



ORACLE[®]



Automatic Storage Management

Oracle Database 11g Release 2

ORACLE[®]

Table of Contents

- Storage management challenges
- ASM overview
- *Extending ASM to manage general purpose files*
 - *ASM Volume Manager*
 - *ASM Cluster File System*
 - *ACFS availability and scalability*
- *ASM new enhancements for managing Oracle database files*
- *Installation and configuration tools*
- ASM as an integrated solution
- ASM reference customers



Challenges of Managing Data Explosion

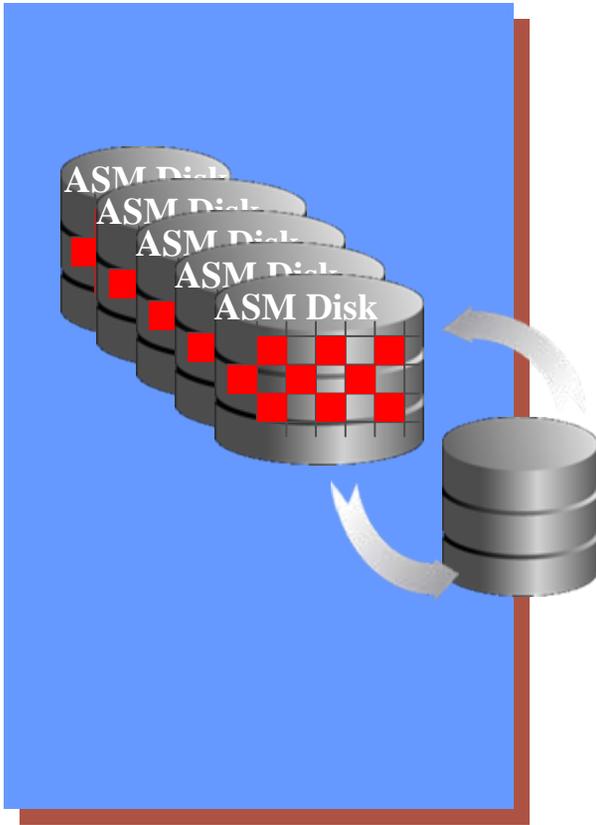


- Complexity in storage and file management
 - Many software layers
 - Lack of a single solution for ALL data
 - Lack of simple and automated tools to manage complexity
- Multi-vendor support logistics and finger pointing
- Complexity of managing performance
 - Eliminating hot spots
- Reducing storage Total Cost of Ownership

Automatic Storage Management (ASM)

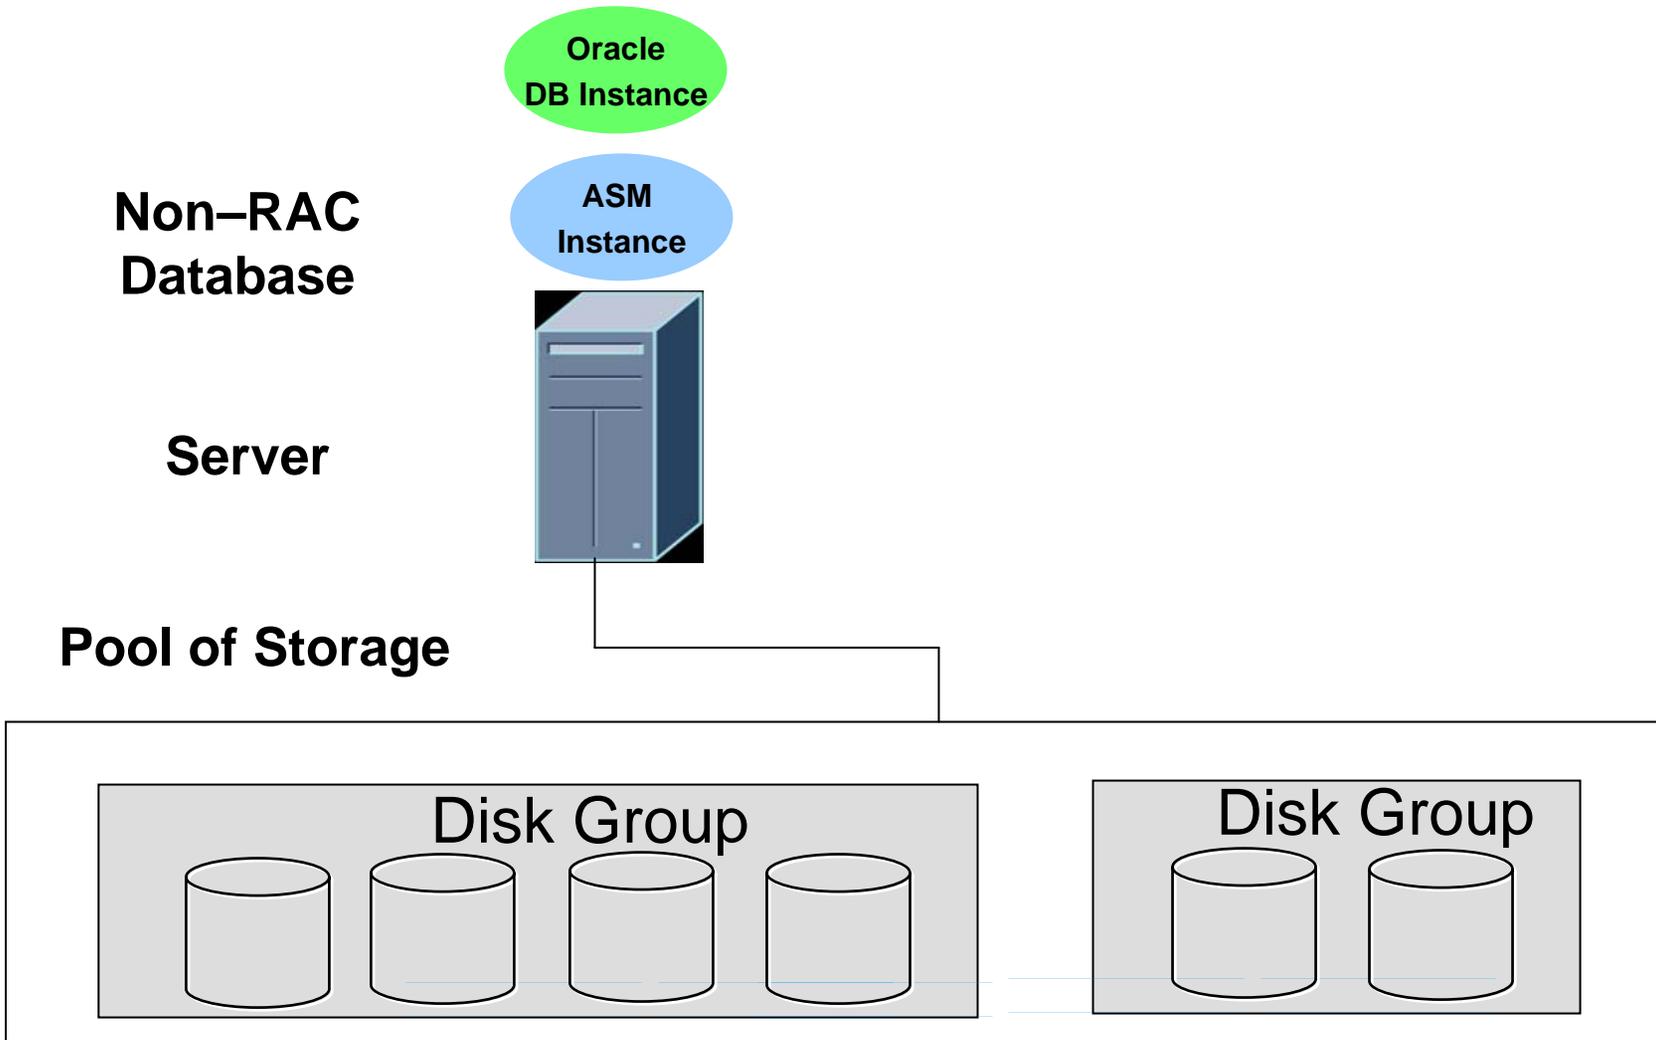
A Volume Manager and File System Integrated in the Database

ASM Disk Group



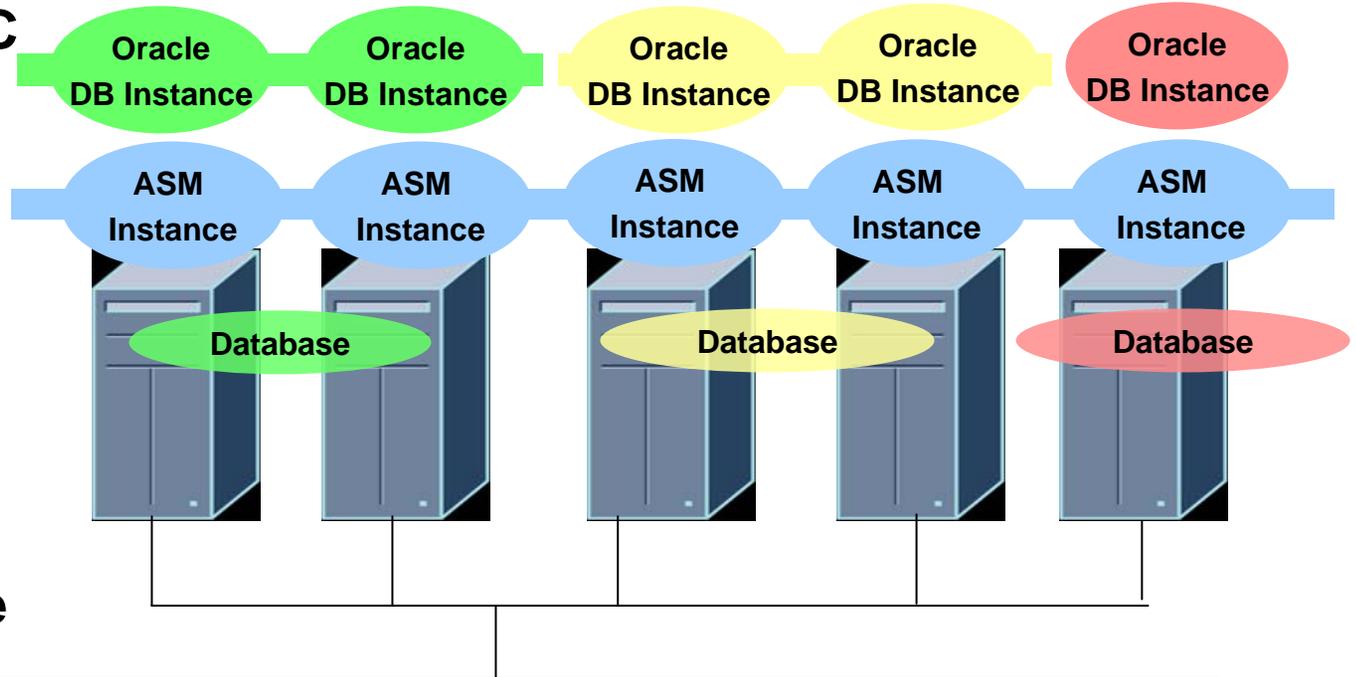
- Easier to manage than file systems
- Reduces storage costs
- Provides best performance
- Stores all data **11.2**

ASM Process Architecture



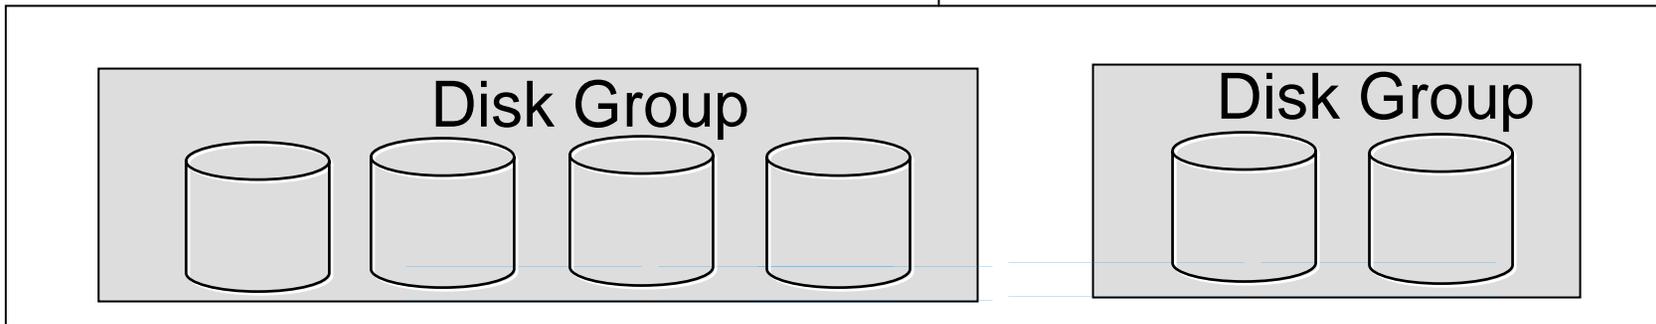
ASM Clusters

RAC or Non-RAC
Databases



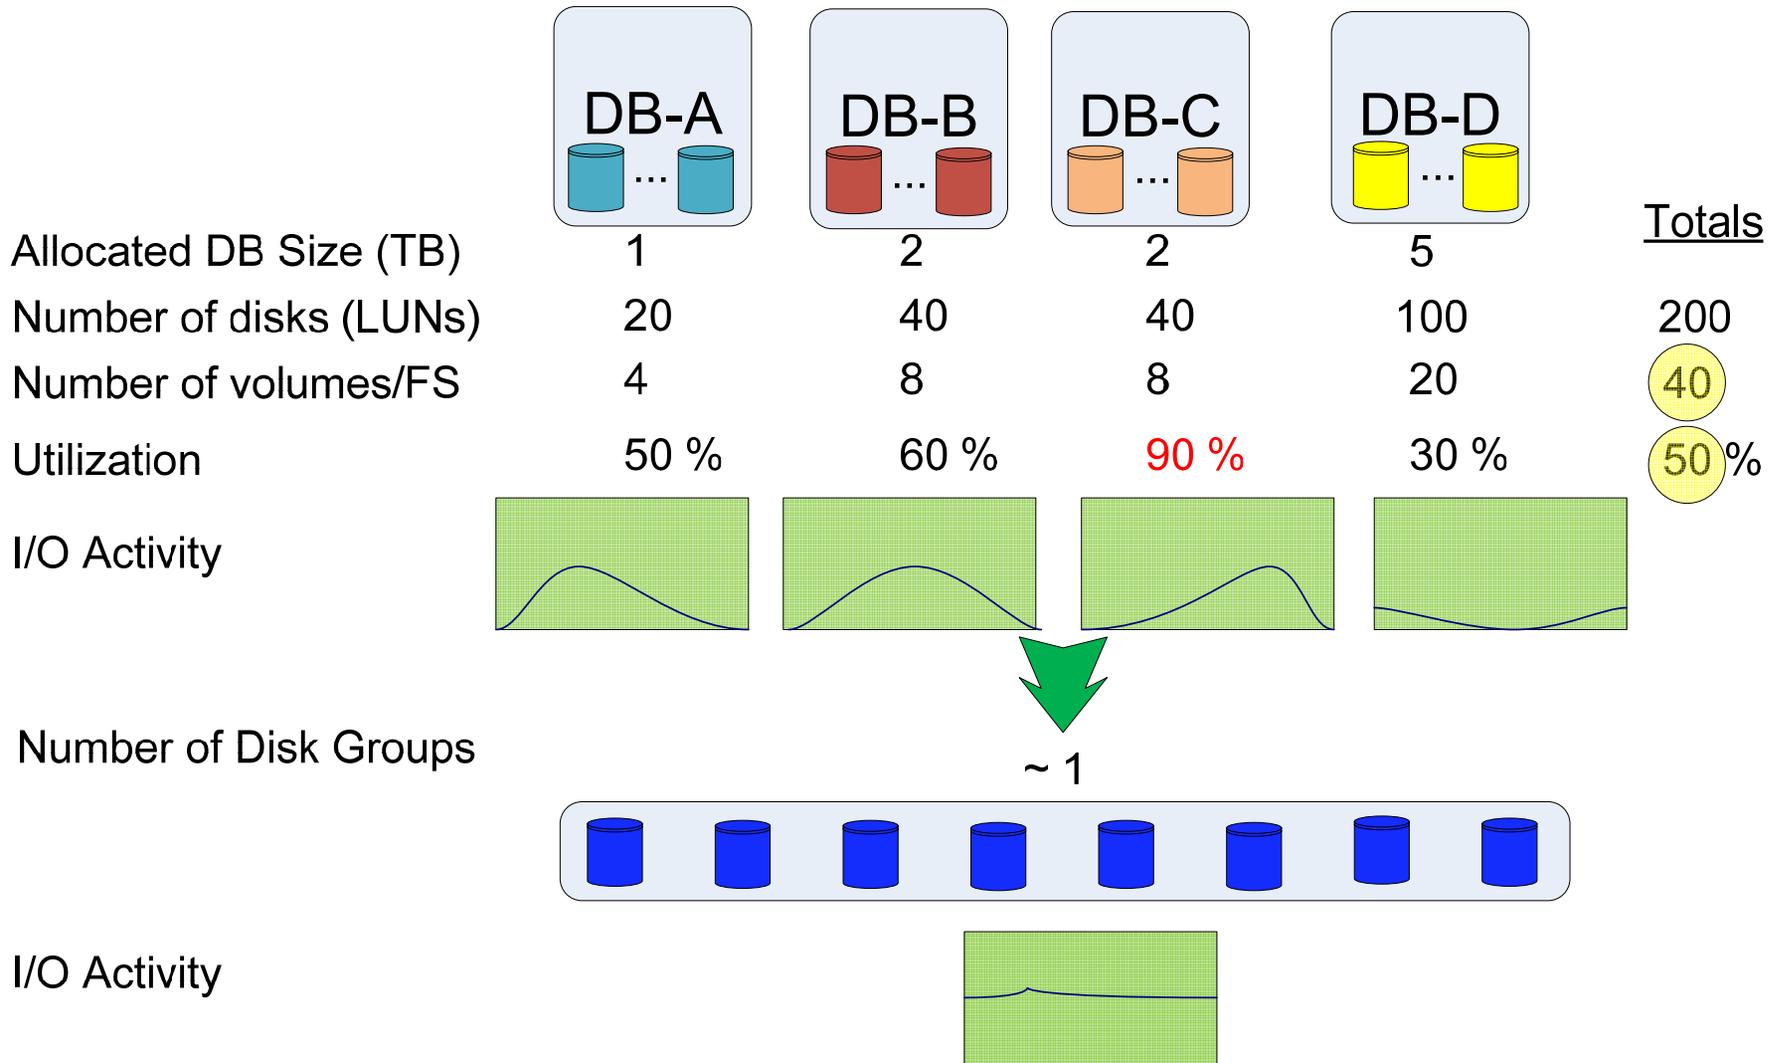
Clustered
Servers

Clustered
Pool of Storage

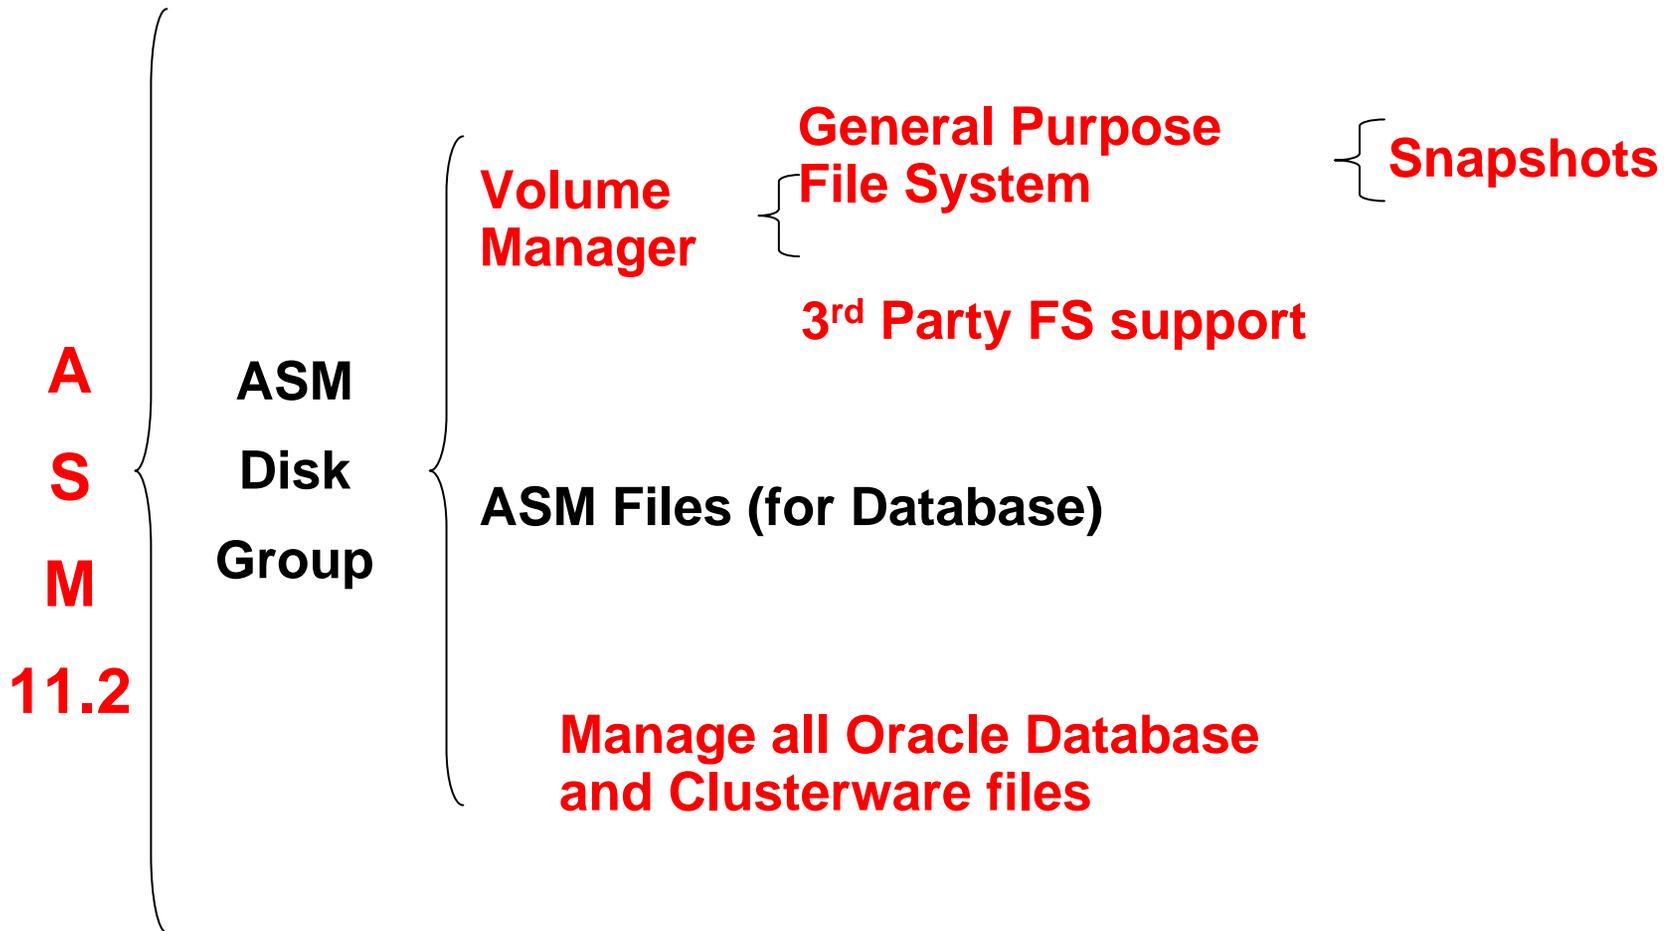


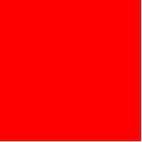
ORACLE

ASM Consolidates Storage & Management



Extending ASM to Support ALL Files

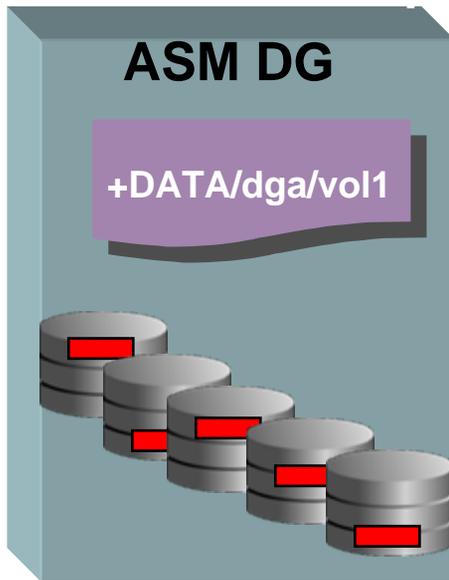




New in
11.2

ASM Dynamic Volume Manager (ADVMM)

Dynamic Volume Manager and Dynamic Volumes



- A new cluster (and single host) volume management service for both Oracle and general volume applications on Linux, Windows, and UNIX platforms
- Loadable kernel driver
- 'Dynamic Volume' is a new ASM file type – 'asmvol'
 - Example: +DATA/dga/vol1
 - Leverage the full set of ASM features
 - Managed using SQL, ASMCMD, EM
- An OS device file is created automatically when a dynamic volume is created
 - Linux: /dev/asm/dga/vol1
 - Windows: asm_dga_vol1

Creating an ASM Dynamic Volume

```
$ sqlplus / as sysasm
SQL> ALTER DISKGROUP data ADD VOLUME volume1 SIZE 500M;

$ asmcmd
ASMCMD> volcreate -d DATA -s 100M volume2

$ ASMCMD> volinfo -a
```

Dynamic Volume SQL*PLUS

Command Examples

- SQL> ALTER DISKGROUP DGROUPA **ADD** VOLUME vol1 SIZE 10G;
RESIZE VOLUME vol1 SIZE 15G;
DROP VOLUME vol1;
ENABLE VOLUME asmvol1;
- SQL> ALTER DISKGROUP ALL **DISABLE** VOLUME ALL;
- SQL> ALTER DISKGROUP DGROUPA **MODIFY** VOLUME asmvol1
USAGE 'acfs';

Extending ASMCMD for Dynamic Volumes

- `asmcmd volcreate` dname volname -s size [-r redundancy] [-w stripe_width] [-c stripe_columns] [-p [hot|cold]] [-i [hot|cold]]
- `asmcmd volresize` dname volname -s newsize [-f]
- `asmcmd voldelete` dname volname
- `asmcmd volenable` [-d dname] [-d dname -n volname]
- `asmcmd voldisable` [-d dname] [-d dname -n volname]
- `asmcmd volset` dname volname [-u usage_message] | [-m mount_path] [-p [hot|cold]] [-i [hot|cold]]
- `asmcmd volinfo` [-d dname] [-n volname]

Managing ADVM Dynamic Volumes

Oracle Enterprise Manager 10g
Database Control
Automatic Storage Management: +ASM_stacg39.us.oracle.com > Logged in As SYS / SYSASM
Disk Group: DEMO

General Performance Templates Files Access Control **Volumes**

Create Enable All Disable All

ASM volumes are typically formatted with Oracle File Systems, which you can use to store non-database files such as executables, Oracle trace files and alert logs, application configuration files, and so on.

Enable Disable Resize Create Oracle File System View Delete

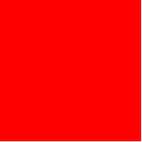
Select All | Select None

Select	Volume	Volume Device	State	Usage	Mount Point	Size (GB)	Disk Group Allocated Space (GB)	Redundancy
<input type="checkbox"/>	DEMO_VOL	/dev/asm/demo/demo_vol	ENABLED	OFS	/ofs_mp/mp_dir1	0.25	0.5	Mirrored
<input type="checkbox"/>	HQ_VOL	/dev/asm/demo/hq_vol	ENABLED	Unknown		0.25	0.75	High
<input type="checkbox"/>	NEDC_VOL	/dev/asm/demo/nedc_vol	ENABLED	OFS	/ofs_mp/mp_dir3	0.25	0.25	Unprotected
<input type="checkbox"/>	SALES_VOL	/dev/asm/demo/sales_vol	ENABLED	OFS	/ofs_mp/mp_dir2	0.25	0.75	High

General Performance Templates Files Access Control Volumes

Database | Setup | Preferences | Help | Logout

Trusted sites



New in
11.2

ASM Cluster File System (ACFS)

Managing non-Oracle Database
General Purpose files

What is ASM Cluster File System (ACFS)

- General purpose scalable file system
 - Journaling, extent based
 - Single node and cluster
 - POSIX, X/OPEN file system solution for UNIX/Linux
 - Windows file system solution for Windows platforms
- Accessible through NAS protocols (NFS, CIFS)
- Leverages ASM technology for volume mgt
- Integrated with Oracle Clusterware for cluster support
- Multi OS platform (Linux and Windows at initial release)
- Integrated with Oracle system mgt tools
 - Oracle installation and configuration
 - Enterprise Manager and ASM Storage mgt tools
 - Native OS File System Management tools

Managing the ACFS File System is Simple

1. Native Linux, UNIX, and Windows OS file system commands and extensions
2. *acfsutil* platform independent commands
3. ASM Configuration Assistant (ASMCA)
4. Enterprise Manager

Creating an ACFS File System

```
$ mkfs -t acfs -b 4k /dev/asm/volume1-62 ; Create fs on ASM volume

$ su - root
# mount -t acfs /dev/asm/volume1-62 /u01/app/oracle/acfsdata/images
# su - oracle_usr

$ acfsutil registry -a /dev/asm/volume1-62 /u01/app/oracle/acfsdata/images

$ mount -t acfs -o all none none ; Mount everything in the registry

$ cd /u01/app/oracle/acfsdata/images
$ cp /home/my_images .
```

Linux-Unix File System APIs

- Standard POSIX and X/OPEN file system APIs and Commands
- access, ar, cat, chattr, chgrp, chmod, chown, cksum, compress, cp, cpio, dd, df, diff, dir, dirname, du, file, find, **fsck**, fuser, grep, gunzip, gzip, link, ln, lockfile, ls, mkdir, mkfifo, **mkfs**, mknod, mktemp, more, **mount**, mv, od, pwd, rcp, rename, rm, rmdir, size, stat, string, sync, tail, tar, unlink, **unmount**, ...

Linux-Unix Extensions

- Create an ACFS file system

```
mkfs [-vf] -t acfs [-b blksz] [-n name ] device [blocks]
```

Example: % mkfs -t acfs /dev/asm/diskgroup/vol1

- Mount an ACFS file system

```
mount [-v] -t acfs [-o options] device dir
```

Example: % mount -t acfs /dev/asm/diskgroup/vol1 /oracle/cluster1/myacfs

- Unmount an ACFS file system

```
umount [-v] device|dir
```

- Check and repair an ACFS file system

```
fsck [-avnf] -t acfs [info] device
```

Example: % fsck -t acfs /dev/asm/diskgroup/vol1

Windows File System APIs

- Windows file management APIs and Commands:
- Assoc, Attrib, Chdir, **acfschkdsk**,
Comp, Copy, Dir, Expand, Fc, Find, Findstr,
Flattemp, **acfsformat**, Freedisk, Mkdir, More,
acfsmountvol, **acfsdismount**, Move,
OpenFiles, Rcp, Rmdir, Replace, Tree, Type,
Xcopy, Process Explorer, ...

Leveraging ASM for Balanced Distribution of Data

All ACFS files get the benefit of ASM extent distribution across disks in a disk group

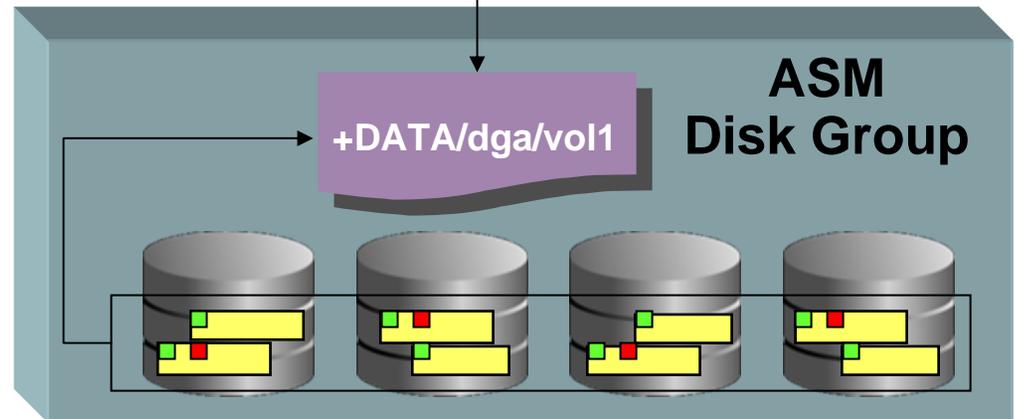
ACFS File System: /usr1

ACFS file A ■

ACFS file B ■

Linux: /dev/asm/dgb/vol1

Win: \\.\asm_dgname_volname



ASM Extent

Leveraging ASM Mirroring and Striping

Files in /usr1 FS are mirrored
leveraging ASM Dynamic Volume
Mirroring

Dynamic Volumes mirroring
attribute is user selectable

- 2-way mirror
- 3-way mirror
- No ASM mirror

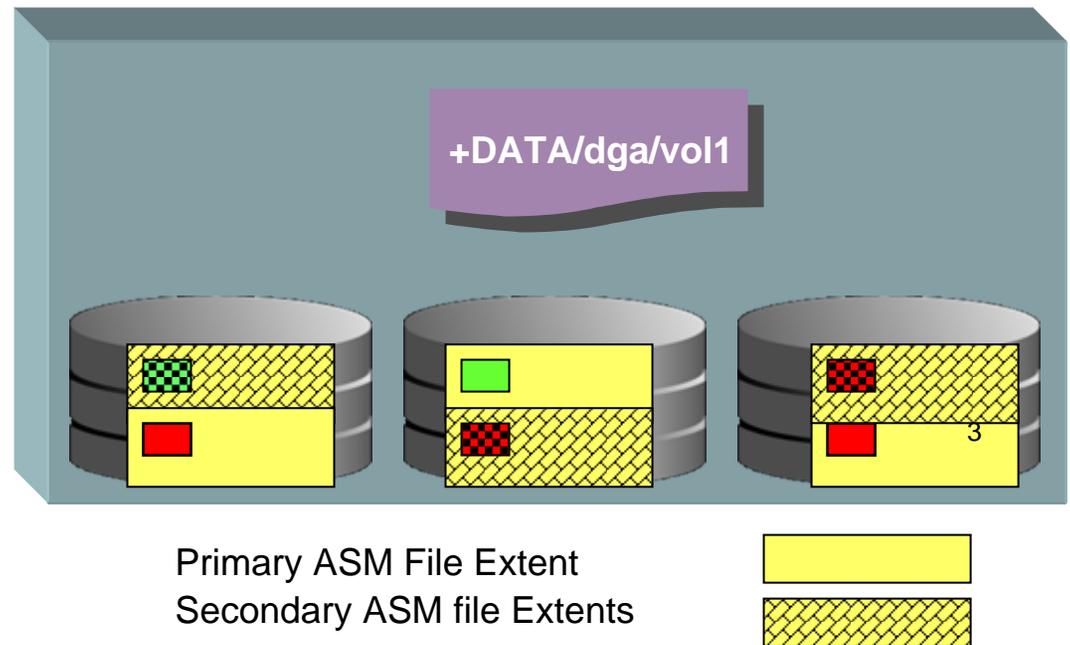
ACFS File System: /usr1

ACFS file A  Mirrored 

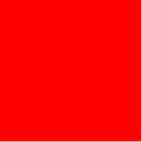
ACFS file B  Mirrored 

Linux: /dev/asm/dgb/vol1

Win: \\.\asm_dgname_volname



ORACLE



ACFS File System Snapshot Data Service

ACFS Read-Only Snapshots

- Dynamic, fast, space efficient, “point in time” copies of ASM file system files
- An enabler for:
 - On-line backups
 - On-line, disk-based, file backup model using snapshots and individual file recoveries
- Up to 63 snapshot images per ASM file system
- Policy based snapshots:
 - Schedule snapshots on an interval basis: every 5 seconds, every 30 minutes, daily, ... with recycling (using EM)
- ACFS CLIs support creation and removal of snapshots
- ACFS Snapshot functions integrated with EM

ACFS Snapshot Command Examples

- Create an ACFS snapshots
 - Syntax: `acfsutil snap create snap_name mountp`
- Delete ACFS snapshot
 - `acfsutil snap delete snap_name mountp`
- **V\$ASM_ACFSSNAPSHOTS**
 - ***FS_NAME*** - *The ACFS mount point*
 - ***DEVICE_NAME*** – *The device name*
 - ***SNAPSHOT_NAME*** – *Snapshot Name which is a valid directory name*
 - ***SPACE_USED_GB*** – *Space used by the snapshot in GB*
 - ***CREATION_TIME*** – *The time at which the snapshot was created*

Managing ACFS using EM - Snapshots

General Snapshots

Summary

A snapshot is a space efficient point-in-time copy of the file system. It can be used as a backup source to recover accidentally deleted or modified files, or for data mining and report applications.

Total Number of Snapshots **3**

Last Snapshot Creation Time **Jun 10, 2008 12:51:28 PM PDT**

Total File System Free Space (GB) **0.1869**

Snapshots

Search

To narrow down the search results or to search for a specific snapshot, enter a search string in the Snapshot name field and click Search. To display all the OFS snapshots for the selected OFS filesystem, clear the search field and click Search.

Name Search

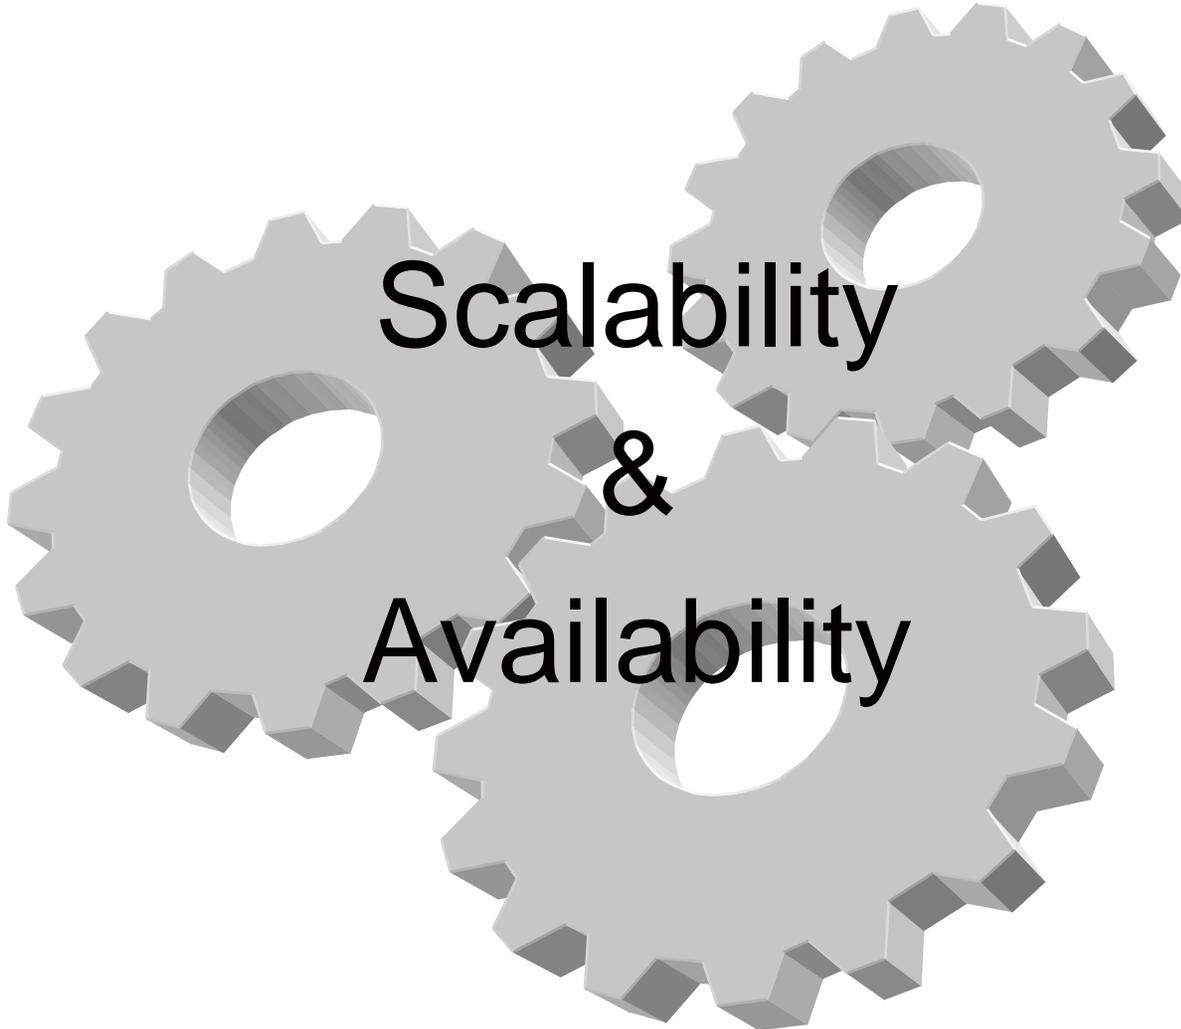
TIP To return exact or case-sensitive matches, double quote the search strings. You can use the wildcard symbol (% , *) in a double-quoted string.

Create

Delete

Select All | Select None

Select	Name	Creation Time
<input type="checkbox"/>	snapshot_20080609_112805	Jun 9, 2008 11:28:14 AM PDT
<input type="checkbox"/>	snapshot_20080610_125103	Jun 10, 2008 12:51:18 PM PDT
<input type="checkbox"/>	snapshot_20080610_125124	Jun 10, 2008 12:51:28 PM PDT

Three interlocking gray gears are arranged in a triangular pattern. The text 'Scalability & Availability' is centered over the gears.

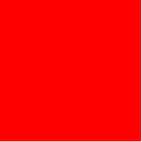
Scalability & Availability

Availability

- ACFS is designed to be ‘always on-line’ during storage configuration changes
 - Leverages all the ASM always on-line attributes
 - ASM mirroring (normal and high redundancy)
- ACFS designed to off-line rather than panic the OS under unusual error conditions (metadata corruption or writes failures)
 - ACFS isolates errors to individual file system or ACFS file, taking the appropriate off-line action

Scalability

- Scales to support large number of cluster nodes
 - Leverages Oracle Clusterware
- Large number of file systems and 100's of thousands of files in a directory
- Scales to very large capacities
 - Small to exabytes for a file system and files
- Fast directory lookup
 - Designed to scale to thousands of files efficiently



New in
11.2

Managing Oracle Database Files

ASM Enhancements

ORACLE

ASM 11g Release 2 New Features

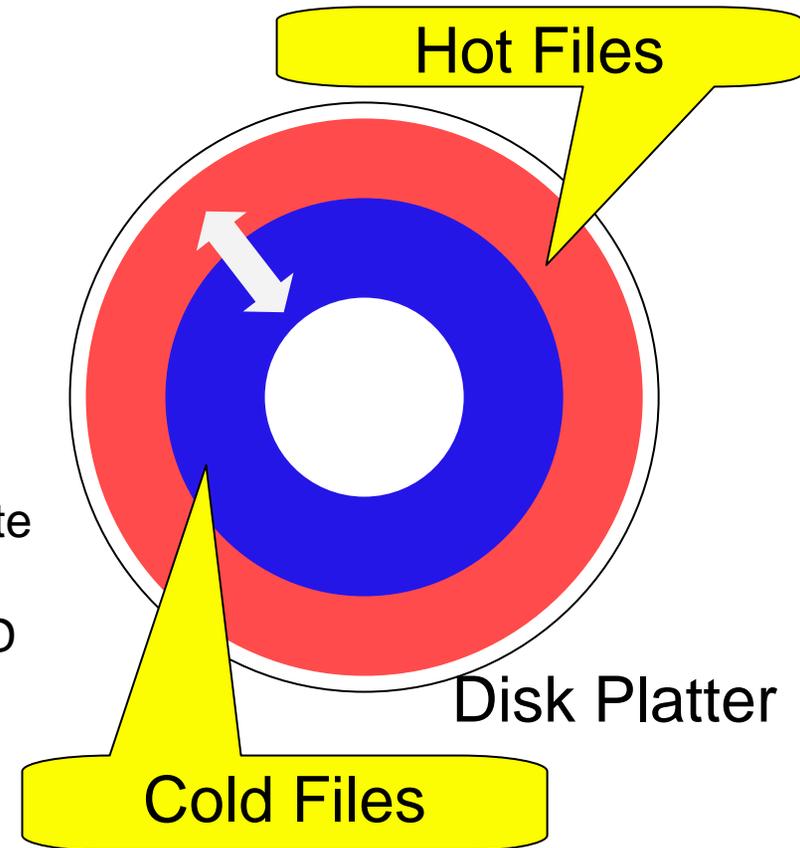
- All files on ASM
 - Oracle Cluster Registry, Voting Disk and SPFILE
- Tunable performance
 - Intelligent Data Placement (IDP)
- Improved management
 - Full featured ASMCMD
 - ASM file access control
 - ASM Disk Group rename
 - Datafile to disk mapping
- Installation & configuration mgt
 - ASM Install & Configuration Assistant (ASMCA)

OCR and Voting Disk on ASM

- OCR is a regular ASM file
 - A new ASM file type
- Voting Disk
 - Stored in selected ASM disks
 - 'crsctl' used to specify a disk group for Voting disks
 - ASM auto creates 1/3/5 Voting Disks based on Ext/Normal/High redundancy on Failure Groups
 - Quorum failure group
- ASM determines redundancy for OCR and Voting Disk based on disk group redundancy level
- ASM SPFILE is also supported on ASM now

Intelligent Data Placement (IDP)

- Policy-based file placement for hot files on high performance regions of disk
- Leverage disk performance regions on disk drives
 - 50% performance difference from outer to inner tracks
- Classify or mark an ASM file to be HOT/COLD
 - Alter diskgroup dname modify file 'xxx' attributes HOT/COLD or based on a template at creation time
 - Rebalance to migrate the 'file' to HOT/COLD IDP region
- IDP regions are dynamic
- New V\$ASMFILE recording IO stats
- The IDP feature better leveraged when ASM disks are whole disks



Complete ASMCMD Functionality

A complete API for the system admin

- Extending ASMCMD to manage:
 - ASM instance (startup/shutdown, init.ora, .)
 - Disk group, disk and failure group (create, mount, add, drop, ...)
 - Attribute (list, set)
 - User/Group (add user, change password, ...)
 - Template (list, add, alter drop)
 - Iostat and df
 - Isof: list files that are open by ASM
Isod: list disks that are open by ASM
 - Dynamic Volume

ASM File Access Control

- Introducing 3 classes of permissions:
 - owner, group, and other
- 3 levels of permissions per class:
 - none, read-only, read-write
- ACL commands apply to OS User names
 - Database users inherit this ACL

- `ALTER DISKGROUP data1 SET ATTRIBUTE 'access_control.enabled' = 'true';`
- `ALTER DISKGROUP data1 SET ATTRIBUTE 'access_control.umask' = '026';`

ASMCMD ACL Commands

Summary of Oracle ASM File Access Control Commands

- **chgrp** Changes the user group of a file or list of files.
- **chmod** Changes permissions of a closed file or list of files.
- **chown** Changes the owner of a closed file or list of files.
- **groups** List the user groups to which a user belongs.
- **grpmod** Adds or removes users from an existing user group.
- **lsgrp** Lists user groups.
- **lsusr** Lists users in a disk group.
- **mkgrp** Creates a new user group.
- **mkusr** Adds a user to disk group.
- **passwd** Changes the password of a user.
- **rmgrp** Deletes a user group.
- **rmusr** Deletes a user from a disk group.

ASM Disk Group Rename

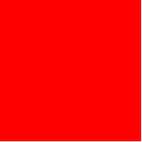
Solution

- *renamedg* tool may be used to rename a disk group
- *renamedg* disk_string is used to discover disks to be renamed in a disk group
- Renaming a disk group is a 2 stage process (for safety)
 1. Generates a config file (each line contains description of change to be made to one disk)
 2. uses conf file to perform change

```
$ renamedg -dgroupname MYOLDDG -newdgroupname MYNEWDG  
$ renamedg -phase one -dgroupname MYOLDDG -newdgroupname  
MYNEWDG -config /tmp/myrendg.conf  
$ renamedg -phase two -config /tmp/myrendg.conf
```

Other Features

- 4k sector size disk I/O support
 - SECTOR_SIZE disk group attributes is specified during DG creation
 - Possible values are: 512b or 4096b (4k)
- Oracle File Mapping support (OFM)
 - Efficient Oracle File Mapping interface
 - View mapping of datafiles to physical devices
 - Dynamic V\$ views used to store mapping info
 - Used with same OFM interfaces



New in
11.2

Installation and Configuration Management

*Designed for
DBAs and System Admins*

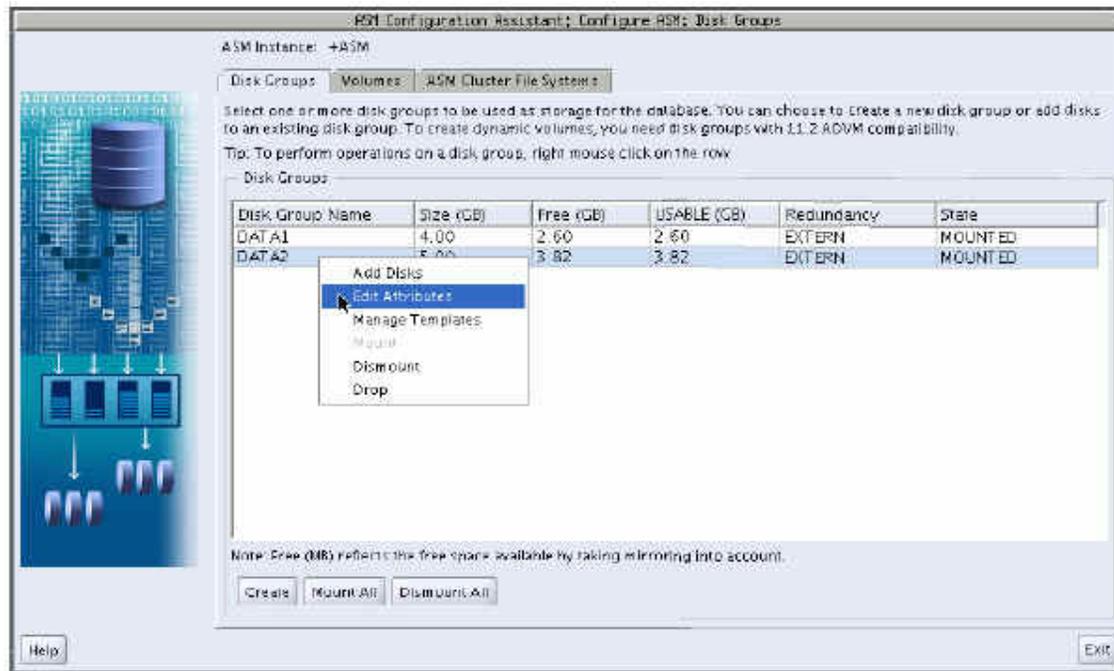
ORACLE®

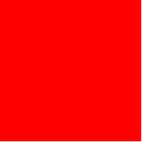
Oracle Grid Infrastructure

- ASM and Oracle Clusterware integrated and installed in a single 'Grid Infrastructure' home
 - Integrated installation, configuration and upgrade
 - Installed using Oracle Universal Installer (OUI)
 - Not a part of database install
 - Option to minimally configure ASM and ACFS
 - Automate creation of Oracle DB Homes on ACFS

ASM Configuration Assistant (ASMCA)

- Manage DB before db install
- One tool to configure ASM, Dynamic Volumes and ACFS
- Upgrade previous versions of ASM to the current version (11g R2)
- Create & manage Dynamic Volumes
- Create and manage ACFS file systems
- Create Oracle DB Homes for rdbms binaries to be installed





In Summary...

One Integrated Solution for ALL Data

Simple

Low Cost

High Performance

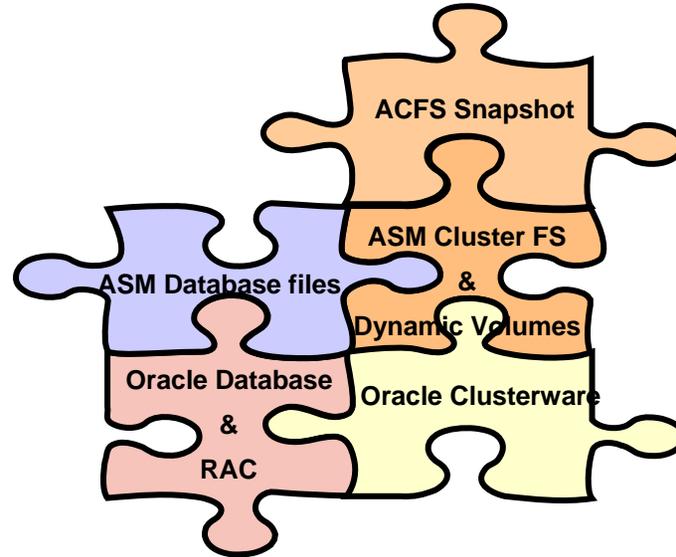
Always On-Line

Scalable

Optimal Utilization

Un-structured
Data

Structured
Data



Cross Platform
Linux, Windows, Solaris,
HP-UX, AIX

One Management
Interface

One Install and
Configure

One Clusterware
Framework

One Vendor for
Support

ORACLE

Value Proposition

- Reduced complexity thru automation
 - No I/O tuning, reshuffling data files
- Simplified solution stack
 - No finger pointing, 3rd party LVM/FS
- Increases utilization and database and application uptime
 - No down-time for storage provisioning or migration
 - Managing storage pool vs. islands
- Single vendor cross-platform integrated solution
 - Bundled and tested as one solution
 - Same cross platform tools to install, configure and manage
- Free

ASM adoption

- De-facto standard for RAC and grid deployments, >65%
- De-facto standard for VLDB deployments
- Large and growing adoption for single instance deployments, >20%
- Thousands of customers using ASM
- One of the most popular features in the database



Some ASM Reference Customers In Production

QUALCOMM®

Overstock.com®
OUTLET SHOPPING

gasNatural

amazon.com®

IKON

Burlington
Not Affiliated with Burlington Industries coat factory



Precision Response Corporation

STARWOOD
HOTELS & RESORTS WORLDWIDE, INC.

publicTRUST

posco

COMIC
RELIEF

DELL™

EA
ELECTRONIC ARTS™

FIDELITY
NATIONAL FINANCIAL™



o/s
Ordnance
Survey®

BPU <> banca
BANCHE POPOLARI UNITE

Deutsche Post World Net

BankWare

OHIO SAVINGS
BANK

ADVANCE AMERICA
CASH ADVANCE CENTERS



THE UNIVERSITY OF
MELBOURNE

Providian

Polk.

HOLMESGLEN
INSTITUTE OF TAFE

The
Chicago
Stock Exchange

OSU OKLAHOMA STATE UNIVERSITY
OKLAHOMA STATE

Cavalier
TELEPHONE &
TV

UNIVERSITETET
I OSLO

ORACLE®

ORACLE®