

Oracle Exadata: A guide for decision-makers

September, 2026, Version 2.11
Copyright © 2026, Oracle and/or its affiliates
Public

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement, nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Introduction

The Oracle Exadata Database Machine (Exadata) is a computing and storage platform specialized for running Oracle databases. The goal of Exadata is higher performance and availability at lower cost by optimizing and integrating hardware and software at all levels and moving database algorithms and intelligence into storage and networking, bypassing the traditional layers of general-purpose servers.

Exadata is a combined hardware and software platform that includes scale-out database servers, scale-out intelligent storage servers, ultra-fast networking, memory acceleration, NVMe flash, and specialized Exadata Software in a wide range of shapes and price points. Exadata storage features high-performance servers to store data and run Exadata Software for data-intensive database processing directly in the shared storage tier.

A brief history

Exadata debuted in 2008 as the first offering in Oracle's family of Engineered Systems for use in corporate data centers deployed as private clouds. In October 2015, Exadata became available in the Oracle public cloud as a subscription service, known as the Exadata Cloud Service, later re-branded as Exadata Cloud Infrastructure in 2022, supporting both Exadata Database Service and Autonomous Database cloud services.

Oracle databases deployed with Exadata Database Service on Exadata Cloud Infrastructure are 100% compatible with databases deployed on Exadata on-premises, enabling customers to transition to the Oracle Cloud with zero application changes. Oracle manages this service, including hardware, network, Linux software, and Exadata software, while customers retain complete control over their databases.

In early 2017, a third Exadata deployment choice became available. Exadata Cloud@Customer is Exadata Cloud Infrastructure deployed on-premises (behind the customer's firewall) and managed by Oracle Cloud experts. Exadata Cloud@Customer is owned and managed by Oracle and acquired by customers through a pay-as-you-go subscription. The Oracle Cloud@Customer program brings all the benefits of the Oracle public cloud while addressing potential network latency, security, and regulatory concerns.

In 2018, Oracle introduced Oracle Autonomous Database – a cloud-based self-driving, self-securing, self-managing database that provides mission critical availability and security while reducing management costs. Oracle Autonomous Database is available on Exadata Cloud Infrastructure and Exadata Cloud@Customer deployments.

In 2019, the release of Exadata X8M enhanced Exadata's performance through the addition of two major technical breakthroughs - persistent memory (PMem) and RDMA (Remote Direct Memory Access) over Converged Ethernet (RoCE). Oracle Exadata X8M used RDMA directly from the database to access persistent memory in smart storage servers, bypassing the entire OS, I/O, and network software stacks. This produced lower latency and higher throughput.

In 2023, Exadata X10M replaced persistent memory with Exadata RDMA Memory (XRMEM), a new memory acceleration tier in storage, due to changes in the vendor landscape for persistent memory. Also significant was the incorporation of AMD processors in all Exadata servers, offering a large increase in the number of available compute cores.

In 2024, the release of Oracle Database 23ai and Exadata System Software 24ai enabled AI vector acceleration, including AI Vector Search offloads to Exadata storage. Another significant capability in these releases is Exadata Exascale, a unique intelligent data architecture for Exadata infrastructure. Exadata Exascale is comprised of hyper-elastic, RDMA-enabled, database-optimized computing and storage resources. Exadata Exascale enables organizations to leverage Exadata for all Oracle Database workloads regardless of size or development lifecycle stage with revolutionary new resource pooling and database snapshot capabilities. In addition, Exascale significantly lowers

Exadata's entry price point in the public cloud while maintaining all the extreme performance, reliability, availability, and security benefits, making Exadata Database Service affordable for any Oracle Database customer.

In late 2024, the release of Exadata System Software 25.1 introduced significant new capabilities: Adaptive Top-K filtering, Vector distance projection, and BINARY vectors. Enhancements have been made to Exascale free space management and volume snapshots, as well as the acceleration of ASM rebalance operations, resulting in increased availability.

In early 2025, the launch of Exadata X11M marked a significant milestone, providing efficiencies in performance, utilization, power, and administration. Exadata X11M adopts the latest AMD processors, providing up to 1.25x more concurrent and faster serial transactions than Exadata X10M. Exadata X11M is available in the Oracle public cloud, offering a significant performance increase over the Exadata X9M generation. Exadata X11M delivers exceptional value through accelerated performance, increased efficiency and sustainability, and support for Oracle Database 23ai, so customers can develop their applications anywhere, run them everywhere, and do more for the same price.

Exadata X11M is now available through our multicloud partnerships with the other major cloud providers – Microsoft Azure, Amazon AWS, and Google Cloud – providing the same performance as Exadata on-premises and Exadata in Oracle public cloud. The introduction of Exadata Database Service in these multicloud partners is a massive milestone for customers who can run their applications and databases in the same cloud regions, ensuring the lowest latency. The service also allows customers to run Oracle RAC, a key high-availability option, in the major clouds, which was not available earlier.

In July 2025, the release of Exadata System Software 25.2 introduced powerful new capabilities and extended Exadata's leadership as the premier platform for Oracle Database. The key focus areas in this release were: Accelerated AI Smart Scan, Extreme performance, Exascale innovations, Secure, elastic, and observable infrastructure.

In September 2025, a low-cost storage tier for storing less frequently accessed data with Exadata X11M Extended (XT) Storage Server was introduced. This storage server uses 1 AMD EPYC 32-core processor with 2 x 6.8 TB performance-optimized flash cards and 12 x 22 TB hard disks. The X11M XT storage server is the first generation of XT servers to include flash which enables Exadata Exascale.

In October 2025, Oracle AI Database 26ai became the long-term support release that replaced Oracle Database 23ai.

In May 2026, Exadata System Software 26ai, release 26.1, introduced the world's first online migration for RDMA-enabled database virtual machines with Exadata VM Live Migration. This groundbreaking capability increases business continuity, operational agility, and infrastructure resilience during planned maintenance by moving virtual machines between Exadata database servers with minimal disruption to running database services, applications, and users.

In June 2026, Oracle introduced Exadata X11 High Capacity (HC) and Extreme Flash (EF) storage servers, together with the X11-Z High Capacity (HC-Z) storage server. These servers expand Exadata's cost-optimized storage options by retaining its intelligent storage capabilities in a streamlined hardware configuration that omits XRMEM by design. X11 HC and EF are available for on-premises and Cloud@Customer deployments, while X11-Z HC-Z is available for on-premises deployments.

In September 2026, Oracle released Exadata System Software 26ai, release 26.2, adding post-quantum, also called quantum-resistant, cryptography and FIPS 140-3 validated cryptographic modules. Built on Oracle Linux 9, the release further strengthens Exadata's security and compliance posture.

Exadata use cases

Exadata is designed to optimally run any Oracle Database workload or combination of workloads, such as an OLTP application running simultaneously with analytics processing. The platform is frequently used to consolidate many databases that were previously running on dedicated general-purpose database servers. Exadata's scale-out architecture is naturally suited to running in the Oracle Cloud, where computing requirements can dynamically grow and shrink.

Historically, specialized database computing platforms were designed for a particular workload, such as data warehousing, and were poor or unusable for other workloads, such as OLTP. Exadata has optimizations for all database workloads, implemented such that mixed workloads share system resources fairly. Exadata Resource Management features automatically prioritize the allocation of system resources, such as favoring workloads servicing interactive users over reporting and batch processing, even if they are accessing the same data.

Long-running requests, typical of data warehouses, reports, batch jobs, and analytics, run many times faster compared to a conventional, non-Exadata database server. Customer references often cite performance gains of 10x or greater. Analytics workloads can also use the Oracle Database In-Memory option on Exadata for additional acceleration, and in-memory databases on Exadata have been extended to take advantage of flash, whose capacity is many times larger than the capacity of DRAM. Exadata's Hybrid Columnar Compression feature is intended to reduce the storage consumption of data warehouses and archival data as well as increase performance by reducing the amount of I/O.

Transactional (OLTP) workloads on Exadata benefit from the incorporation of XRMEM (memory acceleration) and flash into Exadata's storage hierarchy and the automatic tiering of data between XRMEM, flash, and disk storage. Special algorithms optimize response-time-sensitive database operations such as log writes. For the most demanding OLTP, all-flash storage eliminates the latency of disk media completely.

Exadata design concepts

To better understand the design of Exadata, it helps to compare it with a traditional database computing platform, assembled from separate hardware and software components operating independently.

Traditional computing platforms are general-purpose

The hardware components that make up a typical database computing platform are a database server connected over a network to a storage array. The database software runs on the database server and sends or receives data to and from the storage array over the network. The hardware components use standard software protocols to communicate with each other. This separation via standard interfaces is what allows a general-purpose computing platform to run a wide variety of workloads, software, and hardware from different vendors. All the application logic and the processing of data are performed on the database server, to which all the data must be sent. This approach enables using a computing platform for a wide range of software applications, though it will not be optimized for any particular application.

Oracle Database is the focus of Exadata

The goal of Exadata is to create a complete stack of software and hardware tailored to the Oracle Database, that performs processing in the optimal location. Because Exadata is only processing Oracle Database requests, it can take advantage of that focus in all of the software layers. The hardware design includes technologies such as very fast RDMA-enabled Ethernet networking, specialized DRAM caching (XRMEM), and flash storage integrated into the architecture to yield the most advantages to Oracle Database applications. Given the importance of data storage to databases, Oracle Exadata places particular focus on optimizing that aspect of the platform.

Exadata uses unique technologies in the storage layer that easily scale out and parallelize Oracle Database requests. The addition of flash memory and XRMEM to Exadata Storage Servers also opens up a range of possibilities for optimizing performance in the storage layer. For example, as the performance and capacity of flash storage increases at a rapid rate, the network can become a bottleneck for traditional database platforms, whereas Exadata's offloading of database processing into Exadata Storage Servers avoids that problem. The addition of XRMEM in the Exadata storage layer exposes the limitations of traditional platforms even more acutely.

Adding database intelligence to storage

At the time Exadata was conceived, Oracle had several decades of experience developing database software, and was well aware of the limitations and performance bottlenecks imposed by traditional computing platforms. To fulfill the Exadata mission, Oracle needed a storage layer that could easily scale out and parallelize Oracle Database requests. Oracle also recognized the opportunity for storage to cooperate in the processing of database requests beyond just storing and shipping data. For example, rather than send an entire database table across the network to the database server to find a small number of records, such data filtering could be done in storage so that only the resulting records need to be sent across the network.

In summary, Oracle recognized the need for a powerful server that could run intelligent database software and act as a storage array, with a modular design that could easily grow in capacity and performance as the database grew. Building a "database-aware" storage server that could cooperate with database servers in the execution of database requests became a compelling undertaking, enabled by focusing Exadata on what is best for the Oracle Database.

The database-aware Exadata Storage Server, invented by Oracle to replace the traditional storage array, is the foundation of Exadata.

Exadata Exascale

The release of Exadata System Software 24ai (24.1) software includes Exadata Exascale, the world's only intelligent data software architecture. Exascale comprises hyper-elastic, RDMA-enabled, database-optimized computing and storage resources. It is a loosely coupled architecture that uses remote direct memory access (RDMA) and SQL offload, delivering the same low latency and high throughput for which Exadata is known. Exadata Exascale's direct I/O architecture combined with Oracle Database 23ai sends requests directly to the Exascale storage pool, bypassing the intermediate tiers that add latency and bottlenecks in non-Oracle cloud databases.

One of Exascale's ground-breaking features is its reimagined approach to database snapshots and clones on Exadata. It enables space-efficient thin clones from any read/writes database, pluggable database, or previously taken snapshot, significantly boosting developer productivity. Exascale uses redirect-on-write technology that is more efficient than traditional copy-on-write.

Another benefit of Exascale is Exascale Volumes which allows more virtual machines to run using shared Exadata storage rather than local storage on the database servers and taking advantage of the increased number of cores in the latest Exadata generations.

Before Exascale, when used as a Cloud Service, Exadata infrastructure was provisioned starting with two database and three storage servers as the smallest configuration. Additional compute and storage capacity could be added as workloads increased. With Exascale, users can now provision a few cores and gigabytes of storage, making Exadata Database Service on Exascale Infrastructure a low-cost entry point with all the Exadata benefits including pay-per-use elastic scaling without subscribing to dedicated infrastructure. This capability extends all Exadata advantages to every organization in every industry, regardless of size or workload.

Optimizing across the full stack

To maximize the effectiveness of Exadata, Oracle controls the software and hardware components of the platform, so that coordinated improvements can be tightly integrated and made anywhere at any time.

Oracle already had a broad portfolio of software products when Exadata was conceived, covering most of the software layers that are required to run a database platform, such as the Oracle Linux operating system, storage management software, monitoring and administrative tools and virtual machine software, and, of course, Oracle Database and options software.

The initial 2008 release of Exadata (V1) was a joint development between Oracle (software) and Hewlett-Packard (hardware). The second generation (V2) of Exadata switched to hardware from Sun Microsystems, and shortly thereafter, Oracle acquired Sun Microsystems and thus gained ownership of the main hardware components of Exadata.

Owning the main hardware components of Exadata completes Oracle's ability to develop an entire computing platform optimized around the Oracle Database. An additional benefit for customers is the ability to support the entire Exadata platform from one vendor: something impossible with a traditional computing platform of hardware and software components from multiple vendors.

Exadata smart software

With the initial versions proving the value of the Exadata concept, Oracle embarked on a regular release schedule of new software capabilities closely aligned with the release and incorporation of leading-edge hardware components.

Using knowledge of Oracle Database formats and algorithms, plus an understanding of the application workload, these "smart" software enhancements, running in the Oracle Database and all software and firmware in the platform, outperform the capabilities of a traditional database platform. The ability to run Oracle Database routines in the Exadata Storage Server software is a major enabler of this advantage.

Software enhancements specific to Exadata achieve better performance in some areas due to the integration of hardware and software. For example, when an OLTP application commits a transaction to the database software, that request is viewed by Exadata as a critical operation and prioritized accordingly within the network and in the storage servers. The transaction commit request will jump ahead of less urgent messages on the network and in the I/O queue.

Another example is the use of flash memory for caching data accessed by analytics applications. Because a columnar data format is more effective for analytics workloads, when Exadata moves the row-formatted data from disk to flash, it automatically reformats the data into a columnar format.

These examples illustrate how the Oracle Database understands the intent of the application and cascades this understanding to the network and storage software, which then behave accordingly.

Since Exadata's debut in 2008, there have been one to two significant Exadata software releases per year, delivering dozens of "smart" software enhancements. Most of these enhancements are based on a smaller number of essential technical foundations, as explained below.

Technical foundations

Offloading to Storage – refers to the execution of data-intensive database operations within the Exadata Storage Servers, such as data scans, table joins, and filtering of rows and columns. Sending just the description of the operation and getting back filtered results substantially reduces the network traffic between the database servers and storage servers. This avoids the network bottleneck of traditional architectures where data-intensive operations require shipping large amounts of data between storage and database servers. Offloading is possible because Exadata storage is built on standard servers, capable of running database functions in coordination with the database server, simultaneous with storage I/O. Over time, more database functions and more data types have been offloaded. In addition, "reverse offloading" will push an operation back to the database servers if Exadata storage is too busy.

Storage Indexes – enable the avoidance of I/O by tracking column values within relatively small regions of storage. Storage Indexes are automatically maintained and kept in memory on Exadata Storage Servers. If the Storage Index

indicates that an I/O to a region will not find a match, that I/O is avoided, which yields a significant performance benefit. Initially, Storage Indexes track value ranges within a small number of columns. Over time, more columns and more sophisticated value tracking have been added, so that additional classes of I/O operations can be avoided.

Flash and XRMEM Caching – delivers the low latency (fast response) of flash and DRAM while preserving the lower cost of disk for storing large databases for the best I/O performance at the lowest cost. In general, a small percentage of a database is active at any one time. If just the active data is held in flash, for instance, the I/O performance would be equal to all-flash storage at a much lower cost. Exadata monitors the current workloads and keeps the most active data in flash or XRMEM in the optimal format. For example, Exadata knows when an I/O is part of a database backup, and not an indication of an active data block, whereas traditional storage arrays view any I/O as a "hot" block. Flash caching will also reformat rows into columnar format in flash for data being accessed for analytics. Initially, flash caching was only used for reading data, but has since been enhanced to include log writes and all other write I/O. The flash cache is also used as an extension of Oracle's Database In-Memory columnar data store, for significantly larger in-memory databases than DRAM capacity alone. XRMEM adds an even faster cache in storage and substantially improves I/Os per second (IOPS) and latency.

Hybrid Columnar Compression (HCC) – reduces the amount of storage consumed by infrequently updated data, such as data warehouses, that can grow to enormous sizes. Conventional data compression algorithms yield between 2x and 4x compression, whereas HCC averages between 10x and 15x compression due to the greater compressibility of columnar formats. Such a large reduction in the amount of I/O can also substantially improve performance. Initially, HCC tables did not support row-level locking, limiting their use with OLTP applications. In 2016, support for row-level locking was added to HCC on Exadata, improving the performance of mixed workloads with HCC data. The hybrid format of HCC enables Exadata to avoid the performance pitfalls of columnar-only databases.

Resource Management – allocates Exadata system resources, such as CPU, I/O, and network bandwidth, to databases, applications, and users based on priorities. When consolidating many databases on Exadata, Resource Management ensures the appropriate quality of service. I/O Resource Management debuted in Exadata V1, and Network Resource Management was added in Exadata X4.

In-Memory Databases – offer exceptional performance for Analytics workloads, leveraging DRAM on database servers, a complement to Exadata's emphasis on storage and networking. Oracle Database In-Memory became available in 2014 on Exadata, leveraging its fast internal network for In-Memory Fault Tolerance. To support larger in-memory databases, Exadata Storage Servers implement in-memory routines and in-memory data formats in Exadata flash, as an extension of the same in-memory processing that occurs on database servers.

AI Smart Scans – Starting with Oracle Database 23ai, the VECTOR data type is introduced that stores vectors in thousands of dimensions. The vector search functionality enables semantic searches on unstructured data like images, videos, etc. Oracle Database 23ai, in conjunction with Exadata System Software 24ai, introduces AI Smart Scan, a collection of unique Exadata software optimizations capable of improving the performance of AI vector search operations by orders of magnitude. AI Smart Scan processes vector data at memory speed, leveraging ultra-fast Exadata RDMA Memory (XRMEM) and Exadata Smart Flash Cache in the Exadata storage servers, and performs highly scalable vector distance computations and top-K filtering where the data lives, avoiding unnecessary network data transfer and database server processing.

Smart software enhancements

Below is a more detailed listing of software enhancements, grouped by their value to analytics or OLTP workloads and their impact on database availability and security. Similar enhancements cannot be duplicated on conventional platforms because they require modifications to system software and APIs and integration across database software, operating systems, networking, and storage.

For Analytics

- Automatically parallelize and offload data scans to storage
- Filter rows in storage based on 'where' clause
- Filter rows in storage based on columns selected
- JSON and XML offload
- Filter rows in storage based on join with another table
- Hybrid Columnar Compression
- Storage Index data skipping
- I/O Resource Management by User, Query, Service, DB, etc.
- Automatic transformation to In-Memory columnar format from Flash Cache
- Smart flash caching for table scans
- Storage offload index fast full scans
- Storage offload scans on encrypted data, with FIPS compliance
- Storage offload for LOBs and CLOBs
- Storage offload for min/max operations
- Data mining offload to storage
- Reverse offload to DB servers if storage CPUs are busy
- Automatic data columnarization
- Automatic conversion of Data to in-memory formats when loading into Flash Cache

For OLTP

- Exadata RDMA memory data accelerator
- XRMEM commit accelerator (X8M and X9M only)
- Database-aware PCIe flash interface
- Exadata smart flash caching
- Exadata smart flash logging
- Smart write-back flash cache
- I/O Resource Management by DB, user, and workload to ensure QoS
- Exafusion direct-to-wire protocol
- Database intelligent network resource management
- EXAchk full-stack validation
- Full stack security scanning
- Database scoped security
- Cell-to-cell rebalance preserving flash cache and storage index
- Full-stack secure erase
- Instant data file creation
- Control of flash cache size per database
- In-memory OLTP acceleration
- Undo-block remote RDMA read
- Support for 4000 Pluggable Databases per Container Database with Multitenant option

Artificial Intelligence

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

For Availability

- Online RoCE switch software updates with continuous traffic forwarding
- Instant detection of node or cell failure
- Sub-second failover of I/O on stuck disk or flash
- Prefetch OLTP data into secondary mirror Flash Cache
- Offload incremental backups to storage servers
- Instant data file creation
- Prioritized rebalance of critical files
- Cell-to-cell rebalance to preserve Flash Cache population
- Automatic rebalance on predictive disk failures
- Instant detection of node or cell failure
- Sub-second failover of I/O on stuck disk or flash
- Prefetch OLTP data into secondary mirror Flash Cache
- Offload incremental backups to storage servers
- Instant data file creation
- Prioritized rebalance of critical files
- Exadata VM Live Migration
- Cell-to-cell rebalance to preserve Flash Cache population
- Automatic rebalance on predictive disk failures
- Instant detection of node or cell failure
- Automatic rebalance on disks predicted to fail
- Automatic monitoring of CPU, network and memory using Machine Learning
- Automatic identification of underperforming disks
- Automatic Software Updates on an entire "fleet" of Exadata systems with one operation
- Cell software transparent restart
- Online Linux patching (Ksplice)

For Security

- Comprehensive monitoring and auditing functionality at the server, network, database, and storage layers
- Secured access to perform secure lights-out management of (ILOM) database and storage servers
- Audit record of all logons and configuration changes
- Secure RDMA fabric isolation
- Multi-pass secure erase of disks and flash
- Firewall-protected Exadata Storage Servers
- Secure network access
- Secure RDMA fabric isolation

- Comprehensive monitoring and auditing functionality at the server, network, database, and storage layers
- Secured access to perform secure lights-out management of (ILOM) database and storage servers
- Audit record of all logons and configuration changes
- FIPS 140-3 certification PCI-DSS compliance
- Minimal Linux distribution
- SELinux policies and management tools tailored for Exadata strengthen platform security while simplifying system operations
- Fast, hardware-based (AES) encryption/decryption
- Full-stack security scanning
- Database and ASM-scoped security
- Multi-pass secure erase of disks and flash
- Fast, secure eraser of disk and flash (Crypto erase)
- Advanced Intrusion Detection Environment (AIDE) detects and alerts unknown changes to system software
- InfiniBand partitioning
- Support for IPV6
- Secure computing filter to restrict system calls
- Centralized identification and authorization of OS users

For Power Management and Sustainability

- Intelligently turn off unneeded CPU cores
- Intelligently caps power consumption
- Intelligently optimizes power for low usage windows

Database software

Exadata X11M database servers run the Oracle Linux 9 operating system, Oracle Database 26ai, and Oracle Database 19c Enterprise Edition. Exadata system resources can be optionally virtualized using the KVM-based Oracle hypervisor. All Oracle Database options, such as Real Application Clusters, Multitenant, Database In-Memory, Advanced Compression, Advanced Security, Partitioning, Active Data Guard, and others are optionally available with Exadata.

Applications that are certified for a supported version of the Oracle Database are automatically compatible with Exadata. No additional modifications or certifications are required. The same database software that runs on Exadata on-premises will run on Exadata Cloud Infrastructure, Exadata Cloud@Customer, and multicloud deployments in Microsoft Azure, Google Cloud, and Amazon AWS. In addition, on-premises software licenses are eligible for the Bring Your Own License (BYOL) transfer into the Oracle public cloud, Exadata Cloud@Customer, or multicloud deployments. Oracle Autonomous Database is available exclusively on Exadata cloud platforms.

Networking

Exadata provides high-speed networks for internal and external connectivity. A 100 gigabit per second (100 Gbit/s) RDMA-enabled Ethernet fabric is used for internal connectivity between database and storage servers and includes database cluster interconnect traffic. For external client connectivity, 100, 25, and 10 Gbit/s Ethernet ports are available.

Exadata uses a custom-designed, database-oriented protocol over the Ethernet fabric to achieve higher performance. It makes extensive use of Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) to improve

efficiency by avoiding operating system overhead and extra copies when moving data between servers. Exadata also has a direct-to-wire protocol that allows the database to communicate directly to the RoCE network cards.

Exadata takes advantage of RoCE Class of Service in its Network Resource Management feature to prioritize important traffic across the network. In this feature, the Oracle Database software tags network messages that require low latency, such as transaction commits, lock messages, and I/O operations issued by interactive users, prioritizing them over messages issued by less critical high-throughput workloads such as reporting and batch processing. The result is analogous to how an emergency vehicle with its siren on can move more quickly through heavy traffic - high-priority network messages are moved to the front of the server, network switch, and storage queues, bypassing lower-priority messages, and resulting in shorter and more predictable response times.

Management software

For Exadata systems deployed in traditional on-premises configurations, Oracle Enterprise Manager (EM) supports a single pane of glass view of all the Exadata hardware and software components such as database servers, storage servers, and network switches, and monitors the operations running on them. EM integrates with the built-in Exadata management tooling, as well as with customers' existing systems management and helpdesk tools.

The Exadata Cloud Infrastructure, Exadata Cloud@Customer, and Exadata multicloud deployment platforms are managed by Oracle Cloud Infrastructure operations, while customers control and manage the software and databases running on the database servers. Lifecycle operations for Oracle databases are performed using a web browser, command-line interface (CLI), or REST API-driven automation available through the Cloud Control Plane, including provisioning, updating, scaling, and backup.

Hardware

Before Exadata X10M, Exadata was available in two models: one based on 2-socket database servers and the other based on 8-socket database servers. The adoption of high-core count AMD processors with Exadata X10M database servers removed the need for the 8-socket Exadata model, offering the same number of cores in a much more cost-effective configuration. Thus, with Exadata X10M and newer, only one Exadata hardware model is available, and the -2 and -8 suffixes on the names are removed.

The latest Exadata generation, X11M, was introduced in January 2025. The X11M database and storage servers use a 2 Rack Unit (RU) form factor. Exadata X11M drives the efficiency for performance, utilization, power, and administration compared to previous generations.

The Exadata Database Machine base configuration has 2 database servers and 3 storage servers, referred to as a Quarter Rack. For workloads that can be satisfied with less computing and memory than the standard Exadata database server, the X11M-Z Database Server can be deployed in a flexible configuration, saving power and cost. As the database workload and/or data size increases, additional database and storage servers may be added to increase the volume of work performed in parallel, using Exadata's elastic configuration. Multi-rack Exadata configurations can scale large workloads that exceed a single rack.

Exadata database servers

There are two choices for Exadata X11M database servers: 96-core 2-socket (192 cores) and 32-core 1-socket (32 cores) called X11M-Z Database Servers. The X11M-Z database servers provide the flexibility to run smaller workloads in the same standard rack with power efficiency. Exadata X11M rack is compatible with prior RoCE generations of servers and can reside alongside each other.

Exadata storage servers

Starting with Exadata X11M, there are four choices for Exadata storage servers: Extreme Flash (EF), High Capacity (HC), High Capacity-Z (X11M HC-Z), and low-cost X11M XT servers. The X11M Extreme Flash Storage Server is all-

flash storage containing four performance-optimized and four capacity-optimized flash drives for 27.2 TB of Exadata Smart Flash Cache and 122.88 TB of raw flash storage capacity. Each storage server contains 1.25 TB of XRMEM as an acceleration tier in front of Flash Cache to boost performance further. The X11M High-Capacity Storage Server contains twelve 22 TB disk drives with 264 TB total raw disk capacity, 27.2 TB of Exadata Smart Flash Cache, and 1.25 TB of Exadata RDMA Memory. Exadata Smart Flash Cache is managed automatically by Exadata Smart Storage Software.

The X11M High Capacity-Z storage server contains six disks, 22 TB each, for a total of 132 TB of raw storage capacity, 192 GB of system memory, and 576 GB of Exadata RDMA Memory. X11M-Z HC servers provide the flexibility to run compute-intensive workloads with cost-efficient storage for less data-intensive applications.

The X11M XT storage server uses 1 AMD EPYC 32-core processor with 2 x 6.8 TB performance-optimized flash cards and 12 x 22 TB hard disks. The X11M XT storage server is the first generation of XT server to include flash for Exadata Smart Flash Cache which enables Exadata Exascale. Flash Cache takes advantage of the 2 x 6.8 TB performance-optimized flash cards, the same cards as in the EF, HC, and HC-Z storage servers, and automatically caches the hottest data in flash. Optionally, Flash Cache can also be used to accelerate analytic workloads by enabling (and licensing) Smart Scan. Flash Cache takes advantage of the performance-optimized flash cards, the same cards as in the EF, HC, and HC-Z storage servers, and automatically caches the hottest data in flash. This avoids I/O requests to the hard disks, lowers OLTP read and write latency significantly, and delivers 600,000 OLTP read and 500,000 OLTP write IOPS. When enabled, Smart Scan can deliver up to 50 GB/s SQL scan throughput. X11M XT Storage Servers require Exadata System Software 25ai, release 25.2.0 (released in July 2025) or later.

Exadata X11 HC, X11 EF, and X11-Z HC extend the Exadata architecture to storage-server configurations optimized for different capacity and performance requirements. The servers use a streamlined configuration without XRMEM; however, XRMEM is available as a future expansion option. X11 HC features 2 x 32-core AMD EPYC processors, 12 x 22 TB HDDs, and 4 x 6.8 TB flash devices. X11 EF uses the same processor platform, with 4 x 30.72 TB capacity-optimized flash devices and 4 x 6.8 TB flash devices. X11-Z HC features 1 x 32-core AMD EPYC processor, 6 x 22 TB HDDs, and 4 x 6.8 TB flash devices.

In addition to adding storage servers into an Exadata Database Machine Quarter Rack configuration, storage servers may also be acquired with or added to Exadata Storage Expansion racks.

Performance specifications for a Quarter Rack Exadata X11M configuration with EF, HC, HC-Z, or XT storage servers are as follows:

Storage Server	Max Scan Rate from Flash	Max SQL Read IOPS from XRMEM	Max SQL Write IOPS to Flash
X11M Extreme Flash	300 GB/sec	5,600,000	3,000,000
X11M High Capacity	300 GB/sec	5,600,000	3,000,000
X11M-Z High Capacity	150 GB/sec	4,200,000	1,500,000
X11M-eXTended	150 GB/sec	1,800,000 (from flash)	1,500,000
X11 Extreme Flash	300 GB/sec	3,600,000 (from flash)	3,000,000
X11 High Capacity	300 GB/sec	3,600,000 (from flash)	3,000,000
X11-Z High Capacity	150 GB/sec	1,800,000 (from flash)	1,500,000

Table 1. Maximum performance based on a Quarter Rack configuration of 2 database servers and 3 storage servers. XT servers can only be populated after 2 DB Servers and 3 storage servers in a rack.

Memory-level performance with shared storage

Architects of traditional database platforms have always had to cope with technology change affecting the design of their systems. Their goal is to eliminate bottlenecks so that the output of storage moves through the network and is processed by database servers without any slowdown. Solving an imbalance generally involves adding faster or more network connections or database servers. This was before the advent of ultra-fast PCIe flash memory and the NVMe flash interface.

Flash drives started to become mainstream in corporate computing around 2010, used as a cache in front of hard disks or as a replacement for disks entirely. Every year thereafter flash capacity and performance increased significantly. Flash technology has seen tremendous innovations in the last several years, and throughput for individual flash drives has grown significantly. Most storage vendors can only provide a fraction of the throughput of the flash card as the network remains the substantial bottleneck in traditional database systems. For example, one popular storage solution can only provide 175 GB/s of data throughput with 480 of the latest flash drives. Conversely, Exadata intelligently bypasses the network bottleneck by filtering most of the data within the storage and sending only relevant data across the network to the database servers. The addition of Exadata RDMA Memory in Exadata storages increases the throughput even further from up to 1,700 GB/s for flash alone to 8,500 GB/s in a single rack. Offloading to storage in Exadata bypasses this network bottleneck by filtering out unneeded data in storage before sending the remaining data across the network. The addition of XRMEM, which is faster than flash, increases the value of Exadata offloading even further. While adding flash directly to a traditional database server removes the network bottleneck, it also removes the ability to share storage across multiple database servers. Exadata's approach does not suffer this limitation.

Exadata Generation (2-socket)	V1	V2	X2-2	X3-2	X4-2	X5-2	X6-2	X7-2	X8-2	X8M	X9M	X10M	X11M
Date Introduced	Sep-2008	Sep-2009	Sep-2010	Sep-2012	Nov-2013	Jan-2015	Apr-2016	Oct-2017	Apr-2019	Sep-2019	Sep-2021	June-2023	Jan-2025
Last Ship Date	Oct-2009	Oct-2010	Sep-2012	Feb-2014	Mar-2015	Jul-2016	Nov-2017	Jun-2019	Dec-2020	Sep-2022	Nov-2024	N/A	N/A
Data Sheet Specifications	N/A	N/A	N/A	N/A	X4-2	X5-2	X6-2	X7-2	X8-2	X8M-2	X9M-2	X10M	X11M
Operating System	Linux	Linux	Linux	Linux	Linux	Linux	Linux	Linux	Linux	Linux	Linux	Linux	Linux
Disk Storage (raw TB/PB)	168	336	504	504	672	1344	1344	1680	2.35 PB	2.35 PB	3 PB	3.6 PB	3.6 PB
Flash Cache (raw TB)	N/A	5.3	5.3	22.4	44.8	89.6	179.2	358	358	358	358	380	380
Persistent Memory (TB)	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	21	21	N/A	N/A
Exadata RDMA Memory (XRMEM) (TB)	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	17.5 TB	17.5 TB
Extreme Flash (raw TB)	N/A	N/A	N/A	N/A	N/A	179.2	358.4	716.8	716.8	716.8	716.8	1.7 PB	1.7 PB
Compute Cores	64	64	96	128	192	288	352	384	384	384	512	1,536	1,536
Max Memory (Gib)	256	576	1,152	2,048	4,096	6,144	12 TB	12 TB	12 TB	12 TB	16 TB	24 TB	24 TB
RDMA Network Fabric (Gb/sec)	20	40	40	40	40	40	40	40	40	100	200	200	200
Ethernet (Gb/sec)	8	24	184	400	400	400	400	800	800	800	800	800	800

Table 2. Key statistics for each Exadata 2-socket system since initial introduction. Comparisons are based on configurations of 8 database servers and 14 storage servers.

Exadata Generation (8-socket)	V1	V2	X2-8	X3-8	X4-8	X5-8	X6-8	X7-8	X8-8	X8M-8	X9M-8
Date Introduced	N/A	N/A	Sep- 2010	Sep- 2012	Jul- 2014	Nov- 2015	Apr- 2016	Oct- 2017	Apr- 2019	Sep- 2019	Sep- 2021
Last Ship Date	N/A	N/A	Nov-2012	Dec- 2014	Oct- 2015	Mar- 2016	Nov- 2017	Jun- 2019	Dec- 2020	Jun- 2022	Aug- 2024
Data Sheet Specifications	N/A	N/A	X2-8	X3-8	X4-8	X5-8	X6-8	X7-8	X8-8	X8M-8	X9M-8
Operating System	Linux	Linux	Linux	Linux	Linux	Linux	Linux	Linux	Linux	Linux	Linux
Disk Storage (raw TB)	N/A	N/A	504	504	672	1344	1344	1680	2352	2352	3024
Flash Cache (raw TB)	N/A	N/A	5.3	22.4	89.6	89.6	179.2	358.4	358.4	358.4	358.4
Persistent Memory (TB)	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	21	21
(XRMEM) Exadata RDMA Memory	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Extreme Flash (raw TB)	N/A	N/A	N/A	N/A	179.2	179.2	358.4	716.8	716.8	716.8	716.8
Compute Cores	N/A	N/A	96	160	240	288	288	384	384	384	384
Max Memory (GB)	N/A	N/A	4	4	12	12	12	12	12	12	12
RDMA Network Fabric (Gb/sec)	N/A	N/A	40	40	40	40	40	40	40	100	100
Ethernet (Gb/sec)	N/A	N/A	176	176	180	180	180	540	540	540	540

Table 3. Key statistics for each Exadata 8-socket system, initially introduced with Exadata X2-8. X9M-8 was the last 8-socket model. Comparisons are based on configurations of 2 database servers and 14 storage servers.

Elastic Configurations

Prior to the X5-2 generation, Exadata systems were only available in fixed-size configurations of Eighth, Quarter, Half, and Full Rack sizes. The X5-2 Exadata release in January 2015 introduced elastic configurations. An elastic configuration has a customer-specified combination of database servers and storage servers, allowing individual storage or database servers to be added to a Quarter Rack configuration until the physical rack is full. The ratio of database to storage servers can vary, depending on the characteristics of the intended workload. For example, an Exadata system optimized for in-memory database processing would add many database servers, each with maximum memory. Conversely, an Exadata system optimized for a large data warehouse could add many High-Capacity storage servers. Elastic configurations may also be used to scale out earlier-generation Exadata systems using the latest compatible servers. In addition, Exadata Database Machines have always been able to span multiple racks using the built-in network fabric connections. Thus, Exadata's scale-out extends beyond a single physical rack.

Exadata Evolution

Oracle releases a new generation of Exadata every twelve to thirty-six months. At each release, Oracle refreshes most hardware components to the latest CPU processors, memory, disk, flash, and networking technologies. These hardware refreshes result in performance increases with every release. Exadata software innovations, delivered with each generation and periodically in between, consistently enhance performance, availability, security, management, and workload consolidation.

The evolution of Exadata is best understood through the innovations introduced in each generation.

Exadata V1, released in 2008, focused on accelerating data warehousing by delivering the full throughput of storage to the database. Exadata achieved this by moving database filtering operations into storage, instead of sending all data to the database servers and filtering it there. This capability is referred to as Exadata Smart Scan. Exadata V1 also supported a consolidation feature for allocating I/O bandwidth between databases or workloads, called IORM (I/O Resource Manager). Exadata V1 was available in Full Rack or Half Rack sizes, and the choice of High Performance or High-Capacity storage servers, both based on disk drive for storage. The internal network fabric of Exadata was based upon InfiniBand technology.

Exadata V2, released in 2009, added a Quarter Rack configuration and support for OLTP workloads via flash storage and database-aware Flash Caching.

Exadata V2 also introduced Hybrid Columnar Compression to reduce the amount of storage consumed by large Data Warehouses.

Storage Indexes in Exadata V2 increased performance by eliminating the need to read entire regions of storage, based on the storage server's knowledge of the values contained in the region.

Exadata X2-2, the third generation, was released in 2010 and a second model of Exadata, [Exadata X2-8](#), was introduced. The X2-8 and subsequent 8 socket Exadata models featured processors targeted at large memory, scale-up workloads. The use of flash storage beyond caching began in this release with a Smart Flash Logging feature. Support for 10 Gigabit per second (Gb/sec) Ethernet client connectivity was also added.

Exadata X2-2 also encouraged data security through encryption with the incorporation of processor-based hardware decryption, largely eliminating the performance overhead of software decryption.

A Storage Expansion Rack based on Exadata X2-2 was added in 2011 to accommodate large, fast-growing data warehouses and archival databases. All subsequent Exadata generations have included a new Storage Expansion Rack.

Exadata X3-2 and **X3-8** were released in 2012, including a new Eighth Rack X3-2 entry-level configuration. Flash storage capacity quadrupled and OLTP write throughput increased by 20x via the Write-Back Flash Cache feature.

A number of availability enhancements were added, bypassing slow or failed storage media, reducing the duration of storage server brownouts, and simplifying replacement of failed disks.

Exadata X4-2 was released in 2013. Flash capacity doubled and flash compression was added, effectively doubling capacity again. Network Resource Management was introduced, automatically prioritizing critical messages. InfiniBand bandwidth effectively doubled with support for active/active connections.

Exadata X4-8 was released in 2014, adding Capacity on Demand licensing, I/O latency capping, and timeout thresholds to all subsequent models.

Exadata X5-2 and **X5-8** were released in 2015 with major enhancements. Flash and disk capacity doubled. Elastic configurations were introduced to enable expansion one server at a time. Virtualization was added as an option to Exadata along with Trusted Partitions for flexible licensing within a virtual machine. Database snapshots on Exadata storage enabled efficient development and testing. Oracle Database In-Memory on Exadata included Fault Tolerant

redundancy. The High Performance Exadata storage servers were replaced with all-flash (Extreme Flash) storage servers and Exadata became the first major vendor to adopt the NVMe flash interface. Columnar Flash Cache was introduced to automatically reformat analytics data into columnar format in flash. [IPv6](#) support was completed. Exadata Cloud Service was launched on the Oracle Cloud.

Exadata X6-2 and X6-8 were released in 2016. Flash capacity doubled. Exafusion Direct-to-Wire protocol reduced messaging overhead in a cluster and Smart Fusion Block Transfer eliminates redo log write delays when transferring blocks between database nodes. Exadata Cloud@Customer debuted, enabling Oracle Cloud benefits within corporate data centers.

Exadata X7-2 and X7-8 were released in 2017. Flash capacity doubled. Flash cards became hot-pluggable for online replacement. 10 Terabyte (TB) disk drives debuted along with 25 Gb/sec Ethernet client connectivity. Oracle Database In-Memory processing was extended into flash storage, and storage server DRAM was utilized for faster OLTP.

Exadata X8-2 and X8-8 were released in April 2019. Exadata Extended (XT) Storage Server was introduced for low-cost storage of infrequently accessed data. 14 Terabyte (TB) disk drives debuted along with 60% more compute cores in Exadata storage servers. Machine Learning algorithms were added that automatically monitor CPU, network, and memory to detect anomalies such as stuck processes, memory leaks, and flaky networks, and to automatically create (Auto index), rebuild, or drop indexes. Optimizer statistics are also gathered in real-time as DML executes. For enhanced security, Advanced Intrusion Detection Environment (AIDE) was added to detect and alert when unsanctioned changes to system software are made.

Exadata X8M-2 and X8M-8 were released in September 2019. Substantial performance increases resulted from the addition of Intel Optane DC Persistent Memory in Exadata Storage Servers, and a new 100 Gbit/s internal network fabric based on RoCE (RDMA over Converged Ethernet), replacing the previous InfiniBand fabric. These changes increased read I/O throughput by 2.5x and lowered I/O latency by 10x. In addition, a new KVM hypervisor replaced the Xen hypervisor, doubling the amount of memory available to a guest VM.

Exadata X9M-2 and X9M-8 were released in September 2021 and included the latest generation of Intel Optane Persistent Memory and PCIe Gen 4, which led to significant performance gains over the previous generation. OLTP read I/O throughput increased a further 1.6x and the 1 TB/s Smart Scan threshold was crossed within a single rack.

Exadata X10M was released in June 2023 as a 2-socket only model based on AMD processors with 96-core per socket database servers. The high core count and large memory capacity removed the need for the 8-socket Exadata model. Persistent memory was replaced with Exadata RDMA memory, based on DRAM. Disk storage capacity and all-flash storage capacity increased. Database servers increased in size from 1 to 2 RU for better airflow and cooling.

Exadata X11M launched in January 2025 with significant AI, analytics, and OLTP performance gains from faster AMD processors. Its flexible configurations allow one-socket X11M-Z database and X11M-Z High-Capacity storage servers to share racks with standard servers, pairing high compute with high-capacity storage and low compute with extreme-performance all-flash storage. X11M also improved power efficiency, performance, utilization, and administration to support customer sustainability goals. In June 2026, X11 High Capacity, Extreme Flash, and X11-Z High Capacity storage servers expanded performance-, capacity-, and cost-optimized options. They retain Exadata's intelligent storage without XRMEM; HC and EF are available on premises and Cloud@Customer, while HC-Z is on premises only. Introduced in September 2025, low-cost X11M XT supports less frequently accessed data and, unlike earlier XT servers, both ASM and Exascale configurations.

Connect with us

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

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2026, 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, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.