

Was ist dran an einer spezialisierten Data Warehousing platform?

Hermann Bär
Oracle USA
Redwood Shores, CA

Schlüsselworte

Data warehousing, Exadata, specialized hardware proprietary hardware

Introduction

Specialization in IT is as old as the industry itself. And so is the discussion whether or not ‘best-of-breed’ single-purpose solutions are the best approach to solving a problem. A purpose-built system simply must be better. But is this really true? What exactly does it mean when talking about a purpose-built system for data warehousing? Does purpose-built automatically mean that you have to accept a one-trick pony? Oracle does not think so.

Oracle Exadata Database Machine – the optimal platform for all database workloads

The Oracle Exadata Database Machine is an easy to deploy solution for hosting the Oracle Database that delivers the highest levels of database performance available. The Exadata Database Machine is a “cloud in a box” composed of database servers, Oracle Exadata Storage Servers, an InfiniBand fabric for storage networking and all the other components required to host an Oracle Database, running any kind of database workload. *The Oracle Exadata Database Machine is a general-purpose Database Machine.*

Extreme performance is delivered for all types of database applications by leveraging a massively parallel grid architecture using Real Application Clusters and Exadata storage. The Oracle Exadata Database Machine delivers outstanding I/O and SQL processing performance for online transaction processing (OLTP), data warehousing (DW) and consolidation of mixed workloads. *The Oracle Exadata Database Machine combines purpose-built functionality and techniques – for Data Warehousing, OLTP, and consolidation - in a single system.*

Exadata Storage Server – storage built for the Oracle database

The Exadata Storage Server is an integral component of the Exadata Database Machine. One of the key enablers of Exadata’s extreme performance is the Exadata Smart Flash Cache hardware and the intelligent Oracle Exadata Storage Server Software that drives it. It is a combination of several unique features, is purpose-built for various database use case scenarios and takes advantage of both the software and hardware contained in the Exadata Storage Server.

Exadata storage provides database aware storage services, such as the ability to offload database processing from the database server to storage, and provides this while being transparent to SQL processing and database applications. Just the data requested by the application is returned rather than all the data in the queried tables. Exadata Smart Flash Cache dramatically accelerates Oracle Database

processing by speeding I/O operations. The Flash provides intelligent caching of database objects to avoid physical I/O operations and speeds database logging. The Oracle Database on the Database Machine is the first Flash enabled database. Exadata storage provides an advanced compression technology, Hybrid Columnar Compression, that typically provides 10x, and higher, levels of data compression. Exadata compression boosts the effective data transfer by an order of magnitude. The Oracle Exadata Database Machine is the world's most secure database machine. Building on the superior security capabilities of the Oracle Database, the Exadata storage provides the ability to query fully encrypted databases with near zero overhead at hundreds of gigabytes per second. The combination of these, and many other, features of the product are the basis of the outstanding performance of the Exadata Database Machine.

Exadata Storage Server – the hardware

The Oracle Exadata Storage Servers comes with either twelve 600 GB 15,000 RPM High Performance SAS disks or twelve 3 TB 7,200 RPM High Capacity SAS disks. The High Performance SAS disk based Exadata Storage Servers provide up to 3.25 TB of uncompressed useable capacity, and up to 1.8 GB/second of raw data bandwidth. The High Capacity SAS disk based Exadata Storage Servers provide up to 16 TB of uncompressed useable capacity, and up to 1.3 GB/second of raw data bandwidth. When stored in compressed format, the amount of user data and the amount of data bandwidth delivered by each cell significantly increases.

By pushing database processing to the Exadata Storage Servers all the disks can operate in parallel reducing database server CPU consumption while using much less bandwidth to move data between storage and database servers. Each Exadata Storage Server has sufficient CPU and memory resources to handle all smart database operations that are pushed down into the storage tier.

Exadata Storage Servers also contain Flash storage in addition to the traditional spinning disks. Exadata Flash storage is not a disk replacement – Exadata software intelligence determines how and when to use the Flash storage, and how best to incorporate Flash into the database as part of a coordinated data caching strategy. Scale out Exadata storage enables the benefits of flash performance to be delivered all the way to the application. Traditional storage arrays have many internal and network bottlenecks that prevent realizing the benefits of flash. Flash can be added to storage arrays, but they cannot deliver much of the potential performance to applications. Exadata storage delivers close to 1 GB/sec of throughput from each flash card and scales that performance linearly across the 4 cards in every Exadata Storage Server. Traditional storage arrays do not allow flash cards to be added to the system. Their architecture would need to be redesigned to avoid the disk controller limitations.

The hardware components of the Exadata Storage Server were carefully chosen to match the needs of high performance database processing. The Exadata software is optimized to take the best possible advantage of the hardware components and Oracle Database. Each Exadata Storage Server delivers outstanding I/O performance and bandwidth to the database. The CPU cores in the Exadata Storage Server are dedicated to providing features such as Smart Scan SQL processing that is done in the Exadata storage.

Building on the high security capabilities in every Oracle Database, the Exadata storage provides the ability to query fully encrypted databases with near zero overhead at hundreds of gigabytes per second. This is done by moving decryption processing from software into the Exadata Storage Server hardware, leveraging the built-in capabilities of Intel's processors. The Oracle software and the Intel 5600 processors used in the Exadata Storage Server provide Advanced Encryption Standard (AES) support enabling this.

Exadata Storage Server Software

The Exadata Storage Server runs the Exadata Storage Server Software and provides the unique and powerful Exadata software technology of the Database Machine including Smart Scan, Smart Flash Cache, Smart Flash Logging, IO Resource Manager, Storage Indexes and Hybrid Columnar Compression.

Purpose-built for data warehousing: brute-force massively parallel data scans

Exadata smart scan processing integrates database functionality within the storage layer of the database stack; it allows queries, and other database operations, to be executed much more efficiently. Implementing database functionality as close to the hardware as possible, in the case of Exadata at the disk level, can dramatically speed database operations and increase system throughput. With Exadata storage, database operations are handled much more efficiently. Queries that perform table scans can be processed within Exadata storage with only the required subset of data returned to the database server. Row filtering, column filtering and some join processing (among other functions) are performed within the Exadata storage cells. When this takes place only the relevant and required data is returned to the database server.

Storage Indexes are a very powerful capability provided in Exadata storage that helps avoid I/O operations. The Exadata Storage Server Software creates and maintains a Storage Index (i.e., metadata about the database objects) in the Exadata cell. The Storage Index keeps track of minimum and maximum values of columns for tables stored on that cell. When a query specifies a WHERE clause, but before any I/O is done, the Exadata software examines the Storage Index to determine if rows with the specified column value exist in the cell by comparing the column value to the minimum and maximum values maintained in the Storage Index. If the column value is outside the minimum and maximum range, scan I/O for that query is avoided. Many SQL Operations will run dramatically faster because large numbers of I/O operations are automatically replaced by a few lookups. To minimize operational overhead, Storage Indexes are created and maintained transparently and automatically by the Exadata Storage Server Software.

Purpose-built for data warehousing: store large amount of read-mostly data

Exadata provides a very advanced compression capability called Hybrid Columnar Compression (HCC). Hybrid Columnar Compression enables the highest levels of data compression and provides enterprises with tremendous cost-savings and performance improvements due to reduced I/O. Average storage savings can range from 10x to 15x depending on how HCC is used. With average savings of 10x IT managers can drastically reduce and often eliminate their need to purchase new storage for several years. For example, a 100 terabyte database achieving 10x storage savings would utilize only 10 terabytes of physical storage. With 90 terabytes of storage now available, IT organizations can delay storage purchases for a significant amount of time.

HCC is a new method for organizing data within a database block. As the name implies, this technology utilizes a combination of both row and columnar methods for storing data. This hybrid, or best of both worlds, approach achieves the compression benefits of columnar storage, while avoiding the performance shortfalls of a pure columnar format. A logical construct called the compression unit is used to store a set of Hybrid Columnar-compressed rows. When data is loaded, column values are detached from the set of rows, ordered and grouped together and then compressed. After the column data for a set of rows has been compressed, it is fit into the compression unit.

Purpose-built for data warehousing: process advanced analytics in the storage

Data Mining model scoring is offloaded to Exadata. This makes the deployment of data warehouses on Database Machine an even better and more performant data analysis platform. All data mining scoring

functions (e.g., prediction_probability) are offloaded to Exadata for processing. This will not only speed warehouse analysis but reduce database server CPU consumption and the I/O load between the database server and Exadata storage.

Purpose-built for data warehousing and OLTP: smart Flash Cache

Oracle's Exadata Smart Flash Cache features are unique. Exadata Flash storage is not a disk replacement – Exadata software intelligence determines how and when to use the Flash storage, and how best to incorporate Flash into the database as part of a coordinated data caching strategy. Scale out Exadata storage enables the benefits of flash performance to be delivered all the way to the application. Traditional storage arrays have many internal and network bottlenecks that prevent realizing the benefits of flash. Flash can be added to storage arrays, but they cannot deliver much of the potential performance to applications.

One of the key enablers of Exadata's extreme performance is the Exadata Smart Flash Cache hardware and the intelligent Oracle Exadata Storage Server Software that drives it. The Exadata Smart Flash Cache feature of the Exadata Storage Server Software intelligently caches database objects in flash memory, replacing slow, mechanical I/O operations to disk with very rapid flash memory operations. The Exadata Storage Server Software also provides the Exadata Smart Flash Logging feature to speed database log I/O. Exadata Smart Flash Cache is one of the essential technologies of the Oracle Exadata Database Machine that enables the processing of up to 1.5 million random I/O operations per second (IOPS), and the scanning of data within Exadata storage at up to 75 GB/second.

Purpose-built for every kind of workload: Resource Management

The DBRM and I/O resource management capabilities of Exadata storage can prevent one class of work, or one database, from monopolizing disk resources and bandwidth and ensures user defined SLAs are met when using Exadata storage. The DBRM enables the coordination and prioritization of I/O bandwidth consumed between databases, and between different users and classes of work. By tightly integrating the database with the storage environment, Exadata is aware of what types of work and how much I/O bandwidth is consumed. Users can therefore have the Exadata system identify various types of workloads, assign priority to these workloads, and ensure the most critical workloads get priority.

In data warehousing, or mixed workload environments, you may want to ensure different users and tasks within a database are allocated the correct relative amount of I/O resources. For example you may want to allocate 70% of I/O resources to interactive users on the system and 30% of I/O resources to batch reporting jobs. This is simple to enforce using the DBRM and I/O resource management capabilities of Exadata storage.

Conclusion

Businesses today increasingly need to leverage a unified database platform to enable the deployment and consolidation of all applications onto one common infrastructure. Whether OLTP, DW or mixed workload a common infrastructure delivers the efficiencies and reusability the datacenter needs – and provides the reality of grid computing in-house. Building or using custom special purpose systems for different applications is wasteful and expensive. The need to process more data increases every day while corporations are also finding their IT budgets being squeezed. Examining the total cost of ownership (TCO) for IT software and hardware leads to choosing a common high performance infrastructure for deployments of all applications.

By incorporating the Exadata based Database Machine into the IT infrastructure, companies will:

- Accelerate database performance and be able to do much more in the same amount of time.
- Handle change and growth in scalable and incremental steps by consolidating deployments on to a common infrastructure.
- Deliver mission-critical data availability and protection.

Oracle Exadata Database Machine combines powerful purpose-built functionality for data warehousing without any sacrifices for running any kind of database workload. The end of purpose-built data warehousing platforms is looming.

Kontaktadresse:

Hermann Bär
Oracle USA
500 Oracle Parkway, MS 40p742
Redwood Shores, CA 94065
USA

Telefon:	+1 650-506-6833
Fax:	+1 650-506-6833
E-Mail	hermann.baer@oracle.com
Internet:	www.oracle.com