Compliments of Oracle Press

Sneak Preview of New Books from Oracle Press

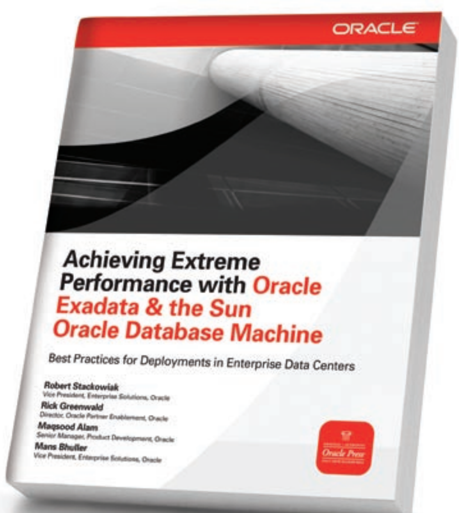

Achieving Extreme Performance with Oracle Exadata & the Sun Oracle Database Machine
Robert Stackowiak, Rick Greenwald, Maqsood Alam, Mans Bhuller

Managing Oracle Fusion Applications
Richard Bingham

Learn more. Mc Graw Hill Do more.
MHPROFESSIONAL.COM

OraclePressBooks.com

ORIGINAL • AUTHENTIC Oracle Press ONLY FROM McGRAW-HILL

Available March 2011

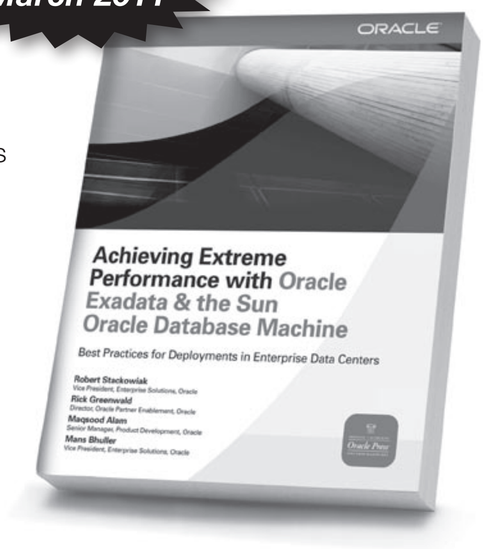

Achieving Extreme Performance with Oracle Exadata & the Sun Oracle Database Machine

Best Practices for Deployments in Enterprise Data Centers

Robert Stackowiak, Rick Greenwald, Maqsood Alam, Mans Bhuller

978-007-175259-6

A comprehensive guide to all the technology in the Sun Oracle Database Machine and the Exadata Storage Server, coupled with best practices that can only be delivered by experienced technologists like this author team.

TABLE OF CONTENTS

Once the Sun Oracle Database Machine is deployed in your data center, you must perform ongoing administrative tasks on the software and hardware components that make up the Sun Oracle Database Machine, as you would with any other system.

Trained personnel perform the initial setup and configuration of the Database Machine at the customer deployment site. The setup process includes the installation and configuration of the operating system and the Oracle software, and the configuration of InfiniBand and Ethernet networks. At the end of the setup process, you can use the Sun Oracle Database Machine to create new database instances and to start migrating your data from non-Exadata platforms, or to start deploying your brand-new applications to the newly built Exadata platform.

The Exadata Storage Server components such as the Exadata Storage Server software and the Linux operating system (OS) need to be administered on an ongoing basis. The administrative tasks include patching and upgrading of the OS and the Exadata Storage Server software, monitoring and tuning of the software components, and maintenance of the hardware components. Apart from the monitoring and maintenance activities, you may also perform tasks related to the setup of advanced Exadata Storage Server software features such as I/O Resource Manager and Exadata security.

This chapter will discuss the different processes and tools available for performing administrative tasks on the Exadata Storage Server and its components. The topics covered in this chapter are

- Exadata Storage Server components and architecture
- Exadata Storage Server administration
- Exadata Storage Server monitoring

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

NOTE

This chapter does not focus on administering the database servers in the Sun Oracle Database Machine. The database servers are managed, maintained, and monitored similarly to a regular Oracle Database 11g Release 2 deployed on Linux-based servers, with Oracle Real Application Clusters (Oracle RAC) and Oracle Automatic Storage Management (Oracle ASM). You would patch the Oracle database and the grid infrastructure software as needed, perform administrative tasks, monitor and tune the databases, back up the database and OS, and perform database upgrades, just as you would on standard Oracle systems. Moreover, you need to manage the operating system and the hardware as you would normally manage any other Linux-based operating systems and Intel-based servers that are deployed in your data center.

Exadata Storage Server Components and Architecture

At a high level, the Sun Oracle Database Machine comprises a database grid, a storage grid, and the InfiniBand network providing the unified fabric for storage and database inter-instance communication. The Sun Oracle servers that compose the database and the storage grids are built using open standards–based hardware and are made of components that are normally found in enterprise class servers such as CPUs, memory, PCIe (Peripheral Component Interconnect Express) slots, hard disks, Host Channel Adapters (HCAs), and network cards. The Oracle Enterprise Linux (OEL) operating system (OS) manages the Sun Oracle servers in both the database and the storage grid.

In the following section, we will discuss the architecture of the different components of the Sun Oracle Database Machine in detail.

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

Database Server Components

The database grid consists of one or more Oracle RAC or single-instance Oracle databases and utilizes Oracle ASM as the storage management layer. Oracle ASM acts as the cluster volume manager for the storage served by the Exadata Storage Servers. Oracle ASM is the only software that can be used to manage the Exadata Storage Servers, and you cannot directly attach or mount the Exadata-provided storage to any other server or storage volume manager.

Oracle ASM provides data placement and management services on the Exadata Storage Servers. It also stripes and mirrors the database extents into chunks called *ASM Allocation Units* (ASM AU, also called ASM extents) and spreads them across the disks served by the Exadata Storage Servers. Oracle ASM provides high availability of the extents and eliminates single points of failure by keeping a mirrored copy of the extent available at all times. The striping of data by ASM provides optimal I/O performance because the I/O requests will be parallelized across all the available Exadata Storage Servers in the grid that houses the striped data.

Oracle ASM will be active during the initial placement of data and also during ongoing management operations, such as the rebalancing of data across the Exadata Storage Servers. The database-server processes initiate I/O requests directly with the Exadata Storage Servers and bypass the Oracle ASM instance, thus avoiding the overhead of using ASM for each database I/O operation. The direct communication of the database with the storage is possible because the database server stores the extent maps of ASM files in the SGA, where the instance can initiate read/write I/O requests directly with the Exadata Storage Servers without requiring the Oracle ASM instance. This method of initiating I/O requests is the same when the ASM files reside on non-Exadata-based storage.

The database server processes and the Oracle ASM instances communicate with the Exadata Storage Servers using the *i*DB protocol. The *i*DB protocol is a data transfer protocol that operates at the disk block level. However, unlike other block-request data transfer protocols such as iSCCI, *i*DB is bit more intelligent and has knowledge about the internal structure of the database storage requests. This knowledge enables *i*DB to pass information about the filter predicates of the SQL statements, which in turn lets the Exadata Storage Servers initiate Exadata Smart Scans, thereby enabling the

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

computation of SQL predicates directly in the Storage Server. When the predicate is not offloaded to be processed by the Exadata Storage Servers, then the *i*DB merely acts like the other block-request protocols.

The *i*DB protocol also handles the I/O bandwidth allocations for Exadata I/O Resource Management (IORM) and provides bandwidth aggregation and failover of the InfiniBand network. The *i*DB communication is facilitated by software processes that reside on the Oracle ASM, the database nodes, and the Exadata Storage Servers, as shown in Figure 5-1.

The Database Resource Manager (DBRM) component of the Oracle database performs management of CPUs and I/O resources. The DBRM communicates with the Exadata Storage Server processes for performing I/O Resource Management for intra-database resource manager plans. IORM utilizes the DBRM intra-database plans generated by the administrators and regulates the I/O utilization among the different consumer groups within the database.

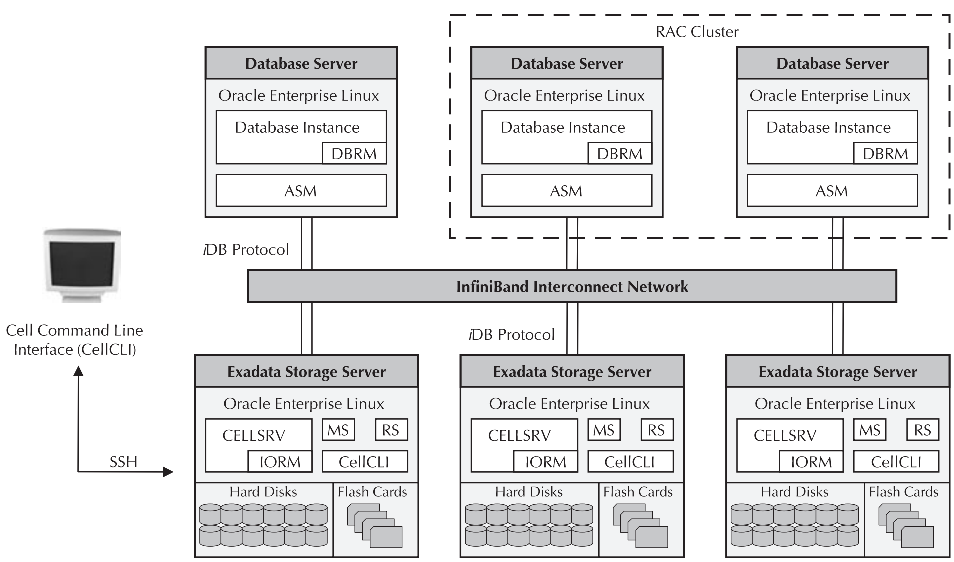

FIGURE 5-1. *Software components of the Sun Oracle Database Machine*

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

Exadata Storage Server Components

The storage grid in the Sun Oracle Database Machine comprises multiple Exadata Storage Servers, with each server providing a network-accessible storage device to be utilized by Oracle ASM. The operating system managing the Exadata Storage Servers is Oracle Enterprise Linux (OEL), and the Oracle software providing the core functions of the Exadata Storage Server is the Oracle Exadata Storage Server software.

The Exadata Storage Server software consists of the following components:

- **Cell Server (CELLSRV)** The CELLSRV process is the main kernel of Exadata Storage Server software and is responsible for serving the *i*DB requests initiated by the database and ASM instances. The CELLSRV process handles simple block I/O requests as well as Exadata Smart Scan requests. The CELLSRV process implements IORM plans and performs throttling of I/O bandwidth based on I/O priorities as defined by the plan. CELLSRV also interacts with DBRM for implementing intra-database resource manager plans.
- **Management Server (MS)** The MS process provides the management and monitoring functions for the Exadata Storage Server software. The MS process is responsible for triggering alerts when exceptions are encountered by the Exadata Storage Server components. The MS process communicates with the CELLSRV process and the operating system utilities in order to check the health of the software and hardware components.
- **Restart Server (RS)** The RS process handles startup and shutdown requests of the Exadata Storage Server software. The RS process monitors the CELLSRV and MS processes, and restarts them when they die abruptly. A backup RS process takes care of restarting the primary RS process if the primary is not running.
- **CellCLI** The Cell Command Line Interface (CellCLI) is the utility that you will use to perform management and administrative functions on the Exadata Storage Server software. More details on CellCLI are covered later in this chapter.

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

Exadata Storage Layout

The components of the Exadata Storage Servers that provide persistent storage services are the 12 SAS- or SATA-based hard disks and the 16 flash disks. The 16 flash disks are created across the four Sun Flash Accelerator F20 PCIe cards that are available in each Exadata Storage Server.

In this section, we will discuss the best practices for architecting these storage devices and for presenting them to Oracle ASM. Once the storage is presented to Oracle ASM, you will create ASM Diskgroups and will eventually create tablespaces in the Oracle database that use the ASM Diskgroups to store database objects.

LUNs, Cell Disks, and Grid Disks As shown in Figure 5-2, each hard disk gets presented to the Exadata Storage Server software and the OS as a Logical Unit (LUN). A LUN is what is visible of the hard disk to the OS. The Exadata Storage Server software will take the LUN and format it to create a *cell disk*. A LUN and a cell disk have a one-to-one relationship.

When the cell disk is created, it reserves a small portion of the available storage on the LUN for a system area. The system area is used to store information about the LUNs, cell disks, and the grid disks, which are discussed next, that will be created on top of the cell disks. The system area on the first two drives is special and is larger than the system area on the remaining ten drives. The system area on the first two drives is used to store the operating system and the Linux file system of the Exadata Storage Server.

Grid disks are created on top of cell disks, and they add another layer of logical abstraction to the physical storage. Grid disks are the entities that get presented to Oracle ASM as ASM Disks. Oracle ASM will create ASM

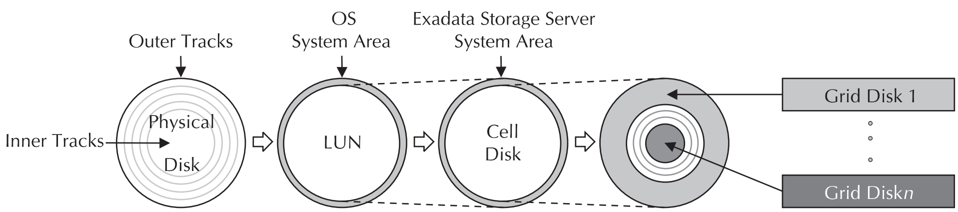

FIGURE 5-2. *Relationship between LUNs, cell disks, and grid disks*

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

Diskgroups across the available grid disks. The ASM Diskgroups are the entities that are presented to the database for creating tablespaces.

You can create multiple grid disks per cell disk. When the grid disk is created, it will carve out a chunk of space from the unused outer tracks of the physical disk. The first grid disk you create on the cell disk will allocate space from the outer tracks and move toward the inner tracks, reserving the number of tracks that correspond to the size of the grid disk. This grid disk provides the fastest performance because the outer tracks have shorter seek times and provide the best read/write performance. The next grid disk you create starts from the tracks where the first grid disk ends, and this process repeats until either you exhaust all the space on the cell disk, or you are done creating the grid disks.

Creating multiple grid disks per cell disk allows you to create multiple pools of storage on the same Exadata Storage Server. The multiple grid disks can be assigned to separate ASM Diskgroups, which can be provisioned to different databases. Even within the same database, you can store your active data on the grid disks created on the outer tracks, and store the less active data on the grid disks created on the inner tracks. The best practice for creating grid disks is to create two grid disks per cell disk. The first grid disk will be used to store user data since it provides the best performance. The second grid disk can be used to store the database Fast Recovery Area (FRA) for storing the database flashback logs, archive logs, and backups.

Interleaved Grid Disks The drawback of creating multiple grid disks per cell disk is that it does not provide the second grid disk the opportunity to be placed on the faster outer tracks. Moreover, the space utilization within a grid disk is such that only half of the space tends to be used for hot data, and the remaining half is used for mirrored blocks or colder data, which does not need the higher performance of the outer tracks. This situation can be mitigated by creating *interleaved grid disks*, shown in Figure 5-3, instead of regular grid disks.

With interleaved grid disks, you would provide the second grid disk the opportunity to be placed toward the faster-performing outer tracks. This mechanism is very useful when you set up the Exadata Storage Server for multiple storage pools by creating multiple ASM Diskgroups. To equalize the performance, all the ASM Diskgroups can use the hot areas of the cell disk.

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

Interleaving is enabled at the cell disk level by setting the INTERLEAVING attribute when the cell disk is created. When you enable interleaving, it divides the cell disk into two equal parts: the first part is the faster-performing ("hot") part and is allocated toward the outer tracks; the second part is the slower-performing ("cold") part and is allocated toward the inner tracks. When you create the interleaved grid disk on the cell disk, half of the interleaved grid disk is placed in the hot area, and the remaining half is placed in the cold area. The process will be repeated when you create multiple interleaved grid disks until you exhaust the space on the cell disk or until you no longer wish to create additional grid disks.

Flash Disks and Exadata Smart Flash Cache Each Exadata Storage Server in the Sun Oracle Database Machine comes preinstalled with four Sun Flash Accelerator F20 PCIe cards. The PCIe-based flash storage technology is primarily used to speed up access to data, since I/Os against flash are much faster than I/Os against the hard disks. Each flash card is divided into four modules, and each module is exposed to the operating system as one LUN; this LUN is categorized as a flash disk LUN. The total number of flash disk LUNs exposed to the operating system on each Exadata Storage Server is 16 (4 flash cards × 4 LUNs = 16 flash disks).

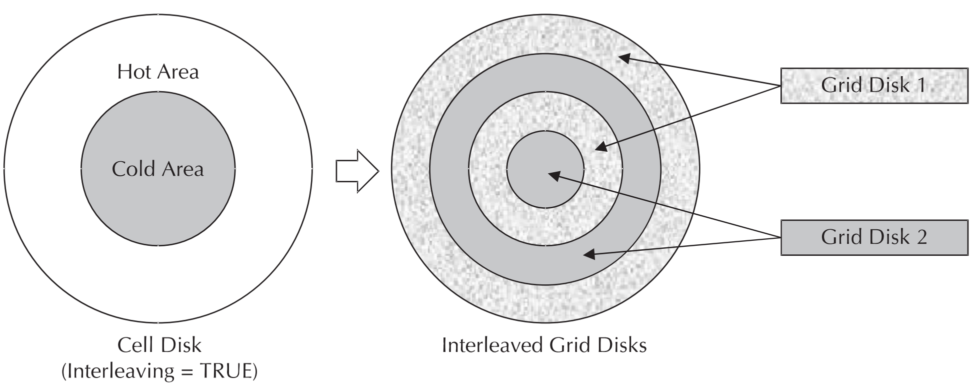

FIGURE 5-3. *Interleaved grid disks*

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

The three main options for managing and utilizing the flash disks in the Exadata Storage Server are highlighted here:

- **Exadata Smart Flash Cache** The first option is to let the Exadata Storage Server software manage the flash disks and configure them as the Exadata Smart Flash Cache. The Exadata Smart Flash Cache stores the frequently accessed (hot) data into the flash cards so the I/O times on hot data are improved dramatically, whereas the cold data resides on the cost-effective hard disk–based storage. When the database server requests read or write operations from the Exadata Storage Server software, it sends additional information about the request that indicates whether the blocks are candidates to be read again. Based upon this additional information, the Exadata Storage Server decides on caching the blocks in the Exadata Smart Flash Cache.

  When the I/O requests are for random reads and index reads, the blocks are likely to be read again and hence will be cached by the Exadata Smart Flash Cache algorithm. In the case of table scans, I/Os involving backups, redo log writes, and ASM mirrored writes, the block most likely will not be read again and hence will bypass the Exadata Smart Flash Cache—unless the object is marked with the KEEP qualifier, as described in Chapter 3 and in the last bullet of this list.

  Exadata Smart Flash Cache is configured by default and is also the best practice for managing the flash disks since the Exadata Storage Server software will be responsible for managing the flash without requiring any user intervention or tuning.

- **User-Managed Flash Disks** The second option for managing the flash disks is to have them partially managed as Exadata Smart Flash Cache and partially managed by the user for manually placing database objects, as shown in Figure 5-4. When the user manages the flash disk, the grid disks are created on a portion of the flash-based cell disks and exposed to Oracle ASM. Oracle ASM creates ASM Diskgroups that reside on the flash disks and thus provide the extreme I/O performance of flash storage. This option allows you to place the hot (frequently accessed) tables directly on a tablespace that is created on the flash-based ASM Diskgroup.

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

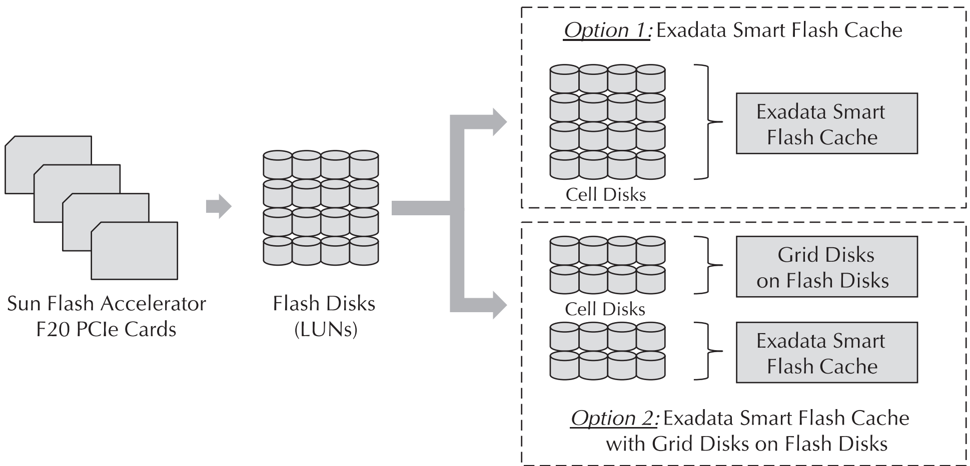

FIGURE 5-4. *Relationship between flash disks, Exadata Smart Flash Cache, and flash-based grid disks*

The drawback of this method is that it requires the user to analyze and identify the tables that qualify as "frequently accessed"; also, when the access pattern of the data changes, the user needs to perform ongoing management of the flash disk space by moving objects in and out of the flash-based ASM Diskgroup.

- **Pinning Objects in Exadata Smart Flash Cache** This option is an extension to the first option, and it allows the user to pin database objects such as tables, indexes, partitions, and LOB columns in the Exadata Smart Flash Cache. To pin the object at the object level, you will specify the CELL_FLASH_CACHE storage attribute. The CELL_FLASH_CACHE attribute tells the Exadata Storage Server software to pin the object in the Flash Cache or to skip caching the object altogether. The Exadata Smart Flash Cache's caching algorithm gives the pinned objects special treatment; it caches the pinned objects less aggressively and allows them to remain in the cache longer than other objects.

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

The CELL_FLASH_CACHE storage attribute can be set to KEEP, DEFAULT, and NONE. When the attribute is set to KEEP, the object will be pinned in the Flash Cache. When set to NONE, the object will never be cached. The DEFAULT attribute invokes the default caching behavior on the object, which is what was discussed in the first option.

The default behavior of the caching algorithm for Exadata Smart Scans is to avoid placing the Smart Scan blocks in the Exadata Smart Flash Cache. However, if you have the CELL_FLASH_CACHE set to KEEP on the table undergoing a Smart Scan operation, the table will be cached in the Flash Cache. Tables can be moved in and out of the Flash Cache with a simple ALTER command, without the need to move the table to different tablespaces, data files, or LUNs as you would have to do with traditional storage created on flash disks.

In this example, the ALTER statement will pin the *stocks* table in the Flash Cache:

```
SQL> ALTER TABLE stocks STORAGE (CELL_FLASH_CACHE KEEP);
```

NOTE
Be aware that changing the CELL_FLASH_CACHE from KEEP to NONE or DEFAULT will not remove the kept blocks from the Flash Cache—this change will simply subject those blocks to the more aggressive aging algorithm used for these alternative settings.

ASM Disks and ASM Diskgroup ASM Diskgroup is the primary storage entity of ASM that is exposed to the Oracle database for storing database files. The ASM Diskgroup comprises one or more ASM Disks, which are the grid disks exposed from each Exadata Storage Server to Oracle ASM. It is important to know that all ASM Disks within an ASM Diskgroup should reside in the Exadata Storage Servers in order for the database to utilize the Exadata-related features such as Exadata SQL offload processing (Exadata Smart Scans).

With Oracle ASM, you can define *failure groups* in an ASM Diskgroup. An ASM failure group comprises the ASM Disks within the ASM Diskgroup that have the tendency to fail together because they share the same hardware. The reason for creating failure groups is to identify the ASM Disks

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

ch05.ps
P:\010Comp\Oracle_Fluff\259-5\ch05.vp

that are candidates for storing mirrored copies of data. You do not want to store the mirror copy of an ASM extent in an ASM Disk that belongs to the same failure group, because if they fail together, you will lose your data. When you are creating ASM Diskgroups in an Exadata environment, the failure groups are created automatically by the ASM instance. This is because the ASM instance knows to group the ASM Disks that belong to the same Exadata Storage Server together in the same failure group.

An Exadata Grid Redundant Array of Inexpensive Disks (Grid RAID) uses ASM mirroring capabilities by specifying a redundancy level when creating the ASM Diskgroup. For the ASM extents to be available during Exadata Storage Server failures, you must use ASM mirroring by creating the ASM Diskgroup with NORMAL or HIGH redundancy. The NORMAL redundancy option will allocate one mirror copy of the ASM extent in a different failure group as the primary copy, whereas the HIGH redundancy option will allocate two mirrored copies of the ASM extent in two separate failure groups. The NORMAL redundancy tolerates a single Exadata Storage Server failure or multiple failures within the same Storage Server at a time. The HIGH redundancy tolerates two Exadata Storage Server failures or multiple failures confined to two Storage Servers at a time in the storage grid. When you choose the redundancy level, ensure that the post-failure capacity and performance due to the reduced number of Exadata Storage Servers in the Database Machine provide an acceptable level of service to the applications.

The example in Figure 5-5 shows two Exadata Storage Servers—*Server A* and *Server B*. Two ASM Diskgroups are created across *Server A* and *Server B*

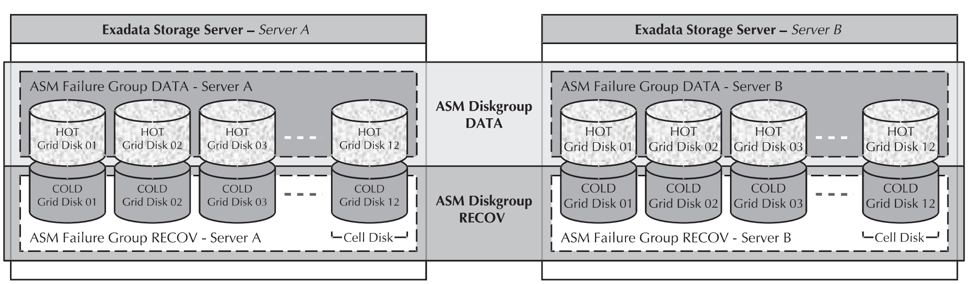

FIGURE 5-5. *ASM Diskgroup architecture*

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

in the grid. The DATA diskgroup is defined on the faster-performing (hot) grid disks, and the RECOV diskgroup is defined on the slower-performing (cold) grid disks. The diskgroups are created with NORMAL redundancy, which implicitly creates two failure groups per each diskgroup. The extents from *Server A*'s failure group will be mirrored to *Server B*'s failure group. If *Server A* fails, then Oracle ASM transparently fetches the mirrored extent from *Server B* and satisfies the I/O request without incurring any downtime.

As you know, the name of the ASM Disk is the same as the Exadata Storage Server grid disk and can be retrieved from the ASM instance by querying the *name* attribute of *v$asm_disk*. By looking at the ASM Disk name, you can pinpoint the Exadata Storage Server and the cell disk the ASM Disk belongs to. This is possible because the ASM Disk name contains the Exadata Storage Server name and the cell disk name as a prefix. The system adds the prefix, which is useful in locating the cell disk to which the grid disk belongs. However, when you have multiple grid disks created per cell disk, then in order to identify whether the grid disk was created on the *hot* or the *cold* areas of the hard disk, you need to add a user-defined prefix to the system-defined name. This is accomplished by using the PREFIX attribute of the CREATE GRIDDISK command. The PREFIX needs to be meaningful to the user and should be defined such that it will help you identify the location of the grid disk within the cell disk.

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

To read more, order your copy today.

Managing Oracle Fusion Applications
Best Practices for Maximizing the Comprehensive Set of Management Tools and Services
Richard Bingham
Written by a member of the Oracle team designing the support strategy for Oracle Fusion Applications, this is the first and only book to explain how to use a comprehensive range of tools, services, and best practices to manage Oracle Fusion Applications

Available March 2011

TABLE OF CONTENTS

Fusion Applications is Oracle's latest and greatest Enterprise Application Suite, offering organizations the best-of-the-best in terms of functions and features, all designed and delivered using the latest technology architectures and tooling.

As most application managers are starting to learn, managing service-oriented architecture (SOA)–based enterprise applications presents a challenge, because while this architecture simplifies development, promotes code reuse, and provides various other benefits, these applications add a range of new management demands. The modular and discrete nature of the Java Platform (J2EE) together with composite-based applications means the work has become decentralized, and when compared to legacy application technologies, the distributed moving parts are much more difficult to visualize and control. The term *loosely-coupled* isn't usually something an application manager likes to read! So while Fusion Applications is built on and natively uses all the latest technologies and models, there are still many opportunities for effective system management, and this book explains precisely where and how to find these.

In addition to the technology challenges, today's expectation of application management includes much more than just keeping the software running. Today, software management is expected to contribute directly to, and be part of, overall operations management. It should deliver noticeable productivity gains and insights, and it should allow its users to complete their tasks better than ever before. Again, this is a core theme included in this book.

The first two chapters provide a detailed overview of Fusion Applications. Chapter 1 begins with a description of all *functional aspects,* including both what products and features are available as well as how the application can be set up to support all kinds of organizational requirements. Chapter 2 then takes a detailed look at the *technical* components of the Fusion Applications stack, explaining their purpose, general architecture, and how they work together to support the application. This chapter also includes a detailed discussion on the flexibility (aka potential complexity) available, from customization levels, to deployment architectural options, and offers an explanation of how the different configurations can affect application management.

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

Chapter 3 moves on to look at *enterprise application management* itself, comparing it to other more traditional roles such as the System Administrator and the DBA. Taking this further, we decompose the role and present its core parts clearly and openly. It is then explained how, when put together properly, these parts can be used to adopt the right approach that helps reduce problems and promote effective results.

Chapter 4 looks at the *lifecycle management* of enterprise applications. It includes careful analyses of several well-established and related lifecycles and uses this to build a completely new model specifically for the management of Fusion Applications.

Having all these basic elements laid out and ready to use, we turn to look at the specific tools, capabilities, and services available for managing Fusion Applications, all presented in the form of a practical *toolbox*. The subsequent five chapters cover the core use cases for application management and provide many detailed descriptions of a wide variety of utilities that can be used for success. Most of the tools described come as part of Fusion Applications, although the discussion also includes some associated products and services that can bring substantial benefits for the application manager. The toolbox is organized in a logical sequence that makes it both easy to read through, as well as useful as a reference for real-world situations.

In Chapter 11, with the know-how about what to do and when and how to do it, we round off the discussion by looking at how to use this knowledge in an intelligent way that helps ensure the initial and ongoing *health of the application*. This ventures further into adopting an effective approach to application management and includes many recommended practices, based on working with hundreds of real applications managers.

The final chapter of the book discusses the future of Fusion Applications and its management, based on where it's starting from, the foundations it puts in place, and the most likely and important evolutions of which the enterprise application manager needs to be aware.

In addition to the text, an extensive appendix is also provided, offering a detailed reference for further reading about all the tools and topics covered. This completes the book and ensures that it forms a comprehensive and central resource to help you understand and manage Fusion Applications.

Copyright © 2011 by The McGraw-Hill Companies, Inc.
This is prepublication content and is subject to change prior to publication.

To read more, order your copy today.

Also Available from Oracle Press

DATABASE

Oracle Database 11*g* Release 2 High Availablity
Jeese
9780071752084 • Mar-11

Oracle Streams 11*g* Data Replication
Despande
9780071496643 • Dec-10

Oracle Database Administration for Microsoft SQL Server DBAs
Malcher
9780071744317 • Sep-10

Oracle Database 11*g* RMAN Backup & Recovery
Hart
9780071628600 • May-10

Oracle Enterprise Manager 10*g* Grid Control Handbook
De Gruyter
9780071634229 • Apr-10

Oracle Database 11*g* PL/SQL Programming Workbook
McLaughlin
9780071493697 • Feb-10

MySQL Database Usage & Administration
Vaswani
9780071605496 • Oct-09

Oracle Essbase & Oracle OLAP
Schrader
9780071621823 • Oct-09

Applied Oracle Security
Knox
9780071613705 • Oct-09

Oracle Database 11*g* Data Guard Handbook
Carpenter
9780071621113 • Aug-09

Oracle Database 11*g*: A Beginner's Guide
Abramson
9780071604598 • Dec-08

Oracle Database 11*g*: The Complete Reference
Loney
9780071598750 • Dec-08

Oracle Enterprise Manager 10*g* Grid Control Implementation Guide
New
9780071492751 • Nov-08

Oracle Database 11*g* PL/SQL Programming
McLaughlin
9780071494458 • Mar-08

Oracle Database AJAX & PHP Web Application Development
Barney
9780071502771 • Feb-08

MIDDLEWARE & JAVA

Java 7 A Beginner's Guide, Fifth Edition
Schildt
9780071606325 • Mar-11

Java 7 The Complete Reference, Eighth Edition
Schildt
9780071606301 • Feb-11

Java 7 Programming
Sarang
9780071633604 • Feb-11

JavaFX A Beginner's Guide
DiMarzio
9780071742412 • Jan-11

Oracle SOA Suite 11*g* Handbook
Jellema
9780071608978 • Sep-10

Quick Start Guide to Oracle Fusion Development
Ronald
9780071744287 • Sep-10

Designing an IAM Framework with Oracle Identity and Access Management Suite
Scheidel
9780071741378 • Aug-10

Oracle Enterprise Manager 10*g* Grid Control Handbook
De Gruyter
9780071634229 • Apr-10

Oracle Fusion Developer Guide
Nimphius
9780071622547 • Jan-10

Oracle WebCenter 11*g* Handbook
Desbien
9780071629324 • Jan-10

JavaServer Faces 2.0: The Complete Reference
Burns
9780071602389 • Dec-09

Oracle JDeveloper 11*g* Handbook
Mills
9780071621823 • Oct-09

Oracle Essbase & Oracle OLAP
Schrader
9780071613705 • Oct-09

Transforming Infoglut! A Pragmatic Strategy for Oracle Enterprise Content Management
MacMillan
9780071602365 • Dec-08

Oracle Enterprise Manager 10*g* Grid Control Implementation Guide
New
9780071492751 • Nov-08

APPLICATIONS

Oracle CRM on Demand Embedded Analytics
Lairson
9780071745369 • Feb-11

Oracle CRM On Demand Combined Analyses
Lairson
9780071745383 • Nov-10

Oracle CRM On Demand Dashboads
Lairson
9780071745345 • Sep-10

PeopleSoft PeopleTools Tips & Techniques
Marion
9780071664936 • Jul-10

Oracle CRM On Demand Deployment Guide
Saenger
9780071717632 • Jun-10

Oracle Self-Service Applications
Cameron
9780071713306 • May-10

Oracle Enterprise Manager 10g Grid Control Handbook
De Gruyter
9780071634229 • Apr-10

Oracle E-Business Suite Development & Extensibility Handbook
Passi
9780071629423 • Oct-09

Oracle Procure-to-Pay Guide
Cameron
9780071622271 • May-09

Oracle General Ledger Guide
Cameron
9780071622295 • May-09

PeopleSoft Developer's Guide for PeopleTools & PeopleCode
Doolittle
9780071496629 • Dec-08

Oracle Enterprise Manager 10*g* Grid Control Implementation Guide
New
9780071492751 • Nov-08

Oracle CRM On Demand Reporting
Lairson
9780071593045 • Sep-08

CERTIFICATIONS

SCJP Sun Certified Programmer for Java 6 Practice Exams (310-065)
Sierra
9780072260885 • Nov-10

NetBeans IDE Programmer Certified Expert Exam Guide (310-045)
Liguori
9780071738804 • Jul-10

OCA Oracle Database SQL Certified Expert Exam Guide (1Z0-047)
O'Hearn
9780071614214 • Dec-09

OCA/OCP Oracle Database 11*g* All-in-One Exam Guide (1Z0-051, 1Z0-052, 1Z0-053)
Watson
9780071629188 • Sep-09

SCJA Sun Certified Java Associate Study Guide (310-019)
Liguori
9780071490030 • Aug-09

OCP Oracle Database 11*g* Administration II Exam Guide (1Z0-053)
Bryla
9780071597098 • Dec-08

OCA Oracle Database 11*g* SQL Fundamentals Exam Guide (1Z0-051)
Watson
9780071597869 • May-08

OCA Oracle Database 11*g* Administration I Exam Guide (1Z0-052)
Watson
9780071591027 • May-08

SCJP Sun Certified Programmer for Java 6 Study Guide (310-065)
Sierra
9780071591065 • Jun-08

OCP Oracle Database 11*g* New Features for Administrators Exam Guide (1Z0-050)
Alapati
9780071496827 • Apr-08

Visit OraclePressBooks.com for more Oracle Press titles, sample chapters, and author podcasts.

Join the Oracle Press Community at

www.OraclePressBooks.com

Find the latest information on Oracle products and technologies. Get exclusive discounts on Oracle Press books. Interact with expert Oracle Press authors and other Oracle Community members. Read blog posts, download content and multimedia, and so much more. Join today!

Join the Oracle Press Community today and get these benefits:

- **Exclusive members-only discounts and offers**
- **Full access to all the features on the site: sample chapters, free codes and downloads, author blogs, podcasts, videos, and more**
- **Interact with authors and Oracle enthusiasts**
- **Follow your favorite authors and topics and receive updates**
- **Newsletter packed with exclusive offers and discounts, sneak previews, and author podcasts and interviews**

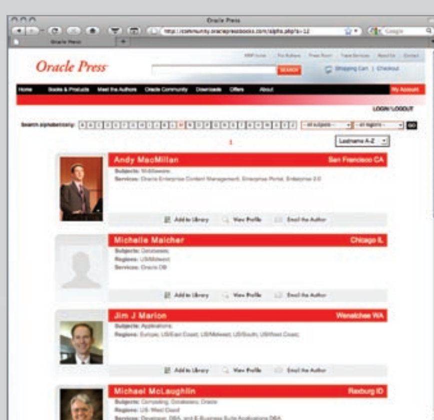

ORIGINAL • AUTHENTIC
Oracle Press
ONLY FROM McGRAW-HILL