consume data at tens to hundreds of gigabytes a second. This is far faster than conventional architectures that use storage arrays can deliver data through their storage heads and the storage network.

The scale-out architecture of the Exadata Database Machine inherently eliminates performance and scalability bottlenecks that are common in storage arrays. Exadata also includes a unique technology that offloads data intensive SQL operations into the Oracle Exadata Storage Servers. By pushing SQL processing to the Exadata Storage Servers, data filtering and processing occurs immediately and in parallel across all storage servers as data is read from disk and flash. Only the rows and columns that are directly relevant to a query are sent to the database servers.

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

Exadata X5-8 can have total of 224 processor cores in the storage servers that can be used to offload the database servers. The CPUs in Exadata Storage Servers do not replace database CPUs. Instead they accelerate data intensive workloads similar to how graphics cards accelerate image intensive workloads.

Optimizing Storage Use and I/O Through Compression

Exadata Storage Servers implement an advanced compression capability called Hybrid Columnar Compression (HCC) that provides dramatic reductions in storage for large databases. HCC provides tremendous cost-savings and performance improvements due to reduced I/O, especially for analytic workloads. The amount of storage savings is data dependent and often ranges from 5x to 20x. Typical storage savings is an industry leading 10x. On conventional systems, enabling high data compression has the drawback of reducing performance. Because the Exadata Database Machine is able to offload decompression overhead into large numbers of processors in Exadata storage, most analytics workloads run faster using HCC than they do without it. HCC delivers the compression and analytic performance benefits of column storage while avoiding the dramatic slowdown that pure columnar stores experience for drilldown operations (single row access).

Two modes of HCC are available. Query optimized compression mode is suitable for read intensive workloads such as Data Warehouses and provides large storage savings and enhanced analytic performance. Archive compression mode provides the highest degree of compression and is targeted at seldom accessed data that is kept online.

On OLTP systems, Hybrid Columnar Compression can be used to compress older, less active data while newer, more active and update intensive data can be compressed using Oracle Advanced Row Compression. Oracle Database 12c 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.

Record breaking I/O Performance from Exadata Smart Flash Cache

New generation of Exadata X5-2 storage servers introduces Extreme Flash Storage Servers. Each Extreme Flash storage server contains eight 1.6 TB state-of-the-art PCI Flash drives. PCI flash delivers ultra-high performance by placing flash memory directly on the high speed PCI bus rather than behind slow disk controllers and directors. Exadata flash uses the latest NVMe (Non-Volatile Memory Express) flash protocol to achieve extremely low I/O overhead. The Extreme Flash Storage Server replaces the previous High Performance disk configuration of the Exadata Storage Server and is backward compatible with all supported Exadata systems.



Fig. 2: Flash Accelerator F160 PCIe Card

Flash performance is often limited and bottlenecked by traditional storage architectures. In contrast, Exadata uses a combination of scale-out storage, InfiniBand networking, database offload, and PCI flash to deliver extremely high performance rates from flash. A single full rack Exadata Database Machine X5-8, with 2 database servers and 14 Extreme Flash storage servers can achieve up to 263 GB per second of data scan bandwidth, and up to 4.1 Million random 8K read and write I/O operations per second (IOPS) when running database workloads. This performance is orders of magnitude faster than traditional storage array architectures, and is also much faster than current all-flash storage arrays. It is important to note that these are real-world end-to-end performance figures measured running SQL workloads with realistic I/O sizes inside a single rack Exadata system. They are not component-level measurements based on low-level I/O tools.

Tiered Disk and Flash Delivers Cost of Disk with Performance of Flash

The second Exadata storage option in X5-8 is the Exadata X5-2 High Capacity Storage Server. This server includes twelve 8 TB SAS disk drives (96 TB total) and four Flash Accelerator F160 NVMe PCIe cards with a total raw capacity of 6.4TB of flash memory. Exadata flash in a High Capacity Storage Server can be used directly as flash disks, but is almost always configured as a flash cache (Exadata Smart Flash Cache) in front of disk since caching provides flash level performance for much more data than fits directly into flash.

The Exadata Smart Flash Cache automatically caches frequently accessed data while keeping infrequently accessed data on disk drives. This provides the performance of flash with the capacity and low cost of disk. The Exadata Smart Flash Cache understands database workloads and knows when to avoid caching data that the database will rarely access or is too big to fit in the cache. For example, Exadata doesn't cache I/Os caused by backups, table scans, or temporary results that will be quickly deleted. In addition to automatic caching, administrators can optionally provide SQL directives to ensure that specific tables, indexes, or partitions are always retained in the flash cache. Tables can be retained in flash cache without the need to move the table to different tablespaces, files or LUNs as is often required with traditional storage. A single full rack Exadata Database Machine X5-8, with 2 database servers and 14 High Capacity storage servers can achieve up to 140 GB per second of data scan bandwidth, and up to 4.1 Million random 8K read I/O operations per second (IOPS) when running database workloads.

Exadata's Smart Flash Cache is designed to deliver flash-level IO rates, and response times for data that is many times larger than the physical flash capacity in the machine by moving active data into flash, while leaving cold data on disk. It is common for hit rates in the Exadata Smart Flash Cache to be over 90%, or even 98% in real-world database workloads even though flash capacity is more than 15 times smaller than disk capacity. Such high flash cache hit rates mean that Exadata Smart Flash Cache provides an effective flash capacity that is many times larger than the physical flash cache. For example, a full rack Exadata Database Machine X5-8 with 2 database servers and 14 High Capacity Storage Servers often has an effective flash capacity equal to the full disk capacity of 1344 TB.

The Exadata Smart Flash cache also caches database block writes. Write caching eliminates disk bottlenecks in large scale OLTP and batch workloads. The flash write capacity of a single full rack Exadata Database Machine X5-8 with 2 database servers and 14 High Capacity Storage Servers exceeds 2.68 Million 8K write I/Os per second. The Exadata write cache is transparent, persistent, and fully redundant. The I/O performance of the Exadata Smart Flash Cache is comparable to dozens of enterprise disk arrays with thousands of disk drives.

To further accelerate OLTP workloads, the Exadata Smart Flash Cache also implements a special algorithm to reduce the latency of log write I/Os called Exadata Smart Flash Logging. The time to commit user transactions or perform critical updates is very sensitive to the latency of log writes. Smart Flash Logging takes advantage of the flash memory in Exadata storage combined with the high speed RAM memory in the Exadata disk controllers to greatly reduce the latency of log writes and avoid the latency

spikes that frequently occur in other flash solutions. The Exadata Smart Flash Logging algorithms are unique to Exadata.

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

Exadata uses only enterprise grade flash that is designed to have high endurance. Exadata is designed for mission critical workloads and does not use consumer grade flash that can potentially experience performance degradations or fail unexpectedly after a few years of usage. The enterprise grade flash chips used in Exadata X5-8 have an expected endurance of 8 years or more for typical database workloads.

The automatic data tiering between RAM, flash and disk implemented in Exadata provides tremendous advantages over other flash-based solutions. Many storage vendors have recognized that the architecture of their traditional storage arrays inherently bottleneck the performance of flash and therefore have developed new flash-only arrays. These flash-only arrays deliver higher performance than traditional arrays but give up the cost advantages of smart tiering of data between disk and flash. Therefore the overall size of data that benefits from flash is limited to the size of expensive flash. Third party flash arrays will also not benefit from Exadata Hybrid Columnar Compression. Data deduplication provided by some flash arrays is very effective for VDI environments but is ineffective for databases.

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

Fault Tolerant and Fastest Database In-Memory Machine

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

Elastic configurations enable users to configure systems with very large system memory. Each 8-socket SMP server on X5-8 can directly access terabytes of memory using its ultra-fast memory interconnect. This often enables in-memory queries to avoid the communication overhead of distributing processing across multiple nodes of a cluster. When in-memory queries are distributed across nodes, the 40Gb per second InfiniBand network provides extremely high throughput and very low latencies.

Exadata implements Oracle Database In-Memory's Fault Tolerant capability, a feature that is unique to Oracle Engineered Systems. On a generic cluster configuration, when a server node fails, the in-memory data on that node is lost, and it takes many minutes to repopulate the in-memory data on a surviving node. During this time analytic queries will run dramatically slower. This means generic configurations will fail to meet business SLAs. However, when deployed on Exadata, the Fault-Tolerant feature of Oracle Database In-Memory can eliminate this slowdown by duplicating any subset of the in-memory data across nodes. If a node fails, queries can transparently use the duplicate copy of data on a surviving node and processing can continue without interruption.

Enterprise-Class Security with Extreme Performance

Exadata Database Machine is the world's most secure database machine. Building on the high security capabilities in the Oracle Database, Exadata 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 near zero overhead at hundreds of gigabytes of user data per second.

The Exadata system is designed and delivered as an integrated whole, and not a collection of components. In traditional database deployments, the customer takes on all the integration tasks for the system – including the task of ensuring the security of each individual software and hardware component, and ensuring that security is maintained across the full product stack. Oracle delivers full stack security in the Exadata Database Machine.

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

Mission Critical High Availability

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

Because of its industry leading availability, the Exadata Database Machine has been

deployed by leading companies for their most critical applications including interbank fund transfers, online securities trading, real-time call tracking, and web-based retailing. Mission Critical availability is not restricted to OLTP workloads; it also applies to warehousing and analytics workloads.

Comprehensive System Management

Oracle Enterprise Manager 12c uses a holistic approach to manage the Exadata Database Machine and provides comprehensive lifecycle management from monitoring to management and ongoing maintenance for the entire system. It provides a unified view of all the hardware and software components such as database servers, Exadata storage, and InfiniBand switches and allows monitoring the operations running on them and their resource utilization. DBAs can drill down from database monitoring screens to the storage layer of the Exadata Database Machine to quickly determine the root cause of application level performance bottlenecks. Lights-out monitoring within Enterprise Manager is optimized for the Exadata Database Machine with predefined metrics and thresholds so that administrators receive timely notifications when issues arise. In addition, hardware incidents are automatically detected and service requests logged to reduce problem resolution time.

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 (DBaaS) Cloud. Multi-database environments inherently have diverse, complex, and unpredictable workloads mixing OLTP, Analytics, and Batch operations with sequential and random access patterns. Exadata's ability to run any type or mix of database workloads with industry leading scalability and performance makes it an ideal platform for multi-database workloads.

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

Exadata implements unique database and I/O resource management. Fine-grained priorities specified for operations at the database level are automatically communicated to Exadata Storage Servers and applied to each I/O operation to ensure that prioritization of database operations applies to both CPU operations and I/O operations.

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

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 the Multitenant Database Option to provide an extremely simple interface for creating new PDB snapshots.

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

All Exadata specific features such as Smart Scan, resource management and Smart Flash Cache work seamlessly on database instances created via Exadata snapshots hence providing an exact test and development environment while using a fraction of valuable storage resources.

Oracle Platinum 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, fault monitoring, software maintenance, and patching are performed remotely by Oracle engineers. Platinum Services provides a higher level of support for all software and hardware within an Engineered System including the Oracle Database. Platinum Services is provided free of 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 single Database Machine Administration team. Database Machine Administrators have full control of all resources in the Exadata Database Machine including storage resources. New database deployments and configuration changes can be implemented by the Database Machine Administrators without coordination across different component management teams that are often overloaded and have differing priorities. Database Machine Administrators can focus on application and business specific enhancements rather than coordinating across component teams, or tuning and triaging of low level configuration issues.

Dramatically Lower Costs

Because of the extreme performance, high storage capacity, and unique compression capabilities delivered by the Exadata Database Machine, workloads that would require very large traditional hardware systems can be run on much smaller Exadata systems. The hardware needed for an application deployed on an Exadata system is often reduced 10X compared to a traditional system.

Exadata provides a huge RAM, flash and disk footprint for large data sets. Raw storage on an Exadata X5-8 can exceed 1344 TB and Hybrid Columnar Compression often compresses data 10X. By intelligently moving active data across storage and memory tiers, Exadata simultaneously delivers the highest performance and the lowest cost.

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

The uniformity of Exadata Database Machine configurations results in large cost savings. Exadata standardizes not just technologies, but also integration, testing, 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 greatly reduced or eliminated. Because all Exadata users run a configuration that is identical to thousands of other users, and is identical to Oracle's internal configurations, it is far less likely that issues will be encountered, and issue resolution is quicker and simpler reducing both operations cost and downtime cost.

Capacity-on-Demand Software Licensing

An X5-8 database server has a substantial amount of compute capacity with eight 18-core processors (144 cores in total).

The Capacity-on-Demand feature allows up to 60% of the cores per server to be turned off during the hardware installation, leaving at least 56 cores enabled for the server. As your workload grows and more cores are needed, Capacity-on-Demand can be used to re-enable cores and license software 8 cores at a time. This pay-as-you-grow approach to software licensing is yet another way in which Exadata helps to align costs with business growth.

Exadata Business Benefits

Beyond the operational benefits of extreme performance, availability, and security at low cost, Exadata also directly benefits the business.

Exadata accelerates time to market for new business applications since the time needed for system configuration, tuning, and testing is largely eliminated. Deployment times are reduced from months to days, and the risk of unexpected system level issues after go-live is greatly reduced. When a new application is deployed, it is common for unanticipated application usage patterns to create performance issues. Exadata's huge I/O, network, and compute throughput can absorb spikes created by unanticipated

workloads without slowing response times of mission critical workloads. Overall Exadata speeds application deployment and reduces risk allowing businesses to innovate faster.

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

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

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

Conclusion

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

EXADATA DATABASE MACHINE X5-8 KEY CAPACITY AND PERFORMANCE METRICS

Metric	Full Rack HC ¹	Full Rack EF ¹	Half Rack HC	Half Rack EF
Flash Metrics				
Maximum SQL Flash Bandwidth ²	140 GB/sec	263 GB/sec	30 GB/sec	56 GB/sec
Maximum SQL Flash Read IOPS ³	4,140,000	4,140,000	1,200,000	1,200,000
Maximum SQL Flash Write IOPS ⁴	2,688,000	4,140,000	576,000	1,131,000
Flash Data Capacity (raw) ⁵	89.6 TB	179.2 TB	19.2 TB	38.4 TB
Effective Flash Cache Capacity ⁶	1344 TB	N/A	288 TB	N/A
Disk Metrics				
Maximum SQL Disk bandwidth ²	25 GB/sec	N/A	5.4 GB/s	N/A
Maximum SQL Disk IOPS ³	36,000	N/A	7,800	N/A
Disk Data Capacity (raw) ⁵	1344 TB	N/A	288 TB	N/A
Combined Metrics				
Data Capacity (usable) ⁷	508TB	66 TB	109 TB	14 TB
Maximum Data Load Rate ⁸	21 TB/hr	21 TB/hr	5 TB/hr	6 TB/hr

Actual system performance varies by application.

¹ EF = Extreme Flash; HC = High Capacity

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

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

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

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

⁶ Effective Flash Capacity is larger than physical flash capacity and takes into account high flash hit ratios due to Exadata's intelligent flash caching algorithms, and the size of underlying disk storage. It is the size of the data files that can often be stored in Exadata and be accessed at the speed of flash memory.

⁷ Usable capacity is measured using normal powers of 2 space terminology with 1 TB = 1024 * 1024 * 1024 * 1024 bytes. It is the actual space available to create a database after taking into account space needed for ASM redundancy, recovering from a drive failure, DBFS disk group, and OS images and binaries.

⁸ Load rates are typically limited by database server CPU, not IO. Rates vary based on load method, indexes, data types, compression, and partitioning.

EXADATA DATABASE MACHINE X5-8 HARDWARE PER RACK

Components

2 to 4 Database Servers, each with:

- 8 x Eighteen-Core Xeon® E7-8895 v3 Processors (2.80 GHz)
- 2 TB Memory (expandable up to 6 TB)
- Disk Controller HBA with 1GB Cache
- 8 x 600 GB 10,000 RPM Disks
- 8 x InfiniBand QDR (40Gb/s) Ports
- 8 x 10 Gb Ethernet Ports based on the Intel 82599 10GbE Controller
- 10 x 1 Gb Ethernet Ports
- 1 x ILOM Ethernet Port
- 4 x Redundant Hot-Swappable Power Supplies

3 to 14 Exadata Storage Servers X5 8 HC or EF each with:

HC	EF
<ul style="list-style-type: none"> • 2 x Eight-core processor for SQL processing • 4 PCI Flash cards each with 1.6 TB (raw) Exadata Smart Flash Cache • 12x 8 TB 7,200 RPM High Capacity disks 	<ul style="list-style-type: none"> • 2 x Eight-core processor for SQL processing • 8 x 1.6 TB NVMe PCI Flash Drives

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

Additional Hardware Components:

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

Included Spare Parts Kit Contains:

- 1 x 1.6 TB NVMe PCI Flash Card and 1 x 4 TB High Capacity disk, or
- 1 x 1.6 TB NVMe PCI Flash drive

EXADATA DATABASE MACHINE X5-8 EXPANSION AND UPGRADES

Expansion & Upgrades

Expansion: Connect any combination of up to 18 Exadata Database Machine racks or Exadata Storage Expansion Racks via the InfiniBand fabric. Larger configurations can be built with external InfiniBand switches. Connected racks can be any combination of X2, X3, or X5 generation hardware.

Upgrade Support Services:

- *Hardware Installation and Software configuration*

EXADATA DATABASE MACHINE X5-8 SUPPORT SERVICES

Components

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

EXADATA DATABASE MACHINE X5-8 ENVIRONMENTAL SPECIFICATIONS

Metric	Full Rack	Half Rack
Height	<ul style="list-style-type: none"> • 78.66" - 1998 mm • 23.62" – 600 mm • 47.24" – 1200 mm 	
Width		
Depth		
Acoustic noise (operating)	8.6 B	8.5 B

Environmentals With High Capacity Disks

Weight	1980.5 lbs (898.3 kg)	1193.1 lbs (541.2 kg)
Maximum power usage	16.2 kW (16.6 kVA)	9.5 kW (9.7 kVA)
Typical power usage ¹	11.4 kW (11.6 kVA)	6.6 kW (6.8 kVA)
Cooling at maximum usage	55,410 BTU/hour	32,323 BTU/hour
	58,457 kJ/hour	34,101 kJ/hour
Cooling at typical usage	38,787 BTU/hour	22,626 BTU/hour
	40,920 kJ/hour	23,871 kJ/hour
Airflow at maximum usage ²	2565 CFM	1496 CFM
Airflow at typical usage ²	1796 CFM	1048 CFM

Environmentals With Extreme Flash Drives

Weight	1826.5 lbs (828.5 kg)	1160.1 lbs (526.2 kg)
Maximum power usage	16.4 kW (16.7 kVA)	9.5 kW (9.7 kVA)
Typical power usage ¹	11.5 kW (11.7 kVA)	6.7 kW (6.8 kVA)
Cooling at maximum usage	55,935 BTU/hour	32,436 BTU/hour
	59,012 kJ/hour	34,220 kJ/hour
Cooling at typical usage	39,155 BTU/hour	22,705 BTU/hour
	41,308 kJ/hour	23,954 kJ/hour
Airflow at maximum usage ²	2590 CFM	1502 CFM
Airflow at typical usage ²	1813 CFM	1051 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

Regulations ³

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

Certifications ³

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

European Union Directives ³

- 2006/95/EC Low Voltage Directive, 2004/108/EC EMC Directive, 2011/65/EU RoHS Directive, 2012/19/EU WEEE Directive

¹ Typical power usage varies by application load.

² Airflow must be front-to-back

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

KEY FEATURES & FUNCTIONALITY

Oracle Database Software (sold separately)

For Database Servers	Oracle Database 11g Release 2 Enterprise Edition and Oracle Database 12c Enterprise Edition Oracle Real Application Clusters, Oracle Partitioning, Oracle Multitenant and other Oracle Database options are available See the release specific documentation for feature support.
For Storage Servers	Oracle Exadata Storage Server Software Licenses are transferable from one system to another, or to a new system.

Oracle Software (included)

For Database Servers	Oracle Linux 6 Update 6 with the Unbreakable Enterprise Kernel 2 Zero-loss Zero-copy Datagram Protocol (ZDP) InfiniBand protocol used to communicate between the Exadata Storage Servers and the Oracle Database which is based on the Reliable Datagram Sockets (RDS) OpenFabrics Enterprise Distribution (OFED)
----------------------	--

Exadata Storage Server Software Features (partial list)

<ul style="list-style-type: none"> • <i>Smart Scan Technology</i> • <i>Smart Flash Cache</i> • <i>Smart Flash Logging</i> • <i>Columnar Flash Cache</i> • <i>Exadata Snapshots</i> • <i>Exafusion Direct to Wire OLTP Protocol</i> 	<ul style="list-style-type: none"> • <i>IO and Network Resource Management</i> • <i>Storage Index</i> • <i>Hybrid Columnar Compression</i> • <i>Smart Scans of Data Mining model scoring</i> • <i>I/O Analysis in AWR reports</i>
--	--

High Availability Features

<ul style="list-style-type: none"> • <i>Redundant power supplies for all servers</i> • <i>Redundant InfiniBand switches</i> • <i>Redundant Power Distribution Units</i> • <i>Oracle Automatic Storage Management: All database files mirrored; disk failures do not interrupt query processing</i> • <i>Oracle Real Application Clusters: database server failures are tolerated</i> 	<ul style="list-style-type: none"> • <i>Oracle Exadata Storage Server Software: storage server failures are tolerated</i> • <i>Backup is performed using Oracle Recovery Manager</i> • <i>Point in time restores are performed using Oracle Flashback Technologies</i> • <i>Oracle Data Guard for protection against disasters</i> • <i>Near instant server death detection</i> • <i>I/O latency capping</i> • <i>In-Memory Fault Tolerance</i>
---	--





Manageability Features

<ul style="list-style-type: none"> • <i>Oracle Embedded Integrated Lights Out Manager (ILOM)</i> • <i>Oracle Enterprise Manager 12c</i> 	
---	--

**CONTACT US**

For more information about [insert product name], visit oracle.com or call +1.800.ORACLE1 to speak to an Oracle representative.

CONNECT WITH US

-  blogs.oracle.com/oracle
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Hardware and Software, Engineered to Work Together

Copyright © 2015, 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. 0416