

# Oracle Database 10g Release 2 Automatic Storage Management New Features Overview

*An Oracle Technical White Paper  
April 2005*

**ORACLE DATABASE 10G RELEASE 2  
AUTOMATIC STORAGE MANAGEMENT  
NEW FEATURES OVERVIEW**

<b>ORACLE DATABASE 10G RELEASE 2 AUTOMATIC STORAGE MANAGEMENT NEW FEATURES OVERVIEW .....</b>	<b>2</b>
<b>INTRODUCTION.....</b>	<b>3</b>
AUTOMATIC STORAGE MANAGEMENT OVERVIEW .....	3
AUTOMATIC STORAGE MANAGEMENT NEW FEATURES .....	3
<b>DATABASE CONSOLIDATION AND CLUSTERING.....</b>	<b>4</b>
DATABASE STORAGE CONSOLIDATION WITH SINGLE INSTANCE AND RAC .....	4
MULTIPLE DATABASE VERSION SUPPORT .....	4
INDEPENDENT ASM ORACLE HOME .....	5
<b>EASY DATABASE MIGRATION TO ASM.....</b>	<b>6</b>
ENTERPRISE MANAGER ASM MIGRATION UTILITY .....	6
<b>MANAGEMENT FLEXIBILITY AND SIMPLICITY .....</b>	<b>7</b>
COMMAND LINE INTERFACE.....	7
INTERACTIVE COMMAND LINE .....	7
LIST OF COMMANDS:.....	8
NON-INTERACTIVE MODE COMMAND LINE: .....	8
ASM FILE SYSTEM VIEW THROUGH XMLDB VIRTUAL FOLDERS .....	8
ASM VIRTUAL FOLDER ACCESS VIA FTP.....	9
QUERYING THE ASM VIRTUAL FOLDER.....	9
<b>ENHANCEMENTS IN DBMS_FILE_TRANSFER UTILITY .....</b>	<b>10</b>
<b>ENHANCED ASM MANAGEABILITY AND ALERTS.....</b>	<b>10</b>
NEW V\$ASM VIEWS FOR ENHANCED MANAGEABILITY .....	10
ASM WAIT ON REBALANCE PROCESS.....	10
<b>SPACE MANAGEMENT.....</b>	<b>11</b>
USABLE_FREE_SPACE .....	11
REQUIRED_MB_FREE .....	11
OFFLINE_DISKS .....	12
REDUNDANCY_LOWERED.....	12
<b>CONCLUSION.....</b>	<b>12</b>

## INTRODUCTION

This paper introduces the new features of Automatic Storage Management (ASM) in Oracle Database 10g Release 2 (10.2). It describes the ASM features, benefits and improvements in simplification and automation for managing your Oracle database files. It discusses capabilities for database consolidation and storage clustering in Oracle Real Application Clusters (RAC) as well as single instance database environments.

This white paper is targeted at a technical audience mainly comprising:

- Database and system administrators
- Architects
- Consultants
- System engineers
- Project managers

## Automatic Storage Management Overview

Automatic Storage Management is a feature of Oracle Database 10g that provides integrated cluster file system and volume management capabilities at no additional cost. ASM lowers Oracle database storage total cost of ownership and increases storage utilization without compromising performance or availability. With ASM, a fraction of the time is needed to manage your database storage environment and files.

ASM eliminates the need for over provisioning and maximizes storage resource utilization facilitating database consolidation. The ASM self-tuning feature evenly distributes database files across all available storage. It delivers equivalent to raw high performance - sustained over time, with the ease of use of a file system. ASM's intelligent mirroring technology enables triple data protection, even on non-RAID storage arrays, empowering reliable low cost storage deployment.

ASM reduces Oracle Database 10g cost and complexity without compromising performance or availability:

- Simplify and automate storage management
- Increase storage utilization and agility
- Predictably deliver on performance and availability service level agreements

## Automatic Storage Management New Features

Automatic Storage Management feature in Oracle Database 10g Release 2 expands the capabilities of ASM and provides features and enhancements in the following categories:

- Database consolidation and clustering
  - Database storage consolidation with Single Instance and RAC

- Multiple database version support
- Independent ASM Oracle home
- Easy database migration to ASM
  - Enterprise Manager Grid Control ASM migration utility (available with EM Grid Control 10g Release 2)
- Management flexibility and simplicity
  - Command Line Interface (CLI)
  - ASM file system view through XMLDB virtual folders
  - Enhanced ASM manageability and alerts
  - Enhancements in DBMS\_FILE\_TRANSFER utility

## **DATABASE CONSOLIDATION AND CLUSTERING**

### **Database Storage Consolidation With Single Instance and RAC**

In Oracle Database 10g Release 2, Oracle Clusterware does not require an Oracle Real Application Clusters (RAC) license. Oracle Clusterware is now available with ASM and single instance Oracle Database 10g allowing support for a shared, clustered pool of storage for single instance Oracle databases.

Multiple single instance databases can leverage ASM and Oracle Clusterware to economically consolidate multiple islands of databases into a single clustered pool of storage managed by ASM. This feature allows customers to optimize their storage utilization by eliminating wasted over-provisioned storage and save money by reducing their overall footprint of database storage.

RAC and single instance databases could not be managed by the same ASM instance in Oracle Database Release 1. This created challenges in implementing storage Grid architectures and consolidated database solutions.

Oracle Database 10g Release 2 enhances the ASM functionality in a clustered environment allowing one ASM instance per node to manage all database instances in the cluster. Therefore, an ASM instance on a given node can now manage storage for single instance or RAC database instances. This feature relieves the customer from maintaining more than one ASM instance needed to serve all the database types that might exist in the cluster, thus obviating the need for DBAs to manage separate storage pools. This is a storage utilization feature that translates to significant cost savings and simplicity in management.

### **Multiple Database Version Support**

ASM feature in Oracle Database 10g Release 2 transparently supports both older and newer software versions of the database. Both forward and backward

compatibility is maintained between all Oracle Database 10g Release 1 and Release 2. Therefore, any combination of 10.1.u.v and 10.2.x.y for either ASM or database instances interoperate seamlessly without administrative intervention. The only requirement is that a database connecting to a 10.2 ASM instance must be at version 10.1.0.3 or later.

When mixing software versions, ASM functionality reverts to the functionality of the earliest version in use. For example, a 10.1.0.3 database instance working with a 10.2 ASM instance does not exploit new features in ASM 10.2. Conversely, a 10.2 database instance working with a 10.1.0.3 ASM instance, will not exploit any of the new 10.2 features of ASM. V\$ASM\_CLIENT has two new columns that gives you information about the software version number and corresponding database compatibility level. Note V\$ASM\_CLIENT exists on both ASM and database instance, therefore the content of the view will be in respect to the instance from where its queried.

The SOFTWARE\_VERSION column of V\$ASM\_CLIENT gives software version number of the database or ASM instance for the selected diskgroup connection. The COMPATIBLE\_VERSION column gives the compatible setting of the database or ASM instance for the selected diskgroups connection.

### **Independent ASM Oracle Home**

For higher availability and manageability, Oracle recommends that the ASM instance be installed in a separate ORACLE\_HOME (ASM\_HOME) from the database ORACLE\_HOME if multiple databases are running on that node. A distinct ASM\_HOME allows for upgrading and patching the database without affecting ASM. Additionally, a distinct ASM\_HOME also creates an infrastructure to support Database Storage Consolidation. Although a separate ASM\_HOME was possible in Oracle Database 10g Release 1, it requires some manual effort to accomplish this task. In Oracle Database 10g Release 2, Oracle Universal Installer (OUI) and Database Configuration Assistant (DBCA) have been enhanced to allow the user to seamlessly create and install an ASM instance in a separate ORACLE\_HOME.

OUI now has options to:

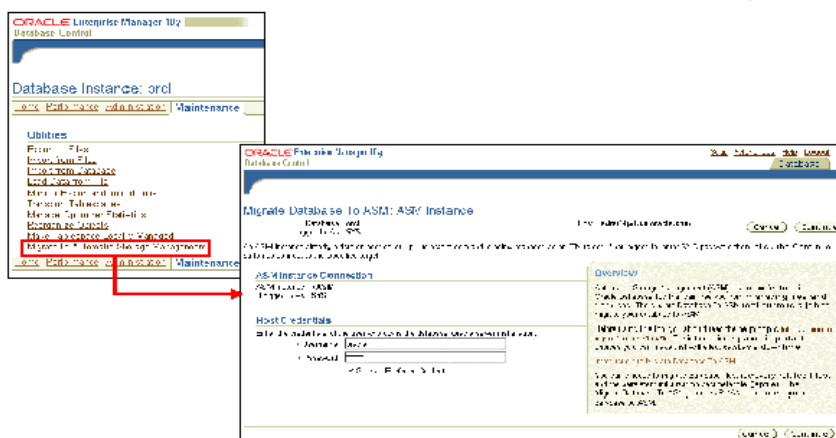
- Install and configure a database that uses ASM for storage management
- Install and configure an ASM instance, without creating a database.
- Install and configure ASM on a system that already has a running database. Subsequently, the DBA can use the EM Migration Utility to migrate the database to ASM.

## EASY DATABASE MIGRATION TO ASM

### Enterprise Manager ASM Migration Utility

The ASM Migration utility is a new feature EM Grid Control 10g Release 2 that provides simplified database migration to ASM based storage. This utility migrates Oracle Database 10g Release 1 and later databases residing on a file system or raw devices to Automatic Storage Management on the same host. The EM Migration utility can be accessed from the Maintenance tab of the Database Instance.

### Migrate Database To ASM Page



The purpose of EM Migration utility is to automate and simplify the database migration process to ASM. The EM Migration utility provides the following benefits:

- Supports migrating all database files, recovery related files, and spfile.
- Provides a wizard walk-through to create migration requests, process and schedule migration requests, as well as report statuses of migration requests.
- Migrates database using the same DBNAME, DBID and SID
- Provides a fully recoverable fallback database, if user decides to go back to non-ASM based storage (assuming that ASM storage and existing storage exist at the same time)
- Monitors the migration process
- Supports online and offline database migration

The EM Migration utility essentially uses Recovery Manager (RMAN) commands to migrate control files, datafiles, and recovery related files from file system to ASM.<sup>1</sup> The migration task is created and submitted as a job via EM job system. The task specification process includes walking through the wizard, which involves real-time validation of user requests. The submitted job includes the following major sub-steps:

- Construct initialization parameters
- Migrate database files with RMAN
- Migrate temp files and online logs with PL/SQL procedures
- Cleanup

## MANAGEMENT FLEXIBILITY AND SIMPLICITY

### Command Line Interface

The ASM Command Line Interface (ASMCMD) utility provides an alternative for managing and monitoring ASM diskgroups.

Previously in Oracle Database 10g Release 1 ASM, SQL\*Plus and EM were the only means to manage and access ASM. ASMCMD utility provides a syntax similar to Unix shell commands, thus simplifying ASM management for those administrators not accustomed to the SQL\*Plus interface or SQL programming.

The ASMCMD interface provides both interactive and non-interactive modes. The interactive mode provides a Unix shell-like environment where the user is prompted to issue the commands listed below. The non-interactive mode executes a single command and exits the utility. The latter is made available for scripting and batch processing purposes.

### Interactive command line

Invoking ASMCMD with no command arguments starts an interactive shell environment. Interactive Mode command line example:

```
% asmcmd
ASMCMD> cd +DATAFILE/ORCL/CONTROLFILE
ASMCMD> ls
EXAMPLE.267.553536489
SYSAUX.257.553536323
SYSTEM.256.553536323
UNDOTBS1.258.553536325
UNDOTBS2.268.553536665
USERS.259.553536325
```

---

<sup>1</sup> For files, such as online logs and temp files, the utility will migrate them by dropping the files, using PL/SQL procedures, and then recreate them in ASM.

### List of commands:

help [command]	Help listing
pwd	Show current dir
cd <dir>	Change directory
find [-t <type>] <dir> <alias>	File search
ls [-lsdrtLaH] [alias]	Directory listing
mkdir <dir1 dir2 . . .>	Directory creation
rm <alias1 alias2 . . .>	File and directory deletion
mkalias <system_alias> <user_alias>	User alias creation
du [-H] [dir]	Disk usage
lsdg [-H] [group]	Diskgroup info
lsct [-H] [group]	Client info
rmalias [-r] <user_alias1 user_alias2 . . .>	User alias deletion

### Non-interactive Mode command line:

```
% asmcmd ls -lH +DATAFILE/ORCL/CONTROLFILE
CONTROLFILE MIRROR FINE MAR 29 15:35 Y
+DATAFILE/ORCL/CONTROLFILE/Current.279.1
CONTROLFILE MIRROR FINE MAR 30 14:24 Y
+DATAFILE/ORCL/CONTROLFILE/Current.280.2
```

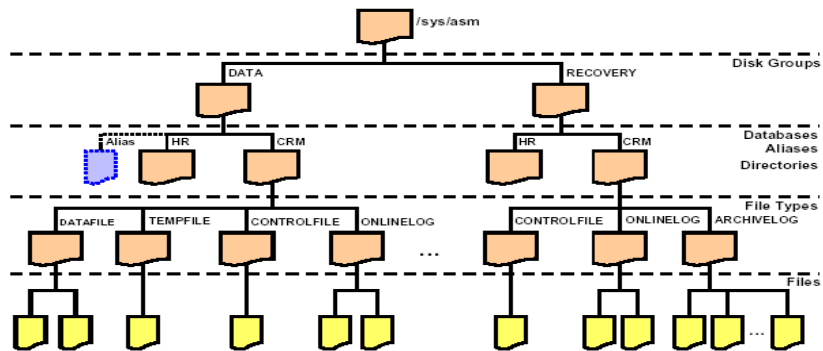
### ASM File System View Through XMLDB Virtual Folders

ASM feature in Oracle Database 10g Release 2 leverages the virtual folder feature in XML DB that provides a mechanism to access and manipulate ASM files and folders via XML DB protocols such as FTP and HTTP/DAV and programmatic APIs. The ASM virtual folder is mounted as /sys/asm within the XML DB hierarchy. The folder is *virtual* in the sense that the ASM folders and files are not actually stored within XML DB. However, any operation on the ASM virtual folder is transparently handled by the underlying ASM component.

The ASM virtual folder is created by default during the installation of XML DB. If the database is not configured to use ASM, then this folder will be empty and no operations will be permitted on it. If ASM is configured, the ASM virtual folder is mounted as /sys/asm within the XML DB hierarchy. Any operation on the ASM virtual folder is transparently handled by the underlying ASM component. The ASM folder contains one sub-folder for every mounted disk group. Each diskgroup folder will contain one sub-folder corresponding to every database name. In addition it may contain additional files and folders corresponding to the aliases that have been created by the administrators.



## ASM Virtual Folder Hierarchy



Note: This is a graphical representation of the ASM GUI folder hierarchy and not the GUI look and feel

### ASM Virtual Folder access via FTP

An FTP tool is also available for copying database files in and out of the ASM diskgroup offering additional flexibility for managing your files. Below is an example of the FTP commands:

```
ftp> open dlsun247 7777
ftp> user sys
ftp> passwd change_on_install
ftp> cd /sys/asm
ftp> ls
ftp> binary
ftp> get t_dbl.f
ftp> put my_dbl.f
ftp> del my_dbl.f
ftp> del myfiles
ftp> quote rm_rf myfiles
ftp> rename myfiles/my_dbl.f myfiles/my_db99.f
```

### Querying the ASM Virtual Folder

The resources in XML DB repository can be queried using the resource and path views. The system-defined views allow users to traverse the XML DB hierarchy as well as search for resources based on XPATH predicates on metadata and content. The resource and path views can be used to traverse the ASM folder and restrict matches based on system and custom metadata predicates.

Example 1: Find ASM files under folder `/sys/asm/dgroupA/myfolder`

```
select * from resource_view
where under_path(res, '/sys/asm/dgroupA/myfolder') = 1;
```

Example 2 : Find ASM files with filename like 'foo%'

```
select * from resource_view
where under_path(res, '/sys/asm') = 1 and
any_path like '%/foo%';
```

## **ENHANCEMENTS IN DBMS\_FILE\_TRANSFER UTILITY**

The DBMS\_FILE\_TRANSFER utility provides a means to copy files between two locations (on the same host or between database servers). ASM leverages this utility to copy files between ASM diskgroups, and is the primary utility used to instantiate an ASM DataGuard database. In Oracle Database 10g Release 2, DBMS\_FILE\_TRANSFER has been enhanced to support all combinations of ASM and non-ASM file transfers as follows:

- ASM to ASM [also in Oracle Database 10g Release 1]
- ASM to OS file [new]
- OS file to ASM [new]
- OS file to OS file [also in Oracle Database 10g Release 1]

These changes now provide DBAs with another method to migrate database files into and out of ASM storage.

## **ENHANCED ASM MANAGEABILITY AND ALERTS**

### **New V\$ASM Views for Enhanced Manageability**

Two new views are provided to minimize the overhead for querying V\$ASM\_DISK and V\$ASM\_DISKGROUP by eliminating the need for expensive disk discovery.

The new V\$ASM\_DISK\_STAT and V\$ASM\_DISKGROUP\_STAT views are identical to V\$ASM\_DISK, V\$ASM\_DISKGROUP, however, the new \_STAT views are polled from memory and therefore do not require disk discovery. Since these new views provide efficient lightweight access, EM can periodically query performance statistics at disk level and aggregate space usage statistics at diskgroup level, without incurring significant processing overhead.

In order to get more accurate real-time statistics, it may be prudent to query the V\$ASM\_DISK and V\$ASM\_DISKGROUP views. However, caution should be exercised when running queries against these views during peak workloads.

### **ASM Wait on Rebalance Process**

An ASM diskgroup rebalance is an asynchronous operation, in that the control is returned immediately to DBA after the operation is started in the background. One can query the status of the ongoing operation from V\$ASM\_OPERATION. However, there are situations when the diskgroup operation needs to be synchronous; i.e., wait until rebalance is completed. ASM 'alter diskgroup' commands, which result in an immediate rebalance, have the option of specifying the option to wait. This allows for accurate scripting that may rely on the space change from a rebalance completing before any subsequent action is taken. For instance, if you add 100GB of storage to a completely full disk group, you won't be able to use all 100GB of storage until the rebalance completes.

ASM feature in Oracle Database 10g Release 2 also introduces the ability to set the ASM initialization parameter ASM\_POWER\_LIMIT to a value of zero. This may be useful in cases where normal or high redundancy is used, and a rebalance, due to a

lost disk, is not desired. Setting `ASM_POWER_LIMIT` to 0 will disable all automatic rebalance operations. However, the DBA can still perform manual rebalance operations when the parameter is set to 0. DBAs can explicitly specify a `POWER` level<sup>2</sup> from 0 through 11 in manual rebalance commands. If `POWER` level is not specified for manual rebalance operations and `ASM_POWER_LIMIT` has been set to 0, a default value of 1 will be used for that particular manual operation. Dynamically setting the `ASM_POWER_LIMIT` parameter to 0 halts any ongoing rebalance operation. An ongoing manual rebalance operation can also be halted by explicitly specifying a power level of 0 in a new invocation of the `ALTER DISKGROUP ... REBALANCE` command. There are now two new `STATE` values possible for `V$ASM_OPERATION`.

- `HALT`- Administrator specified/init param value of "power 0"
- `ERRS`- Rebalance halted in error (e.g. out of space)

Note, in RAC environments, that a row displaying these states will only be produced on the instance currently performing the rebalance operation. However, a rebalance wait clause will still wait irrespective of which node is currently performing the rebalance. Therefore, it is highly recommended to query from `GV$ASM_OPERATION` to get a cluster-wide view of this information.

A new column, `UNBALANCED`, has been added to `V$ASM_DISKGROUP` that will indicate if a diskgroup is in progress or requires a rebalance. For example, if a rebalance is issued with power level 0, then this will indicate that a rebalance operation was pending, and the `UNBALANCED` column will indicate 'Y'. Using this information, the DBA can perform a diskgroup rebalance when necessary.

## SPACE MANAGEMENT

### USABLE\_FREE\_SPACE

ASM feature in Oracle Database 10g Release 2 introduces a new column in `V$ASM_DISKGROUP` called `USABLE_FREE_SPACE` to indicate the amount of free space that can be "safely" utilized taking mirroring into account (ASM in Oracle Database 10g Release 1 did not have a column to take mirroring into account). The column provides a more accurate view of usable space in the diskgroup.

### REQUIRED\_MB\_FREE

Along with `usable_free_space`, a new column has been added to `V$ASM_DISKGROUP` to more accurately indicate the amount of space that is required to be available in a given diskgroup in order to restore redundancy after

---

<sup>2</sup> Power level this is different than `Power_limit`. Power level is specified at the alter diskgroup command level, whereas `power_limit` is defined as instance-wide `init.ora` parameter setting.

one or more disk failures. The amount of space displayed in this column takes mirroring into account.

### **OFFLINE\_DISKS**

A new column called OFFLINE\_DISKS has been added to V\$ASM\_DISKGROUP to display a count of currently offline disks.

### **REDUNDANCY\_LOWERED**

A new column called REDUNDANCY\_LOWERED has been added to V\$ASM\_FILE to indicate whether a file's redundancy is reduced or lower than what is defined. The valid values for this column are 'Y' and 'N', which indicate Yes and No, respectively. Redundancy for a file becomes reduced when one or more data extents in that file are not mirrored at the level specified by the administrator. In case of unprotected files, one or more data extents could be missing altogether. The scenarios that lead to lowering of redundancy are listed below, along with suggested actions that administrators could take to restore redundancy back to the original level (in other words, to restore the value of REDUNDANCY\_LOWERED for the file to 'N', following a rebalance).

#### **Causes of reduced redundancy      Possible solutions to restore redundancy to the original level**

Disk failure	Add another disk of the same or larger size. If one or more extents are missing altogether (because this file was unprotected), it needs to be restored from backup. Also, if disk is being added, wait until rebalance completes. Ensure that the newly added disk is in the appropriate failure group.
Disk running out of space	Add more space (disks) or delete old files and rebalance.
Insufficient number of failure groups	Increase the number of failure groups. At least two failure groups are required for normal redundancy, and at least three failure groups are required for high redundancy.

### **CONCLUSION**

ASM feature in Oracle Database 10g Release 2 extends the automation and simplicity in database storage management required for realizing the Grid Computing vision. The new ASM features further help lower your cost of ownership while improving the overall database quality of service. ASM eliminates the need for 3<sup>rd</sup> file system and volume management software for the database files minimizing cost, improving storage utilization and consolidation that empowers

reliable low cost computing across Oracle Database 10g single instance and RAC environments.



Oracle Database 10g Release 2  
Automatic Storage Management New Features Overview  
April 2005  
Authors: Ara Shakian, Nitin Vengurlekar  
Contributing Authors: Paul Manning

Oracle Corporation  
World Headquarters  
500 Oracle Parkway  
Redwood Shores, CA 94065  
U.S.A.

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200  
[oracle.com](http://oracle.com)

Copyright © 2005, Oracle. All rights reserved.  
This document is provided for information purposes only and the contents hereof are subject to change without notice.  
This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.  
Oracle, JD Edwards, and PeopleSoft are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.