

Oracle Exadata Storage Expansion X11M


The Oracle Exadata Database Machine (Exadata) is engineered to deliver higher performance and capacity at lower costs than other platforms for all types of modern database workloads, including Online Transaction Processing (OLTP), Analytics and Data Warehousing (DW), In-Memory Analytics, Artificial Intelligence (AI), Internet of Things (IoT), financial services, gaming, and compliance data management, as well as efficient consolidation of mixed database workloads. The Oracle Exadata Storage Expansion Rack is engineered to be the simplest and fastest way to add large amounts of storage capacity to an Exadata Database Machine. A natural extension of the Exadata Database Machine, the Exadata Storage Expansion Rack can be used to satisfy the requirements of the largest mission-critical databases.

Engineered For Fast and Reliable Deployment

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

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. The database servers, storage servers, and network are pre-configured, pre-tuned, and pre-tested by Oracle experts. Extensive end-to-end testing and validation ensure all components, including database software, operating system, hypervisor, drivers, and firmware, work seamlessly together and that there are no performance bottlenecks or single points of failure. The Exadata Storage Expansion Rack is built using the same secure, highly available, and scale-out architecture as the Exadata Database Machine, enabling additional storage to be added when needed, rather than having to size for peak capacity requirements.



Oracle Exadata Storage
Expansion X11M

Key Features

- Grow the storage capacity of RoCE based Oracle Exadata Database Machines
- 2x 100 Gb/sec active-active RoCE Network
- Includes from 4 to 19 Oracle Exadata Storage Servers
- Mirrored usable disk storage capacity of up to 2,075TB per rack before compression
- Mirrored usable flash storage capacity of up to 966TB flash capacity per rack before compression
- Up to 1,216 CPU cores dedicated to SQL processing in storage
- Up to 23.75 TB of Exadata RDMA Memory Acceleration per rack
- Complete redundancy for high availability

Extreme Performance and Capacity

The Exadata Storage Expansion X11M rack enables you to grow the Exadata storage capacity and bandwidth of any RoCE-based Exadata Database Machine. RoCE (RDMA over Converged Ethernet) is the internal fabric that connects servers to other servers and racks to other racks in an Exadata system configuration. Exadata Storage Expansion is designed for database deployments that require very large amounts of data, including historical or archive data, backups, documents, images, XML, LOBs, etc. The expansion rack is straightforward to configure as there are no LUNs or mount points to set up. Storage is configured and added to a database, online with simple commands.

Extreme System Scalability and Growth with Elastic Configurations

The Exadata Storage Expansion X11M offers you more flexibility than ever to grow. Exadata Storage Expansion X11M can be configured and purchased as small as a Quarter Rack with four High Capacity or Extreme Flash Storage Servers, and additional storage servers can be added one or more at a time up to a maximum of 19. With the flexibility of adding between 4 and 19 storage servers, there is a configuration that fits any application. In addition to upgrading from a small to a large Exadata Storage Expansion X11M, Oracle continues to use a building-block approach to connect the Exadata Storage Expansion X11M to the Exadata Database Machine X11M using the integrated RoCE network. Exadata Storage Expansion X11M can be coupled to Exadata Database Machine X11M Elastic Rack systems in almost any combination. Up to 14 Exadata Database Machine X11M racks and Exadata Storage Expansion X11M racks can be easily connected via RoCE cables and internal switches. A single rack of Exadata Database Machine X11M with nine database servers and nine High Capacity storage servers¹, along with 13 fully populated Exadata Storage Expansion X11M racks², has a raw disk capacity of 67.5 Petabytes (67,584 TB) and 16,384 CPU cores dedicated to SQL processing. Larger configurations can be built with additional RoCE switches.

As new Exadata Storage Expansion X11M racks are connected to a RoCE-based Exadata Database Machine, the storage capacity and performance of the system grow. The system can be run as a single system or logically partitioned to consolidate multiple databases and clusters. Scaling out is easy with Exadata Database Machine and Exadata Storage Expansion Racks. Automatic Storage Management (ASM) and Exadata Exascale dynamically and automatically balance the data across Exadata storage servers online, evenly spreading the I/O load across the racks, fully utilizing all the hardware, and easily integrating the expansion rack into the configuration. The I/O Resource Manager can also be used to apportion I/O bandwidth to different databases and users of the system to deliver on business service level targets.

- **Key Benefits**
- Pre-configured, pre-tested system optimized for all database applications
- Uncompressed I/O bandwidth of up to 9,500 GB/second per rack from SQL
- Engineered scale-out storage architecture pre-configured to easily expand system capacity and performance, online
- Simple upgrade to meet the needs of any size application
- Over 15 Petabytes of user data can be stored in a rack using the included Hybrid Columnar Compression
- Extreme Scalability by connecting multiple Exadata Database Machine X11M racks or Exadata X11M Storage Expansion Racks. Up to 14 racks can be connected by simply adding RoCE cables and internal switches. Larger configurations can be built with external RoCE switches

¹ Exadata elastic configuration with 9x Exadata X11M Database Servers and 9x Exadata X11M High Capacity Storage Servers

² Exadata Storage Expansion Rack with 19x Exadata X11M High Capacity Storage Servers

Extreme Performance by Offloading Data Intensive Processing

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

The scale-out architecture of the Exadata Database Machine X11M provides high performance and scalability and 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, flash, and memory. Only the rows and columns directly relevant to a query are sent to the database servers. For example, suppose a query is executed to identify the customers who placed sales orders over \$1000 in the month of March. In that case, 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 return 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. Each storage server has CPUs used to offload database processing. A single Exadata Storage Expansion X11M rack with 19 Storage Servers has up to 1,216 processor cores that offload SQL processing from the database servers. The CPUs in the storage servers do not replace database CPUs. They accelerate database-intensive workloads similar to how graphics cards accelerate image-intensive workloads.

Optimize Storage Consumption and I/O Using Compression

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

With Hybrid Columnar Compression, Exadata enables the highest levels of data compression possible with Oracle databases and provides substantial cost-savings and performance improvements due to reduced I/O, especially for analytic workloads. Storage savings is data-dependent and often range from 5x to 20x. Average storage savings is an industry-leading 10x. Exadata 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.

Related Products

- Exadata Cloud Infrastructure
- Exadata Cloud@Customer
- Oracle Exadata Storage Expansion Rack X11M
- Oracle Exadata Storage Server X11M
- Oracle Exadata Database Server X11M
- Oracle Database Enterprise Edition 19c, 21c, and 23ai
- Exadata Database Service
- Oracle Autonomous Database
- Real Application Clusters
- Partitioning
- Multitenant
- Database In-Memory
- Advanced Compression
- Advanced Security
- Active Data Guard
- GoldenGate
- Real Application Testing
- OLAP
- Enterprise Manager
- Oracle Linux
- Oracle Linux Virtualization

Related Services

The following services are available from Oracle:

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

Two modes of Hybrid Columnar Compression are available. **Warehouse compression** mode suits read-intensive workloads and provides large storage savings and enhanced analytic performance. **Archive compression** mode offers the highest degree of compression and targets data that is seldom accessed but must remain online.

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

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

Groundbreaking RDMA-based Network Fabric

The Exadata X11M release uses the same ultra-fast cloud-scale networking fabric that was introduced in Exadata X8M, RDMA over Converged Ethernet (RoCE). RDMA (Remote Direct Memory Access) allows one computer to directly access data from another without Operating System or CPU involvement for high bandwidth and low latency. The network card directly reads/writes memory with no extra copying or buffering and with very low latency. RDMA is an integral part of the Exadata high-performance architecture and has been tuned and enhanced with each new generation of Exadata, underpinning several Exadata-only technologies such as Exafusion Direct-to-Wire Protocol and Smart Fusion Block Transfer.

The Exadata X11M release implements a dual port PCIe Gen 5 network interface card capable of 2x 100 Gb/sec active-active RoCE network for a total throughput of 200 Gb/sec. By leveraging the RoCE network, **Oracle Database on Exadata can perform read I/O directly from memory in the shared storage servers.**

Shared Exadata RDMA Memory Acceleration

Exadata RDMA Memory (XRMEM) in the Exadata Storage Servers is leveraged as a shared read accelerator. The XRMEM Data Accelerator is a memory cache tier in front of Flash Cache, enabling orders of magnitude lower latency accessing remotely stored data. By utilizing RDMA to access memory remotely, **XRMEM Data Accelerator bypasses the network and I/O stack, eliminating expensive CPU interrupts and context switches, and reducing latency by more than 14x**, from 200 microseconds to as low as 14 microseconds. Smart Exadata System Software also ensures data is mirrored across storage servers, which provides additional fault tolerance. Exadata's unique end-to-end integration between Oracle Database and Exadata Storage Servers automatically caches the hottest data blocks efficiently between the buffer cache in database servers and XRMEM and Flash Cache in storage servers. XRMEM is a shared storage tier across all of the storage nodes, which means the aggregate performance of this cache can be dynamically used by any database instance on any database server. This is a significant advantage over general-purpose storage architectures, which preclude sharing storage resources across database instances.

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

Record-breaking Analytics I/O Performance

Each X11M High Capacity (HC) and Extreme Flash (EF) Storage Server includes four 6.8 TB performance-optimized flash drives, offering 27.2 TB of Exadata Smart Flash Cache. Exadata delivers ultra-high performance by placing all the flash devices directly on the high-speed PCIe interface rather than behind slow disk controllers. Exadata HC and

EF Storage Servers include 1.25 TB of DDR5 Exadata RDMA Memory as a data acceleration tier in front of Flash Cache.

For data warehouse environments that require the highest performance, Exadata X11M HC and EF storage servers can achieve **up to 9.5 TB/s scan throughput**³ from a single Exadata Storage Expansion Rack.

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

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

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

Flash in the HC Storage Server can be used directly as flash disks, but is almost always configured as a flash cache (Exadata Smart Flash Cache) in front of disk storage and behind the Exadata RDMA Memory. Exadata Smart Flash Cache is used with the XRMEM Data Accelerator to automatically cache frequently accessed data while keeping infrequently accessed data on disk, delivering the high I/O rates and fast response times of flash with the large capacity and low cost of disk. Exadata uniquely understands database workloads and knows when to avoid caching data that negatively affects overall performance. For example, if large write I/Os caused by backups or large table scans are likely to disrupt higher-priority OLTP or scan operations, those large I/Os will bypass the flash cache and go straight to disk. Otherwise, Exadata System Software will utilize additional spare flash capacity and I/O bandwidth to optimize performance by caching these I/Os. Administrators can also manually provide SQL directives to ensure that specific tables, indexes, or partitions are preferentially retained in the flash cache.

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

Exadata Smart Flash Cache also caches database block writes using Exadata Write-Back Flash Cache technology. Write caching eliminates disk bottlenecks in large-scale OLTP and batch workloads. The flash write capacity of a single Exadata Database Machine X11M rack reaches 13 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 memory, flash, and disk in Exadata provides tremendous advantages over other flash-based solutions. Many storage vendors have developed flash-only arrays to achieve higher performance than traditional disk-based arrays. However, they cannot match the cost advantages of Exadata's smart data tiering between disk and flash. Traditional flash arrays lack Exadata's unique database-aware storage optimization. In addition, generic data deduplication provided by some flash arrays may be effective for workloads such as Virtual Desktop Infrastructure environments but are not for databases. In addition to utilizing its integrated and optimized hardware architecture, Exadata delivers superior performance by offloading data-intensive processing to unique algorithms in storage which have been specifically optimized for Oracle Database.

Extreme Flash Storage Server: Lowest Latency, All Flash

Exadata **Extreme Flash (EF) Storage Server** is the foundation of a database-optimized all-flash Exadata Database Machine. Each EF Storage Server contains four capacity-optimized 30.72 TB flash drives for an aggregate, raw storage capacity of 122.88 TB and increasing the usable storage capacity by 2.4x⁴. Exadata EF storage servers deliver

³ Exadata X11M Storage Expansion Rack with 19x Exadata X11M High Capacity or 19x Exadata X11M Extreme Flash Storage Servers

⁴ Compared to X9M Extreme Flash Storage Servers

⁵ Oracle Exadata Storage Expansion X11M / Version 1.1

extreme high performance, flash capacity, and low latency for database workloads with the highest performance demands. In addition, X11M EF storage includes four 6.8 TB performance-optimized flash drives for a total raw capacity of 27.2 TB, increasing the **size of the Smart Flash Cache by 11.5x⁵** and is used to satisfy read and write requests. 1.25 TB of DDR5 Exadata RDMA Memory (XRMEM) is available in front of flash to boost performance further.

High Capacity-Z Storage Server: High Performance Storage Accelerates Smaller Database Workloads

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

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

Extreme Backup and Recovery Performance

A key strength of the Exadata Storage Expansion rack is when used as a destination for Exadata Database Machine backups. A full database backup can be created at up to 48 TB/hour when backing up uncompressed data that is being written to mirrored disk in an Exadata Storage Expansion rack. Taking incremental backups increases the effective backup rate to **hundreds of terabytes per hour**. Combining incremental backups and Hybrid Columnar Compression increases the effective backup rate to **petabytes per hour**.

A disk backup on an Exadata Storage Expansion rack is used directly without loss of performance and without having to do a restore. This unique backup capability is only available when backing up to an Exadata Storage Expansion rack. It is the fastest and simplest way to back up and recover your Oracle Exadata Database Machine.

Mission Critical High Availability

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

Exadata in an MAA configuration is recognized by the analyst firm IDC as a system that **delivers at least 5-nines availability** (99.999%) and is categorized in the IDC AL4 fault-tolerant market 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 Database Machines implement a unique RDMA-based **sub-second node failure detection**, leading to the virtual elimination of application brownout conditions.

⁵ Exadata X7 – X9M Extreme Flash Storage Servers allocated 2.32 TB to Flash Cache of the total 51.2 TB raw flash capacity per server.

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

Disk and flash devices occasionally exhibit high latency I/O operations due to internal recovery of failed sectors, internal firmware reboots, or wear leveling. These extended I/O operations can cause stalls in mission-critical OLTP databases. With Exadata's **unique I/O Latency Capping**, Oracle Exadata System software automatically redirects read I/O operations to an ASM-mirrored copy of the data when the latency of a read I/O is much longer than expected. Similarly, I/O Latency Capping automatically redirects high latency write I/O operations to a healthy flash device, eliminating outliers during write operations. Exadata System Software uses machine learning techniques to predict components susceptible to failure and takes proactive action to take such components out of service gracefully. If disks fail, ASM performs a rebalance operation for the data resident on the disk, while applications continue to access the database with no interruption. Exadata allows hot swapping of disks, fans, power supplies, and PCIe flash cards to avoid downtime. Exadata System Software takes rebalance further by preserving the flash cache population and storage indexes when moving data between storage servers to maintain consistent application performance. On rare occasions, when there are outliers within the networking subsystem, Exadata redirects the I/O issued by the database server to another storage server.

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

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

Highest Level of Service

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

Oracle Platinum Services is available exclusively for Oracle's Engineered Systems. Platinum Services provides fault monitoring, faster response times, and expedited escalation to development. With Platinum Services, Oracle support engineers perform software maintenance and patching remotely. Platinum Services covers all software and hardware within an Engineered System, including the Oracle Database – the highest level of support ever for a full-stack software/hardware platform. Platinum Services is provided at no extra charge to Exadata customers.

IT Agility

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

Dramatically Lower Costs

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

Exadata provides a huge memory, flash, and disk footprint for large data sets. Raw disk storage on a fully populated Exadata Storage Expansion⁷ can reach 5 PB, while raw flash storage can be up to 2.3 PB. In addition, Hybrid Columnar Compression may also increase the effective storage and memory capacity by an average factor of 10. By intelligently moving active data across disk, flash, and memory tiers, Exadata simultaneously delivers the highest performance and the lowest cost.

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

The uniformity of Exadata Database Machine configurations results in large cost savings. **Exadata standardizes not just technologies but also integration, testing, security, hardening, tuning, and support.** Customers deploy Exadata systems much faster and with less labor than traditional systems. Low-level tuning, integration, and maintenance is reduced or eliminated. All Exadata users run a configuration that is identical to thousands of other users and are identical to Oracle's internal configurations, making it far less likely that issues will be encountered. When issues occur, the resolution is simpler as customers work with one supplier – Oracle, as the entire system – hardware, firmware, operating system, hypervisor, and database layers are all owned and supported by Oracle. The “one-hand-to-shake” support model enables faster problem resolution times and reduces downtime and associated costs, further increasing economic benefits.

Exadata Business Benefits

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

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

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

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

- Real-time business data analysis
- Faster financial closes

⁷ Fully populated Exadata Storage Expansion Rack with 19x Exadata X11M High Capacity Storage Servers or 19x Exadata X11M Extreme Flash Storage Servers

- 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. Coupled with cost savings, ease of management, and enhanced supportability result in greater business agility and efficiency.

The Exadata Storage Expansion Rack is the simplest, fastest way to add large amounts of additional storage capacity to an Exadata Database Machine. A natural extension of the Exadata Database Machine, the Exadata Storage Expansion Rack can be used to satisfy the requirements of the largest mission critical databases.

Given what can be achieved with Exadata, it is no surprise it is the new global standard for running Oracle Databases – on-premises or in the cloud.

Exadata Storage Expansion X11M Key Capacity and Performance Metrics

Metric	Quarter Rack Storage Expansion		Single Exadata Storage Servers			Storage Expansion Rack Maximum Configuration	
	High Capacity (HC)	Extreme Flash (EF)	High Capacity (HC)	Extreme Flash (EF)	High Capacity-Z (HC-Z)	High Capacity (HC)	Extreme Flash (EF)
Number of Storage Servers	4	4	1	1	1	19	19
Flash Metrics							
Maximum SQL flash and Exadata RDMA Memory bandwidth²	2,000 GB/s	2,000 GB/s	500 GB/s	500 GB/s	250 GB/s	9,500 GB/s	9,500 GB/s
Maximum SQL read IOPS³	11,200,000	11,200,000	2,800,000	2,800,000	1,400,000	53,200,000	53,200,000
Maximum SQL write IOPS⁴	4,000,000	4,000,000	1,000,000	1,000,000	500,000	19,000,000	19,000,000
Performance-optimized PCI Flash capacity (raw)⁷	108.8 TB	108.8 TB	27.2 TB	27.2 TB	13.6 TB	489.6 TB	516.8 TB
Capacity-optimized PCI Flash storage (raw)		491.5 TB		122.9 TB			2,334.7 TB
Disk Metrics							
Maximum SQL disk bandwidth²	7.2 GB/s		1.8 GB/s		0.9 GB/s	32.4 GB/s	
Maximum SQL disk IOPS³	10,000		2,600		1,300	46,800	
Data capacity (raw)⁵	1,056 TB		264 TB		132 TB	5,016 TB	
Combined Metrics							
Data Capacity (usable) Normal Redundancy⁶	408.2 TB	190.0 TB	102 TB	47.5 TB	51 TB	2,075.8 TB	966.0 TB
Data Capacity (usable) High Redundancy⁶	320.1 TB	149.0 TB	80 TB	37.2 TB	40 TB	1,520.7 TB	707.6 TB
Maximum data load rate⁸	10 TB/hour	10 TB/hour	2.5 TB/hour	2.5 TB/hour	2.5 TB/hour	47.5 TB/hour	47.5 TB/hour

Actual system performance varies by application.

¹ EF = Extreme Flash; HC = High Capacity; HC-Z = High Capacity-Z

² Bandwidth is peak physical scan bandwidth achieved running SQL, assuming no database compression. Effective user data bandwidth is higher when 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. Usable capacity is measured using normal powers of 2 space terminology with 1 TB = 1024 * 1024 * 1024 * 1024 bytes.

⁶ Actual space available for a database after mirroring (ASM normal redundancy) while also providing adequate space (one disk on Quarter and Half Racks and two disks on a Full Rack) to reestablish the mirroring protection after a disk failure in the normal redundancy case.

⁷ Effective Flash Capacity is larger than the physical flash capacity and takes into account the high flash hit ratios due to Exadata's intelligent flash caching algorithms, and the size of the 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.

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 Storage Expansion X11M Hardware

Quarter Rack Storage Server	Cores (Total)	Exadata RDMA Memory (Total)	Available Storage High Capacity (Raw)	Available Storage Extreme Flash (Raw)
4 x HC Storage Servers	256 CPU cores for SQL processing	5 TB	<ul style="list-style-type: none"> 48x 22 TB High Capacity Drives, and 16x 6.8 TB performance-optimized NVMe PCI Flash Cards 	
4 x EF Storage Servers	256 CPU cores for SQL processing	5 TB		<ul style="list-style-type: none"> 16x 30.72 TB capacity-optimized NVMe PCI Flash drives. and 16x 6.8 TB performance-optimized NVMe PCI Flash Drives for EF Quarter Rack
Additional Hardware Components included in both configurations: <ul style="list-style-type: none"> 42U Rack 3x 36-port 100 Gb/s RoCE switches Ethernet switch for administrative connectivity 2x Redundant Power Distribution Units (PDUs) 				

Exadata Storage Expansion X11M Connectivity and Upgrades

Connection to Exadata Database Machine	Upgradability
<p>Connect any combination of up to 14 Exadata Database Machine X11M racks or Exadata Storage Expansion X11M Racks via the RoCE Network Fabric.</p> <p>Larger Configurations can be built with external RoCE switches. Connected racks can be X11M, X10M, X9M or X8M generation hardware</p>	<p>After the initial quarter rack, additional HC, EF, HC-Z or combination of HC, EF and HC-Z storage servers can be added one at a time or as many as needed up to a maximum configuration (19 storage servers)⁸.</p> <p>Hardware Components included with the upgrade:</p> <ul style="list-style-type: none"> • RoCE and Ethernet cables and adapters to connect all the components • 12 x 22 TB High Capacity SAS Drives, 4 x 6.8 TB performance-optimized NVMe PCI Flash Cards, and 1.5 TB DDR5 DRAM (1.25 TB allocated to Exadata RDMA Data Accelerator and the remaining 256 GB allocated to Exadata System Software) for each additional HC storage server, or • 4x 30.72 TB capacity-optimized NVMe PCIe Gen 4 flash drives, 4 x 6.8 TB performance-optimized NVMe PCI Flash Cards, and 1.5 TB DDR5 DRAM (1.25 TB allocated to Exadata RDMA Data Accelerator and the remaining 256 GB reserved for Exadata System Software) for each additional EF storage server, or • 12 x 22 TB High Capacity Drives for each additional XT storage <p>Upgrade Support Services:</p> <ul style="list-style-type: none"> • Hardware Installation and Software Configuration

Exadata Storage Expansion X11M Environmental Specifications

Metric	Quarter Rack	Maximum Configuration	Single Server
Height		78.74 in. (2000 mm)	3.42 in. (86.9 mm)
Width		23.66 in. (601 mm)	17.52 in. (445.0 mm)
Depth		47.13 in. (1197 mm)	30.51 in. (775.0 mm)
Acoustic Noise (operating)	9.4 B	9.9 B	8.6 B
Environmentals With High Capacity Disks			
Weight	917.3 lb (416.1 kg)	2210.8 lb (1002.8 kg)	74 lb (33.6 kg)
Maximum Power Usage	5.3 kW (5.4 kVA)	20.5 kW (20.9 kVA)	1 kW (1 kVA)
Typical Power Usage ¹	3.7 kW (3.8 kVA)	14.4 kW (14.7 kVA)	0.7 kW (0.7 kVA)
Cooling at Maximum Usage	17,951 BTU/hour	70,055 BTU/hour	3,474 BTU/hour
	18,939 kJ/hour	73,908 kJ/hour	3,665 kJ/hour
Cooling at Typical Usage	12,566 BTU/hour	49,038 BTU/hour	2,431 BTU/hour
	13,257 kJ/hour	51,735 kJ/hour	2,565 kJ/hour
Airflow at Maximum Usage ²	831 CFM	3243 CFM	161 CFM
Airflow at Typical Usage ²	582 CFM	2270 CFM	113 CFM
Environmentals With Extreme Flash Disks			
Weight	861.3 lb (390.7 kg)	1944.8 lb (882.1 kg)	60 lb (27.2 kg)
Maximum Power Usage	4.7 kW (4.8 kVA)	18.1 kW (18.5 kVA)	0.9 kW (0.9 kVA)
Typical Power Usage ¹	3.3 kW (3.4 kVA)	12.7 kW (12.9 kVA)	0.6 kW (0.6 kVA)

⁸ Maximum number of storage servers varies. Refer to Oracle Exadata Configuration Assistant.

Cooling at Maximum Usage	16,204 BTU/hour	61,756 BTU/hour	3,037 BTU/hour
	17,095 kJ/hour	65,153 kJ/hour	3,204 kJ/hour
Cooling at Typical Usage	11,343 BTU/hour	43,229 BTU/hour	2,126 BTU/hour
	11,967 kJ/hour	45,607 kJ/hour	2,243 kJ/hour
Airflow at Maximum Usage ²	750 CFM	2859 CFM	141 CFM
Airflow at Typical Usage ²	525 CFM	2001 CFM	98 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 Maximum configuration added as examples of elastic configurations. 19x Storage Servers.</p>			

Exadata Storage Expansion X11M Regulations and Certifications

Regulations ^{1,2,3}	Product Safety:	UL/CSA 60950-1, EN 60950-1, IEC 60950-1 CB Scheme with all country differences
		UL/CSA 62368-1, EN 62368-1, IEC 62368-1 CB Scheme with all country differences
	EMC	
	Emissions:	FCC CFR 47 Part 15, ICES-003, EN55032, KS C 9832, EN61000-3-11, EN61000-3-12
	Immunity:	EN55024, KS C 9835
Certifications ^{2,3}	North America (NRTL), CE (European Union), International CB Scheme, HSE Exemption (India), BSMI (Taiwan), KC (Korea), RCM (Australia), VCCI (Japan), UKCA (United Kingdom)	
European Union Directives ³	2014/35/EU Low Voltage Directive. 2014/30/EU EMC Directive, 2011/65/EU RoHS Directive, 2012/19/EU WEEE Directive	
<p>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.</p>		

Exadata Storage Expansion X11M Support Services

<p>Components:</p> <ul style="list-style-type: none"> • Hardware Warranty: 1 year with a 4 hour 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 and Solaris 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 Exadata RDMA Memory Data Accelerator
- Unique Exadata RDMA Memory Commit Accelerator (X8M and X9M only)
- Unique Database Aware PCI Flash
- Unique Exadata Smart Flash Caching
- Unique Exadata Smart Flash Logging
- Unique Smart Write-back Flash Cache
- Unique I/O Prioritization by cluster, workload, DB or user to ensure QOS
- Unique Exafusion Direct-to-Wire Protocol
- Unique Database Intelligent Network Resource Management
- Unique Exachk full-stack validation
- Unique Full-stack security scanning
- Unique Database scoped security
- Unique Cell-to-Cell Rebalance preserving Flash Cache and Storage Index
- Unique Full-Stack Secure Erase
- Unique Instant Data File Creation
- Unique Smart Fusion Block Transfer
- Unique Control of Flash Cache Size per Database
- Unique In-Memory OLTP Acceleration
- Unique Undo-Block Remote RDMA Read
- Unique Support for 4096 Pluggable Databases per Container Database with Multitenant Option

Exadata and Database Software Features – Artificial Intelligence

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

Exadata and Database Software Features – High Availability

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

Manageability Features

- Oracle Embedded Integrated Lights Out Manager (ILOM) with upgrade pre-staging optimizations
 - Oracle Enterprise Manager Exadata Plug-in
 - Unique Active AWR includes storage stats for end-to-end monitoring
 - Real-time Insights server metric streaming
 - IPv6 Support for Ethernet Connections
 - Capacity on Demand
 - Cell software transparent restart
 - Flash and disk life cycle management alert
 - Automatic Disk Scrub and Repair
 - Trusted Partitions for Oracle Linux Virtualization
 - Automated VLAN Creation
 - Oracle Exadata Deployment Assistant
 - Separate Management Switch and Connectivity
 - Exacli 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, andCisco RoCE and management switches
-

Oracle Database Software (available separately):

- **For database servers:** Oracle Database 19c Enterprise Edition, Oracle Database 21c Enterprise Edition, and Oracle Database 23ai Enterprise Edition. Oracle Database Options such as Oracle Real Application Clusters, Oracle Partitioning, Oracle Multitenant, Oracle Active Data Guard. See the release specific documentation for feature support. Oracle Grid Infrastructure 19c or higher is required.

For storage servers: Oracle Exadata System Software. Licenses are transferable from one system to another, or to a new system.

Oracle Software (included):

For database servers: Oracle Linux 8 with the Unbreakable Enterprise Kernel 7. Zero-loss Zero-copy Datagram Protocol (ZDP) RoCEv2 protocol used to communicate between the Exadata Storage Servers and the Oracle Database which is based on the Reliable Datagram Sockets (RDS) OpenFabrics Enterprise Distribution (OFED)

Connect with us

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com). Outside North America, find your local office at: [oracle.com/contact](https://www.oracle.com/contact).

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2025, Oracle and/or its affiliates. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

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