



# Oracle 10g Automatic Storage Management on HP storage




Klaus Grupe  
HP Oracle Competence Center, EMEA  
Hewlett-Packard



# EMEA HP/Oracle Competence Center




- Located at HP in Boeblingen/Germany, Munich/Germany, Sophia Antipolis/France and Oracle Reading/UK
- Oracle and HP employees in one team
- Founded in spring 1994
- Technology consulting for partners / customers
- Evaluation and tests of new products/features
- Technology transfer to and from US labs
- Technical pre-sales assistance
- Customer specific Oracle Database, Oracle9i AS and Oracle E-Business Suite benchmarks



# Agenda

- Oracle 10g Automatic Storage Management (ASM) Overview
- ASM Architecture
- HP Enterprise Virtual Array (EVA) and ASM
- ASM set up and Configuration on Linux
- ASM and EVA Best Practice
- ASM Administration and Monitoring
- Summary

# Where DBAs Spend Their Time




Source: IOUG 2001 DBA Survey

# DBA Storage Challenges

- Databases are rapidly growing in size
  - DBAs have to manage thousands of data files
- Decreasing windows of scheduled downtime
- Storage resources involve many organizations
  - Database Admin
  - Systems Admin
  - Storage Admin
  - Network Admin
- Management of database storage is increasing in complexity and cost

# Automatic Storage Management (ASM)



Automatic Storage Management


- Portable, high performance cluster file system
- Data is spread across disks to balance load
- Integrated mirroring across disks
- Removes need for third-party volume manager and file system

# Oracle Data stored in ASM


- Data files
- Redo logs
- Spfile
- Datapump dumpsets
- Temp files
- Recovery files
- .....
- No user files

# The Operational Stack

## TODAY



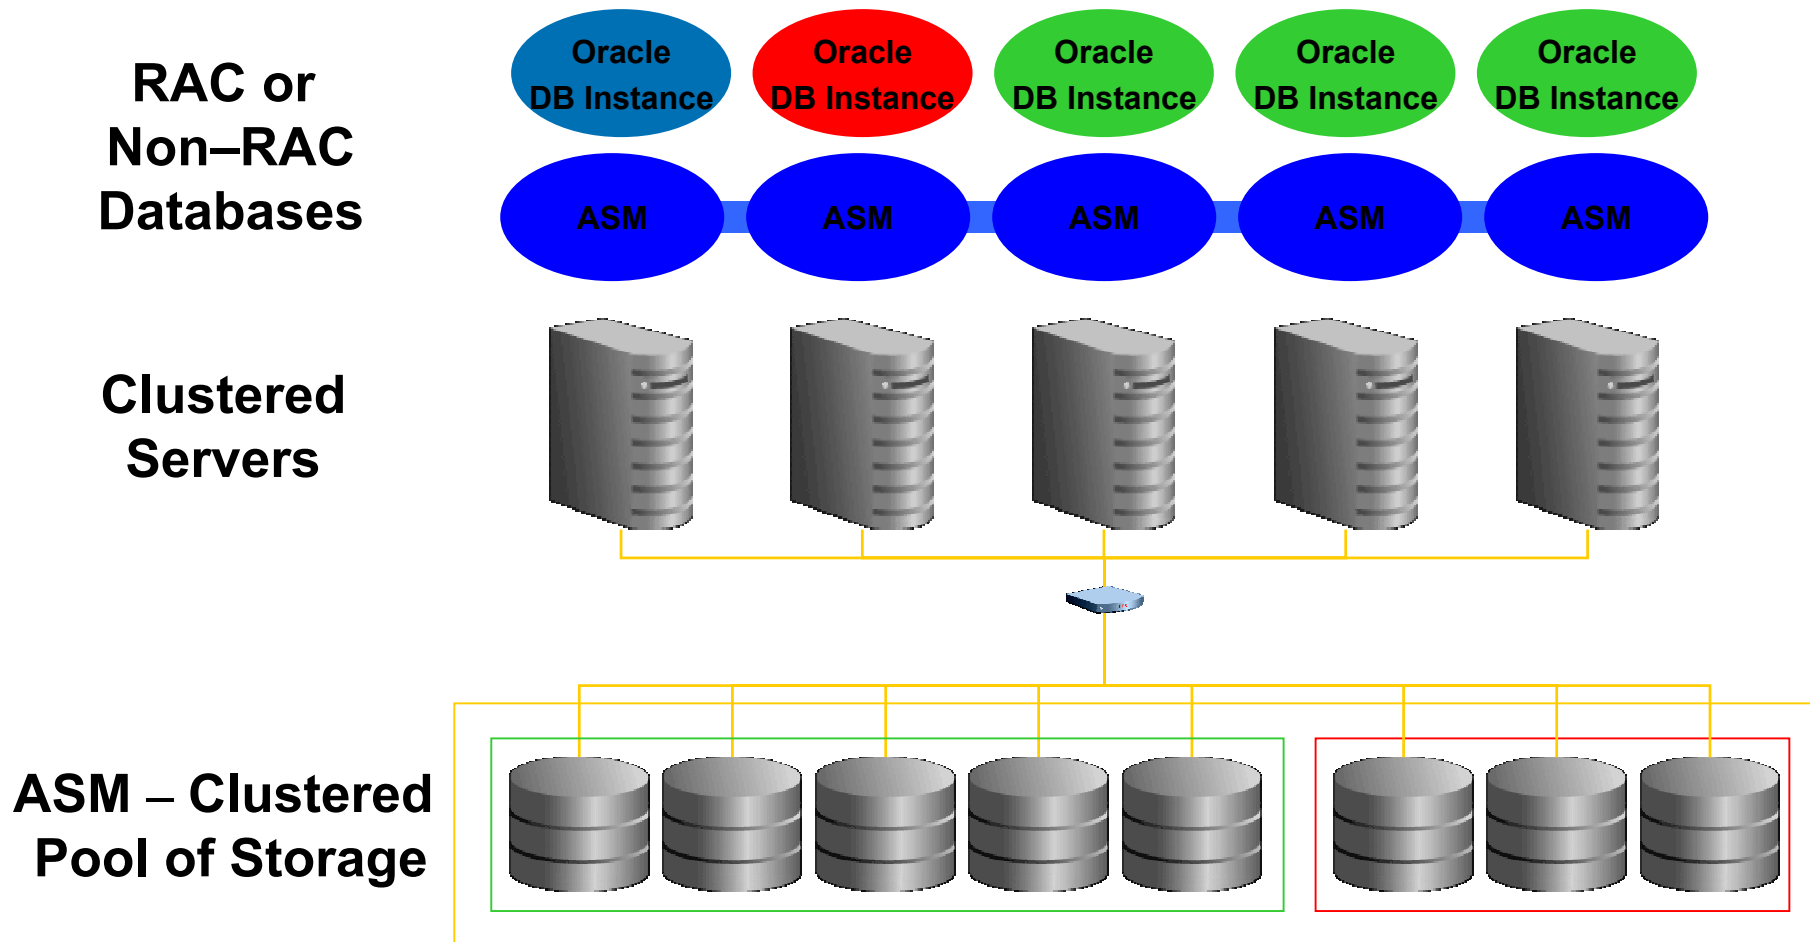
## ASM



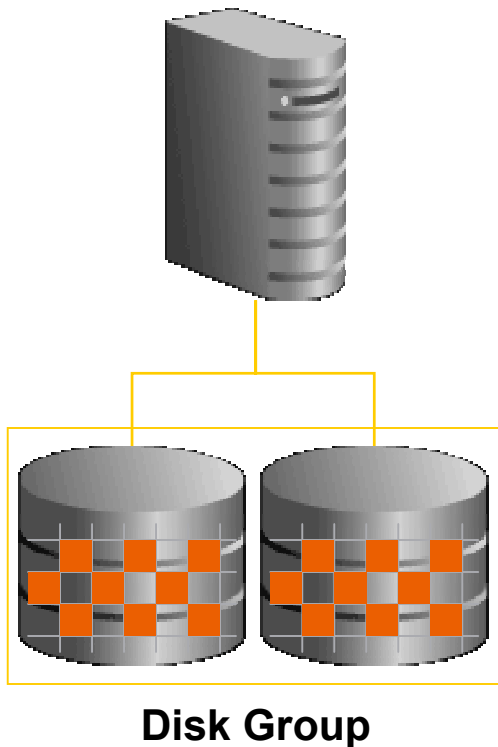
**“The best way to lower mgmt costs is to remove complexity”**



# ASM Enables Consolidated Clustered Storage



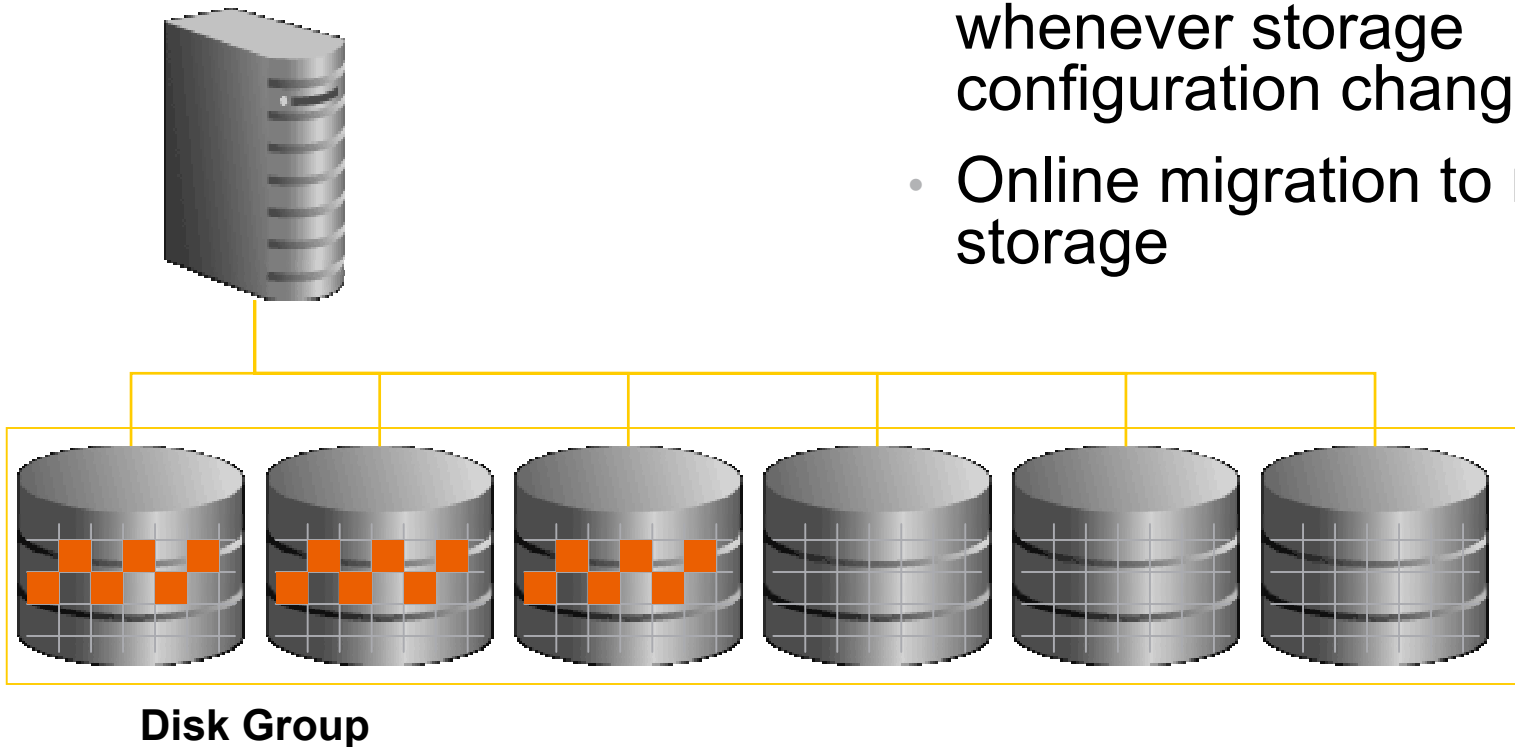
# ASM Disk Groups



- A pool of disks managed as a logical unit
- Partitions total disk space into uniform sized megabyte units
- ASM spreads each file evenly across all disks in a disk group
- Coarse or fine grain striping based on file type
- Disk groups integrated with Oracle Managed Files

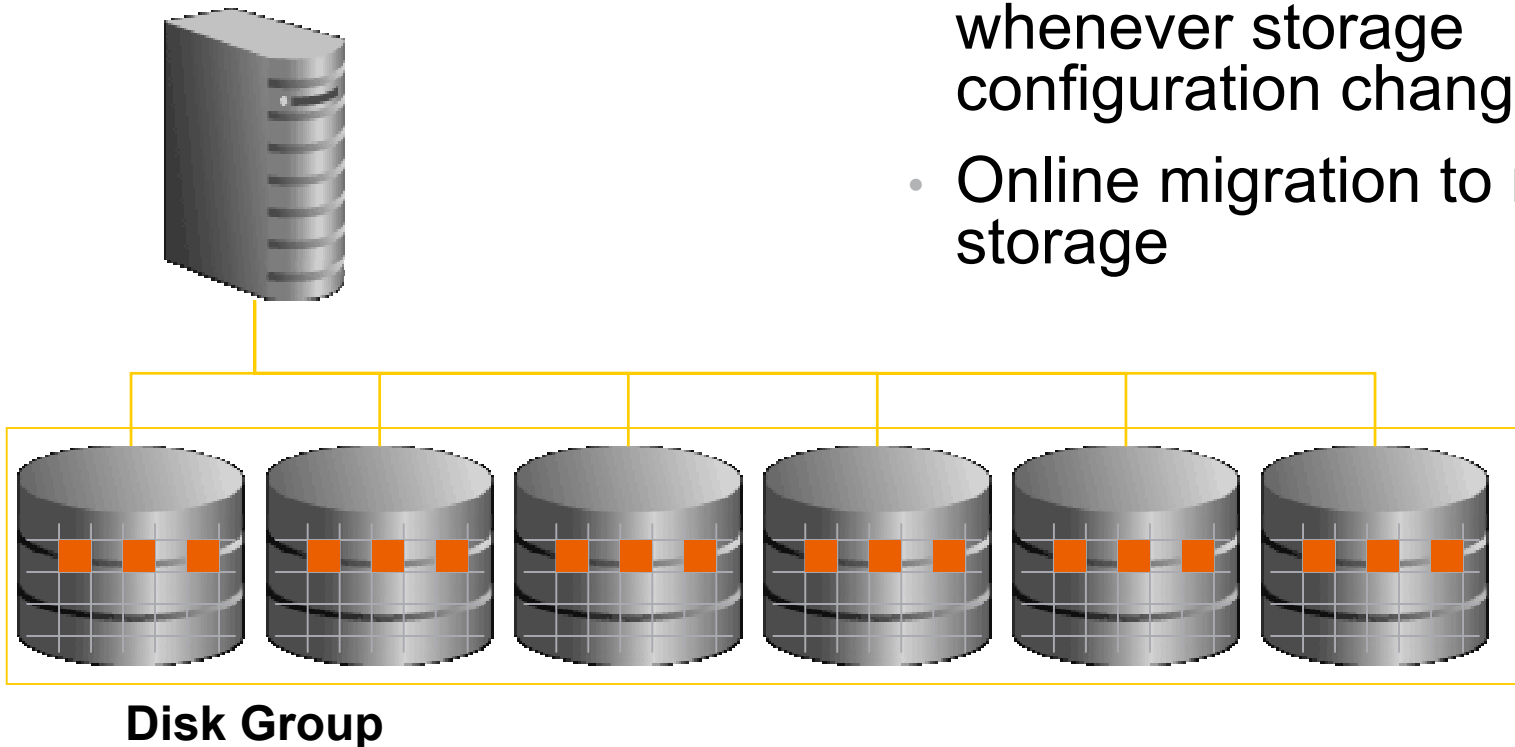
# ASM Dynamic Rebalancing

- Automatic online rebalance whenever storage configuration changes
- Online migration to new storage




# ASM Dynamic Rebalancing

- Automatic online rebalance whenever storage configuration changes
- Online migration to new storage



# ASM Dynamic Rebalancing

- Automatic online rebalance whenever storage configuration changes
- Online migration to new storage




**Disk Group**

# ASM Dynamic Rebalancing



- Automatic online rebalance whenever storage configuration changes
- Online migration to new storage



**Disk Group**

# ASM Mirroring

- 3 choices for disk group redundancy
  - External: defers to hardware mirroring
  - Normal: 2-way mirroring
  - High: 3-way mirroring

# ASM - Integrated with Key Storage Hardware



- ASM-Lib is an Open Interface extending ASM
- Vendor library for disk discovery and I/O
- Efficient I/O interface
- Data description allows end-to-end validation (HARD Initiative)
- Several participating storage vendors and others are in the process of joining



# ASM Integration with Storage Partners

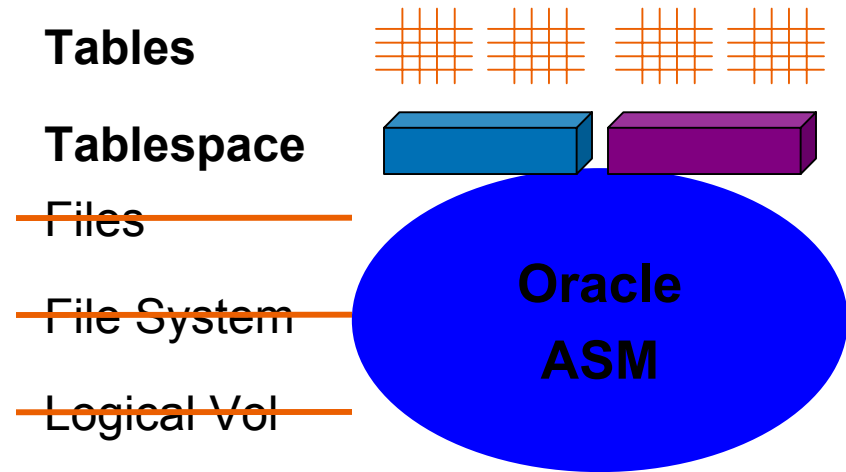
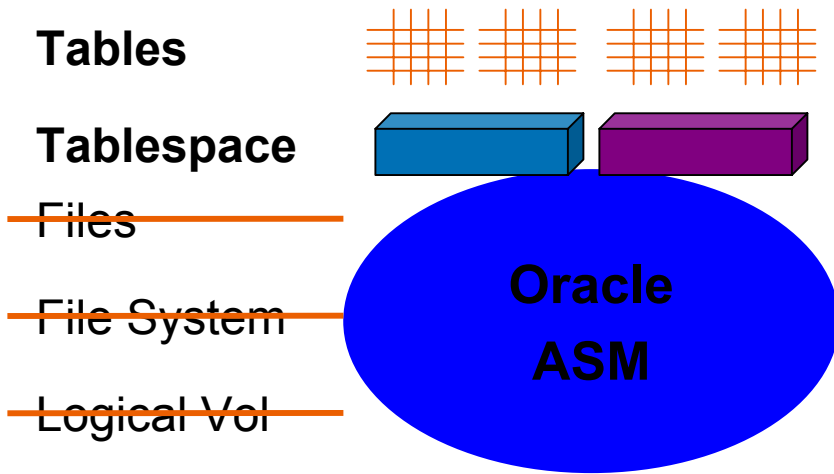


- Developing joint best practice white papers
- Proof of concepts and performance testing
- Extending API definition and interaction
- Interoperability testing
  - Multi-path / channel fail over
  - Snap shot / third mirrors
  - Remote mirroring



# The Operational Stack a closer Look

## ASM

## ASM



Disk Group



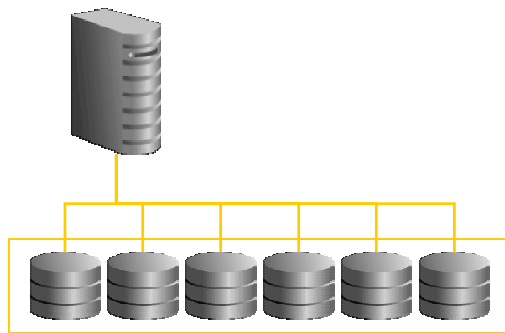
# Storage Virtualization

- Host based
  - Multiple storage systems
  - Multiple RAID disk controller
  - Multiple nodes
- Raid Controller based
  - Very efficient
  - Off loads host
  - Several reconfigurations are transparent for host

# Storage Reconfiguration

- ASM


- Add or remove disks from ASM disk group dynamically
- Rebalancing host based
- Takes additional host resources
- ASM rebalancing possible over multiple controllers



Disk Group

- EVA


- Add or remove disks from EVA disk group dynamically
- Rebalancing controller based
- Small number of ASM disks to manage
- Raid5 available



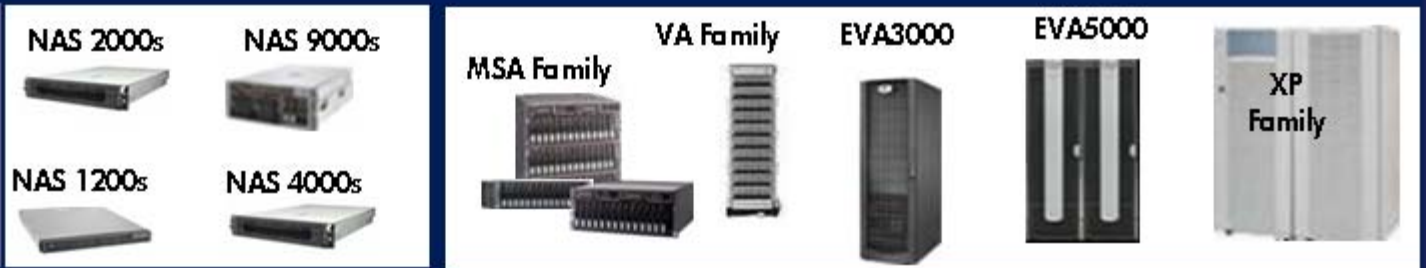
# HP StorageWorks product portfolio



## Infrastructure



## Storage arrays & NAS



## Tape & Optical storage



## Software

### Storage area management



- Device management
- Provisioning

### Data management

- Protect & Recover
- Media lifecycle Management



### Availability management


- Data replication and migration
- Multi-pathing/failover



### Storage virtualization

- Heterogeneous storage services
- Enterprise scale

# Virtualization on the EVA



Physical disks are pooled in **Disk Groups**

**Virtual Disks** are carved out of a Disk Group

A Virtual Disk is presented to a host as a **LUN**

# Disk Groups

- Controller chooses the best drives for Disk Group automatically

⇒ Following several rules to guarantee data accessibility in case of multiple failures

- Always trying to distribute the disks vertically over the shelves
- Redundant Storage Sets are distributed vertically as well
- VRaid 1 Mirror Sets lie on different shelves



# Virtual Disk

A virtual disk is a slice of storage space from a disk group that is virtualized across all mechanisms in the disk group. Can be VRAID 0, 1 or 5.

- One virtual disk can be presented to multiple hosts.
- The maximum size of a virtual disk is 2047 GB. (Minimum is 1GB)
- Dynamic online expansion of Virtual Disk  
Not supported by some Operating Systems






# EVA Virtualization Effects

- It's not traditional storage!
  - Data location is not static
    - Data movement occurs when physical capacity changes
    - Back-end disk access pattern may not be correlated with host access pattern
  - All LUNs are striped across all disks
    - No need for additional (host) striping
    - All LUNs (and database tables) are load balanced
  - Redundancy data is distributed across all disks
    - No need for dedicated “spares”

# ASM Sample Setup 1




- Direct attached SCSI disks
- E.g. MSA 30
- Redundancy: normal




**Disk Group**

# ASM Sample Setup 2



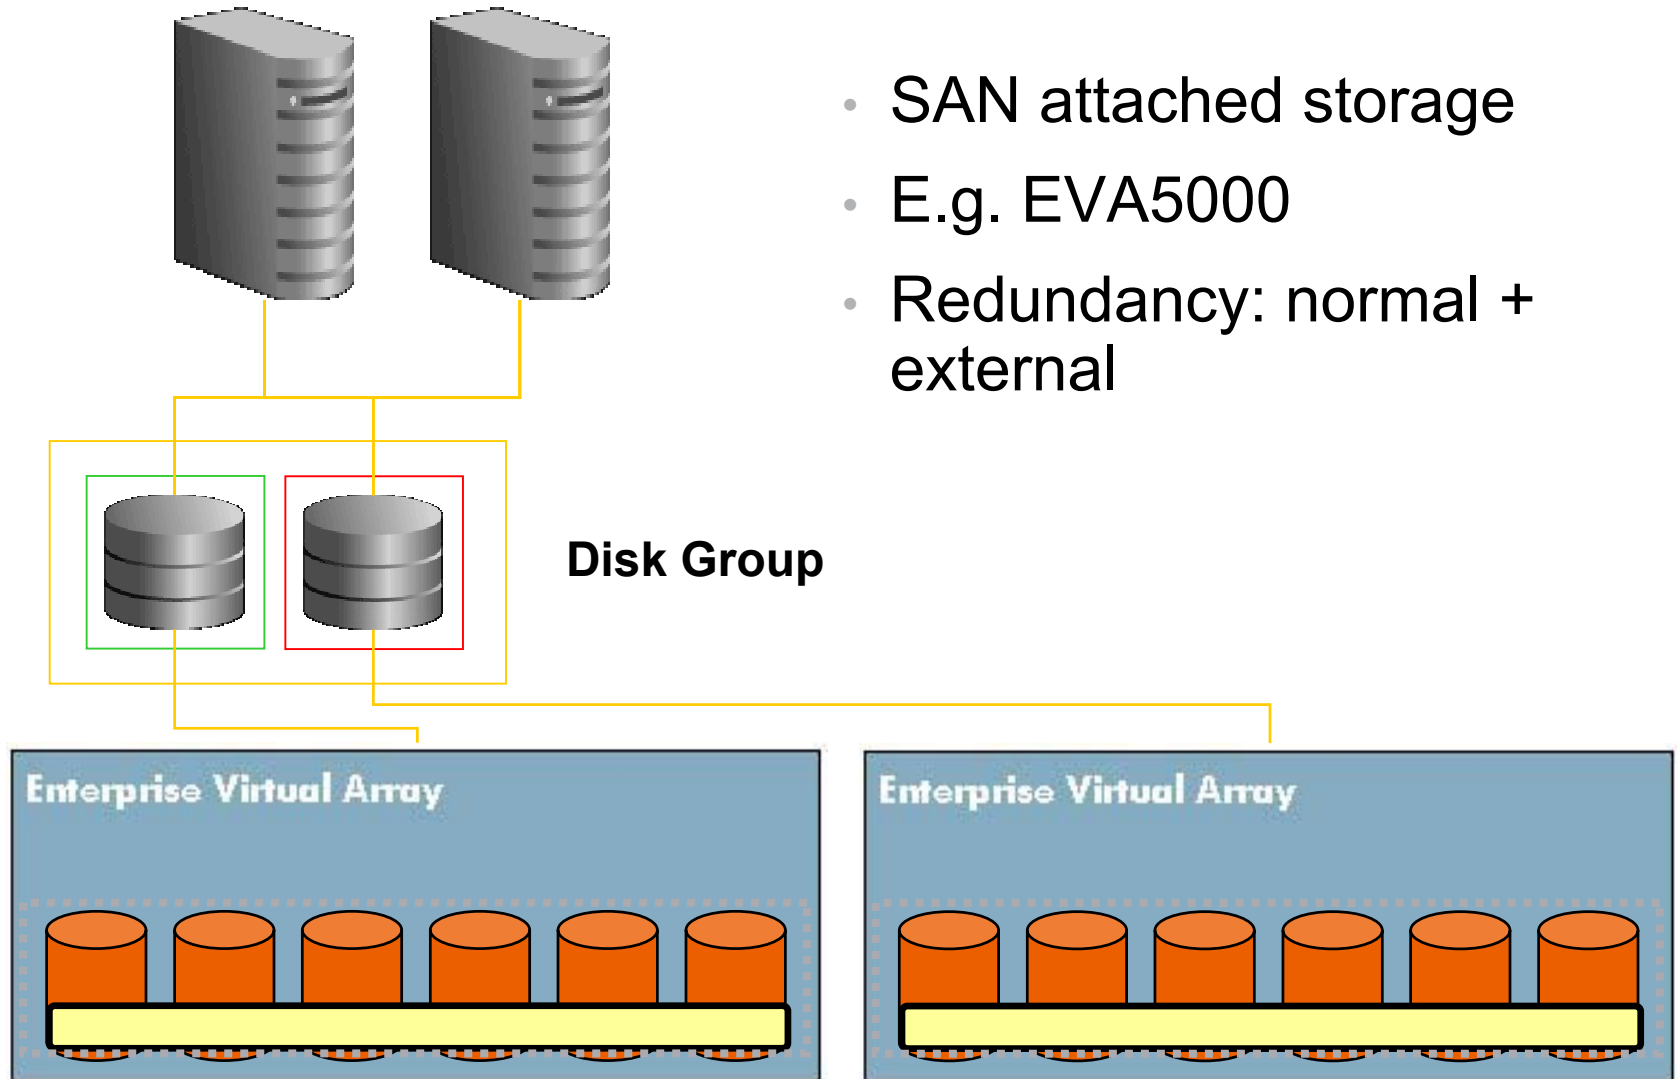
- SAN attached storage
- E.g. EVA3000
- Redundancy: external

# ASM Sample Setup 3



- SAN attached storage
- E.g. EVA3000
- Redundancy: external
- 2 disk groups  
VRaid1 VRaid5

# ASM Sample Setup 4



# ASM set up and Configuration

- Define ASM disks
  - Raw devices or ASM Library
- Create an ASM instance
- Define ASM disk groups
- Create the database
- Multiple databases are possible
  - Decide about ASM disk group design
  - Impact on manageability

# ORACLE ASM Library for Linux

- Oracle recommends to use ASMLib
- Provides native async IO
- Provides disk recovery
- Future enhancements under discussion
- Single interface to define ASM disks
- No raw device handling required
- Reconfiguration of ASM disks more dynamic
- DBCA restriction in Oracle 10g R1(see OTN)

# ASM Library set up

- oracleasm utility
  - Install rpm packages (available on OTN)
  - Configure oracleasm
    - /etc/init.d/configure
  - Create ASM disks
    - /etc/init.d/oracleasm createdisk /dev/sdb
  - Additional oracleasm commands are
    - listdisks, scandisks, querydisk, deletedisk
- Disk recovery
  - asm\_diskstring (ASM parameter)
  - disk recovery path (terminology used by DBCA)
    - ORCL:\*



# ASM raw device set up on Linux

- Raw device bindings
  - /etc/sysconfig/rawdevices
    - /dev/raw/raw8 /dev/sda8
    - /dev/raw/raw9 /dev/sda9
- Make devices available to Oracle
  - chown oracle:dba
- Disk discovery
  - asm\_diskstring (ASM parameter)
  - disk recovery path (terminology used by DBCA)
    - /dev/raw/\*

# ASM Library set up

```

c:\ Telnet saptux1
[root@saptux1 /]#
[root@saptux1 /]#
[root@saptux1 /]#
[root@saptux1 /]#
[root@saptux1 /]#
[root@saptux1 /]#
[root@saptux1 /]#
[root@saptux1 /]#
[root@saptux1 /]#
[root@saptux1 /]#
[root@saptux1 /]#
[root@saptux1 /]# oracleasm
Usage: /etc/init.d/oracleasm {start|stop|restart|link|enable|disable|configure|c
reatedisk|deletedisk|querydisk|listdisks|scandisks|status}
[root@saptux1 /]# oracleasm listdisks
D
[root@saptux1 /]# oracleasm querydisk D
ASM disk "D" is labeled for ASM disk ""
[root@saptux1 /]# oracleasm querydisk /dev/sdb
Disk "/dev/sdb" is marked an ASM disk
[root@saptux1 /]# oracleasm querydisk /dev/sdc
Disk "/dev/sdc" is marked an ASM disk
[root@saptux1 /]# oracleasm status
Checking if ASM is loaded:  OK  ]
Checking if /dev/oracleasm is mounted:  OK  ]
[root@saptux1 /]#

```

# ASM Library and RAC

- ASMLib local file system `/dev/oracleasm`
- Can be checked with `df -ha`
- Stores ASM disk configuration
  - `/dev/oracleasm/disks/DISK1`
- Stored on every node in a RAC environment
- Create an ASM disk only on one node in RAC
  - Writes header on shared disk
- Execute `oracleasm scandisks` on all other nodes
  - Makes local ASM configuration available on other nodes

# ASM set up and Configuration Tools

- oracleasm (ASM Library)
- Database Configuration Assistant (DBCA)
- Enterprise Manager (EM)
- sql commands



# DBCA Sample Session



**X Database Configuration Assistant, Step 7 of 16 : Storage Options**


Select the storage mechanism you would like to use for the database.

Cluster File System  
Use cluster file system for database storage.

Automatic Storage Management (ASM)  
Automatic Storage Management simplifies database storage administration and optimizes database layout for I/O performance. To use this option you must either specify a set of disks to create an ASM disk group or specify an existing ASM disk group.

Raw Devices  
Raw partitions or volumes can provide the required shared storage for Real Application Clusters (RAC) databases if you do not use Automatic Storage Management and a Cluster File System is not available. You need to have created one raw device for each datafile, control file, and log file you are planning to create in the database.

Specify Raw Devices Mapping File




# Oracle Storage Options on Linux

- Single node
  - raw devices, ext2, ext3, reiserfs, ASM
- RAC
  - raw devices, ocfs, ASM

**X Database Configuration Assistant, Step 8 of 16 : Create ASM Instance**

In order to use Automatic Storage Management(ASM), you need to have an ASM instance running on your machine. There are no ASM instances running on this machine. Use this page to specify parameters for a new ASM instance which will be created when you click Next.



The default settings for creating an ASM instance work for most installations. If you would like to make changes to the defaults, use the **ASM Parameters...** button.

**ASM Parameters...**

The new ASM instance has its own **SYS** user for remote management. Specify the password for that user.

SYS password:

Confirm SYS password:

Choose the type of parameter file that you would like to use for the new ASM instance.

Create initialization parameter file (IFILE)

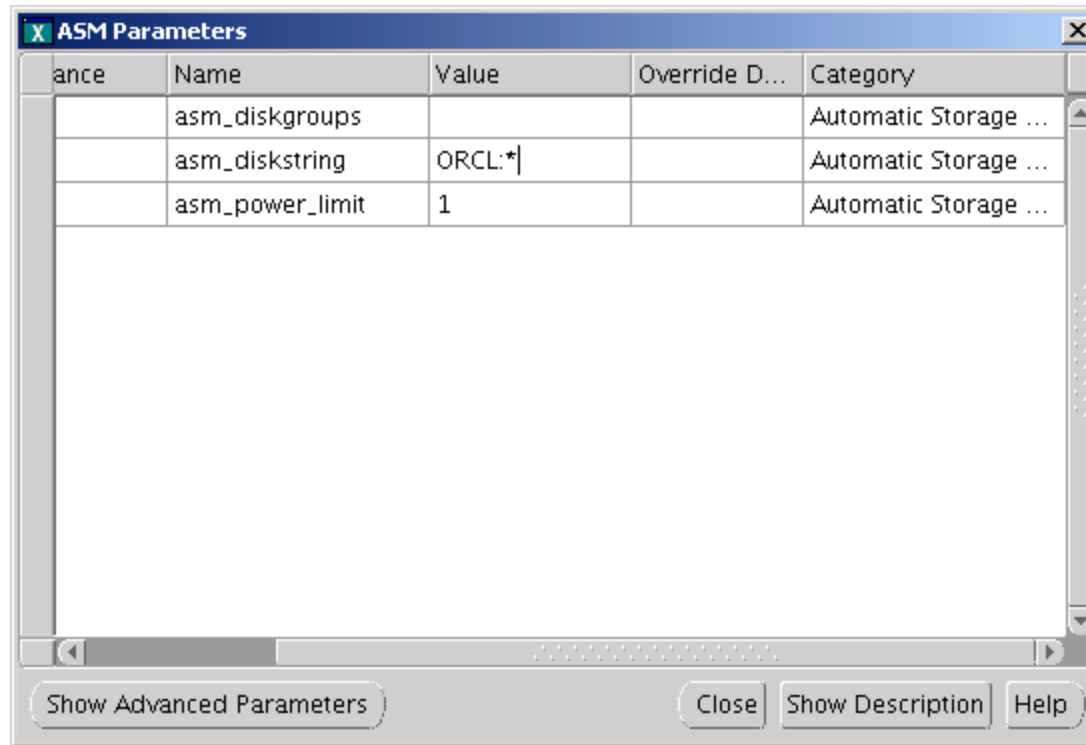
Initialization Parameter Filename:

Create server parameter file (SPFILE)

Server Parameter Filename:

Cancel Help < Back Next >





The image shows a dialog box titled "ASM Parameters" with a table of parameters. The table has four columns: "Name", "Value", "Override D...", and "Category". The "Name" column contains "asm\_diskgroups", "asm\_diskstring", and "asm\_power\_limit". The "Value" column contains empty, "ORCL:\*\|", and "1" respectively. The "Category" column contains "Automatic Storage ..." for all three. At the bottom of the dialog, there are four buttons: "Show Advanced Parameters", "Close", "Show Description", and "Help".

Name	Value	Override D...	Category
asm_diskgroups			Automatic Storage ...
asm_diskstring	ORCL:*\		Automatic Storage ...
asm_power_limit	1		Automatic Storage ...

Database Configuration Assistant, Step 8 of 15 : ASM Disk Groups


Select one or more disk groups to be used as storage for the database. You can choose to create a new disk group or add disks to an existing disk group.

Available Disk Groups

Select	Disk Group Name	Size (MB)	Free (MB)	Redundancy
--------	-----------------	-----------	-----------	------------

Create New Add Disks

Cancel Help Back Next



**Create Disk Group**

Disk Group Name:

Redundancy

High       Normal       External

Select Member Disks

Show Candidates     Show All

<input type="checkbox"/>	Disk Path	Header Status	ASM Name	Size (MB)
<input checked="" type="checkbox"/>	ORCL:disk1	PROVISIONED		5599

Note: If you don't see disks which you believe should be available, you may need to change the disk discovery path.

**X Database Configuration Assistant, Step 8 of 15 : ASM Disk Groups**


Select one or more disk groups to be used as storage for the database. You can choose to create a new disk group or add disks to an existing disk group.

Available Disk Groups

Select	Disk Group Name	Size (MB)	Free (MB)	Redundancy
<input checked="" type="checkbox"/>	DG1	5599	5549	EXTERN

Create New   Add Disks

Cancel   Help   Back   Next



Database Configuration Assistant, Step 9 of 15 : Database File Locations

Specify locations for the Database files to be created:

Use Database File Locations from Template


Use Common Location for All Database Files

Database Files Location:  Browse...

Use Oracle-Managed Files

Database Area:  Browse...

Multiplex Redo Logs and Control Files...

 If you want to specify different locations for any database files, pick either of the above options and use the Storage page to specify each location.

File Location Variables...

Cancel Help Back Next



# Oracle Managed Files (OMF)

- ASM
- File systems
- Makes management easy
  - create tablespace crm;
  - Files are managed by Oracle
- Related database parameter
  - db\_create\_file\_dest
  - db\_create\_online\_log\_dest\_n
  - db\_recover\_file\_dest

Expand All Collapse All

Select Name	Physical Size (KB)	Logical Size (KB)
[-] DG1		
[-] GP		
[-] CONTROLFILE		
[-] DATAFILE		
[-] EXAMPLE.267.1	154624	153608
[-] SYSAUX.257.1	319488	317448
[-] SYSTEM.256.1	452608	450568
[-] UNDOTBS1.258.1	47104	46088
[-] UNDOTBS2.268.1	26624	25608
[-] USERS.259.1	6144	5128
[+] ONLINELOG		
[+] PARAMETERFILE		
[+] TEMPFILE		
[-] spfilegp.ora	1024	2.5
[-] gp_db_files		
[-] user.dbf	6144	5128

Create Alias Create Directory Delete Rename

General Performance Templates Files



# Enterprise Manager Sample Session





Severity	Target Name	Target Type	Category	Name	Message	Alert Triggered	Last Value	Time
(No Alerts!)								

### Related Alerts


Severity	Target Name	Target Type	Category	Name	Message	Alert Triggered	Last Value	Time
(No Alerts!)								

### Job Activity

Jobs scheduled to start no more than 7 days ago

Status	Submitted to the Cluster Database	Submitted to Any Member Instance
(No Job Activity)		

### Critical Patch Advisories

 Patch Advisories **0**  
 Patch Advisory information may be stale. Oracle MetaLink credentials are not configured.




Affected Oracle Homes **0**

Oracle MetaLink Credentials [Not Configured](#)

### Related Links

- |  |                                |  |
|--|--------------------------------|--|
| <a href="#">Advisor Central</a>          | <a href="#">Alert History</a>  | <a href="#">All Metrics</a>              |
| <a href="#">Blackouts</a>                | <a href="#">Deployments</a>    | <a href="#">iSQL*Plus</a>                |
| <a href="#">Jobs</a>                     | <a href="#">Manage Metrics</a> | <a href="#">Metric Collection Errors</a> |
| <a href="#">Monitoring Configuration</a> |                                |  |

### Instances

Name 	Status	Alerts	Policy Violations	Performance Findings
<a href="#">gp_gp1</a>		<a href="#">3</a> <a href="#">2</a>	<a href="#">1</a>	<a href="#">10</a>
<a href="#">gp_gp2</a>		<a href="#">0</a> <a href="#">0</a>	<a href="#">0</a>	n/a


Automatic Storage Management: saptux1.dem.cpqcorp.net\_+ASM1 > Disk Group: DG1

### Disk Group: DG1


General Performance Templates Files

Name **DG1**  
State **MOUNTED**  
Redundancy **EXTERN**  
Total (GB) **11.00 GB**  
Free (GB) **8.93 GB**  
Pending Operations [0](#)

#### Disk Group Usage (GB)



#### Disk Group Usage History (GB)



### Member Disks

View

Select	ASM Disk Name	By Failure Group	Path	Read/Write Errors	State	Size (GB)	Used (GB)	Used (%)
<input checked="" type="radio"/>	<a href="#">DD</a>	DD	ORCL:DD	0	NORMAL	11.00	2.07	18.85

General Performance Templates Files

Automatic Storage Management: saptux1.dem.cpqcorp.net\_+ASM1 > Disk Group: DG1 > Add Disks

### Add Disks

Show SQL Cancel OK

Select Member Disks Only Candidate Disks

Select	Path	Header Status	Label	ASM Disk Name	Size	Size Unit	By Failure Group	Force Usage
<input checked="" type="checkbox"/>	ORCL:D	PROVISIONED	D		11264	MB		<input type="checkbox"/>

Show SQL Cancel OK

[Database](#) | [Setup](#) | [Preferences](#) | [Help](#) | [Logout](#)


Automatic Storage Management: saptux1.dem.cpqcorp.net\_+ASM1 > Disk Group: DG1

# Disk Group: DG1


General Performance Templates Files

Name **DG1**  
State **MOUNTED**  
Redundancy **EXTERN**  
Total (GB) **11.00 GB**  
Free (GB) **8.93 GB**  
Pending Operations **1**

## Disk Group Usage (GB)



## Disk Group Usage History (GB)



## Member Disks

View

Select	ASM Disk Name	By Failure Group	Path	Read/Write Errors	State	Size (GB)	Used (GB)	Used (%)
<input checked="" type="radio"/>	<a href="#">DD</a>	DD	ORCL:DD	0	NORMAL	11.00	2.07	18.85
<input type="radio"/>	<a href="#">D</a>	D	ORCL:D	0	NORMAL	11.00	0.00	.02

General Performance Templates Files

Automatic Storage Management: [saptux1.dem.cpqcorp.net\\_+ASM1](#) > Pending Operations: DG1

## Pending Operations: DG1

Data Retrieved July 2, 2004 9:53:22 AM CEST Refresh Real Time: Manual Refresh Refresh

Operation Type	Status	Desired Power	Actual Power	Operation Rate (Units per minute)	% Complete	Remaining Time (minutes)
REBAL	RUN	1	1	796	36.67	0

Database Control Database


Automatic Storage Management: saptux1.dem.cpqcorp.net\_+ASM1 > Disk Group: DG1

# Disk Group: DG1


General Performance Templates Files

Name **DG1**  
 State **MOUNTED**  
 Redundancy **EXTERN**  
 Total (GB) **22.00 GB**  
 Free (GB) **19.92 GB**  
 Pending Operations [0](#)

## Disk Group Usage (GB)



## Disk Group Usage History (GB)



## Member Disks

View

Select	ASM Disk Name	By Failure Group	Path	Read/Write Errors	State	Size (GB)	Used (GB)	Used (%)
<input checked="" type="radio"/>	DD	DD	ORCL:DD		0 NORMAL	11.00	1.06	9.66
<input type="radio"/>	D	D	ORCL:D		0 NORMAL	11.00	1.01	9.21

General Performance Templates Files

# ASM sql commands

- create diskgroup dg1 external redundancy disk '/dev/raw/raw1' name DISK1;
- create diskgroup dg2 external redundancy disk 'ORCL:DISK2' name DISK2;
- alter diskgroup dg2 add disk 'ORCL:DISK3' name DISK3;

**Make sure the asm\_diskstring parameter is set accordingly**

# Adding Disk Space

- Two options to make space available
  - Modify the size of an existing virtual disk (LUN)
  - Add a new virtual disk (LUN)
- System components involved
  - Disk sub system
  - Linux operating system
  - ASM disk
  - ASM diskgroup
- Can be done off-line or online



hp StorageWorks  
command view eva

Appliance:  
SMAD239FK34D006 16.41.101.115

- Dirk
- Linux
  - ASM\_D\_Copy
  - linux-data
  - linux-data2
    - ACTIVE
  - linux-data2-small
  - linux-data3
    - ACTIVE
  - linux-data3-small
  - linux-new
- hp-ux
- hp-ux-ora8
- hpux-ipf
- tru64
- w2k
- w2k3
- ca-snap
- lvm-mirr1
- lvm-mirr2
- tristan-perf
- Hosts

### Vdisk Active Member Properties

Save changes Create snapshot Create Snapclone ?

General Presentation Data Replication

**Identificati**

Name:   Good

Family Name:

World Wide Name:  -2004 11:36:53

6005-08b4-0000

UUID:  ed write-back

6005-08b4-0001-0220-0000-f000-0e0d-0000

Read:  On

**Microsoft Internet Explorer**

? You are requesting a change in the size of this Vdisk. If the hosts to which this Vdisk is presented are running operating systems that cannot handle such a change, the volume expansion is not guaranteed to happen. Are you sure you wish to continue?

OK Cancel

**Attributes**

Type:	Original
Disk Group:	big20x72GB
Capacity Req:	<input type="text" value="12"/> GB
Capacity Used:	11 GB
Redundancy:	Vraid1
Write Protect:	<input type="text"/> No



Appliance: SMAD239FK34D006 16.41.101.115

- HSV Storage Network
  - ping
    - Virtual Disks
      - Dirk
      - Linux
        - ASM\_D\_Copy
        - linux-data
        - linux-data2
          - ACTIVE
        - linux-data2-small
        - linux-data3
          - ACTIVE
        - linux-data3-small
        - linux-new
      - hp-ux
      - hp-ux-ora8
      - hpux-ipf
      - tru64
      - w2k
      - w2k3
      - ca-snap

### Vdisk Active Member Properties



Identification	
Name:	Active
Family Name:	linux-data3
World Wide LUN Name:	
6005-08b4-0001-0220-0000-f000-0ebd-0000	

Condition/State	
Operational State:	<input checked="" type="checkbox"/> Operation in progress

Date/Time	
Created:	01-Jul-2004 11:36:53

Cache Policies	
Write:	Mirrored write-back
Read:	On

Attributes	
Disk Group:	big20x72GB
Capacity Req:	12 GB
Capacity Used:	Unknown GB
Redundancy:	Vraid1
Write Protect:	No

Comments



Appliance: SMAD239FK34D006 16.41.101.115

- Dirk
- Linux
  - ASM\_D\_Copy
  - linux-data
  - linux-data2
    - ACTIVE
  - linux-data2-small
  - linux-data3
    - ACTIVE
  - linux-data3-small
  - linux-new
- hp-ux
- hp-ux-ora8
- hpux-ipf
- tru64
- w2k
- w2k3
- ca-snap
- lvm-mirr1
- lvm-mirr2
- tristan-perf
- Hosts

### Vdisk Active Member Properties

Save changes Create snapshot Create Snapclone ?

General Presentation Data Replication

Identification	
Name:	ACTIVE
Family Name:	linux-data3
World Wide LUN Name:	6005-08b4-0001-0220-0000-f000-0ebd-0000
UUID:	6005-08b4-0001-0220-0000-f000-0ebd-0000

Condition/State	
Operational State:	<input checked="" type="checkbox"/> Good

Date/Time	
Created:	01-Jul-2004 11:36:53

Cache Policies	
Write:	Mirrored write-back
Read:	On

Attributes	
Type:	Original
Disk Group:	big20x72GB
Capacity Req:	12 GB
Capacity Used:	12 GB
Redundancy:	Vraid1
Write Protect:	No

# Adding Disk Space Online (Option1)

- Modifying the size of an existing virtual disk (LUN)
  - Linux operating system
    - Remove disk from the diskgroup first
      - alter diskgroup dg1 drop disk DISK2
    - Change size of LUN on the EVA
    - Remove device from Linux and add it back
      - echo “scsi remove-single-device a b c d” > /proc/scsi/scsi
      - echo “scsi add-single-device a b c d” > /proc/scsi/scsi
    - Need to be done on all nodes for RAC
  - ASM disk
    - /etc/init.d/oracleasm deletedisk DISK2
    - /etc/init.d/oracleasm scandisks on other nodes (for RAC only)
    - /etc/init.d/oracleasm createdisk DISK2 /dev/sdc
    - /etc/init.d/oracleasm scandisks on other nodes (for RAC only)

# Adding Disk Space Online

- ASM diskgroup
  - alter diskgroup dg1 add disk 'ORCL:DISK2' name DISK2;

Example based on ASMLib

# Adding Disk Space Online (Option2)

- Adding a new virtual device (LUN) online
  - EVA
  - Linux operating system
    - `hp_rescan -a`
    - Need to be done on all nodes for RAC
  - ASM disk
    - `/etc/init.d/oracleasm createdisk NDISK /dev/sdc`
    - `/etc/init.d/oracleasm scandisks` on other nodes (for RAC only)
  - ASM diskgroup
    - `alter diskgroup dg1 add disk 'ORCL:NDISK' name NDISK;`

Example based on ASMLib

# Adding Disks for Performance

- General steps involved
  - Make new disks available
  - Make disks or volumes known to the operating system
  - Create ASM disks
  - Add ASM disks into ASM disk group

# Adding Disks for Performance on EVA



- Make disks available
- Move disks into EVA disk group
- EVA takes advantage of the additional disks automatically (rebalancing)
- Transparent for
  - Operating system (LUN size is not changed)
  - ASM (still same virtual disks)
  - Oracle database
- All online
- This scenario would also work for ocfs because it is completely transparent to the host





Appliance:  
SMAD239FK34D006 16.41.101.115

- [-] big20x72GB
  - Disk 001
  - Disk 002
  - Disk 003
  - Disk 004
  - Disk 005
  - Disk 007
  - Disk 009
  - Disk 010
  - Disk 011
  - Disk 013
  - Disk 017
  - Disk 019
  - Disk 020
  - Disk 021
  - Disk 022
  - Disk 023
  - Disk 024
  - Disk 028
  - Disk 027
  - Disk 030
- [-] small8x36GB
  - Disk 008
  - Disk 014
  - Disk 015

## Disk Group Properties

Save changes **Add disks** Locate Delete ?

General **Vdisks**

### Identification

**Name:** big20x72GB

**UUID:**  
6005-08b4-0001-0256-0000-d000-0057-0000

### Capacity

<b>Total:</b>	1366.03 GB
<b>Available:</b>	
<b>Vraid0</b>	401.99 GB
<b>Vraid1</b>	201.03 GB
<b>Vraid5</b>	321.6 GB
<b>Total disks:</b>	20

### Condition/State

<b>Operational state:</b>	<input checked="" type="checkbox"/> Good
<b>Leveling state:</b>	<input checked="" type="checkbox"/> Inactive
<b>Leveling progress:</b>	n/a
<b>RSS Disk state:</b>	Mirrored

### Disk failure protection

<b>Requested level:</b>	None
<b>Actual level:</b>	None

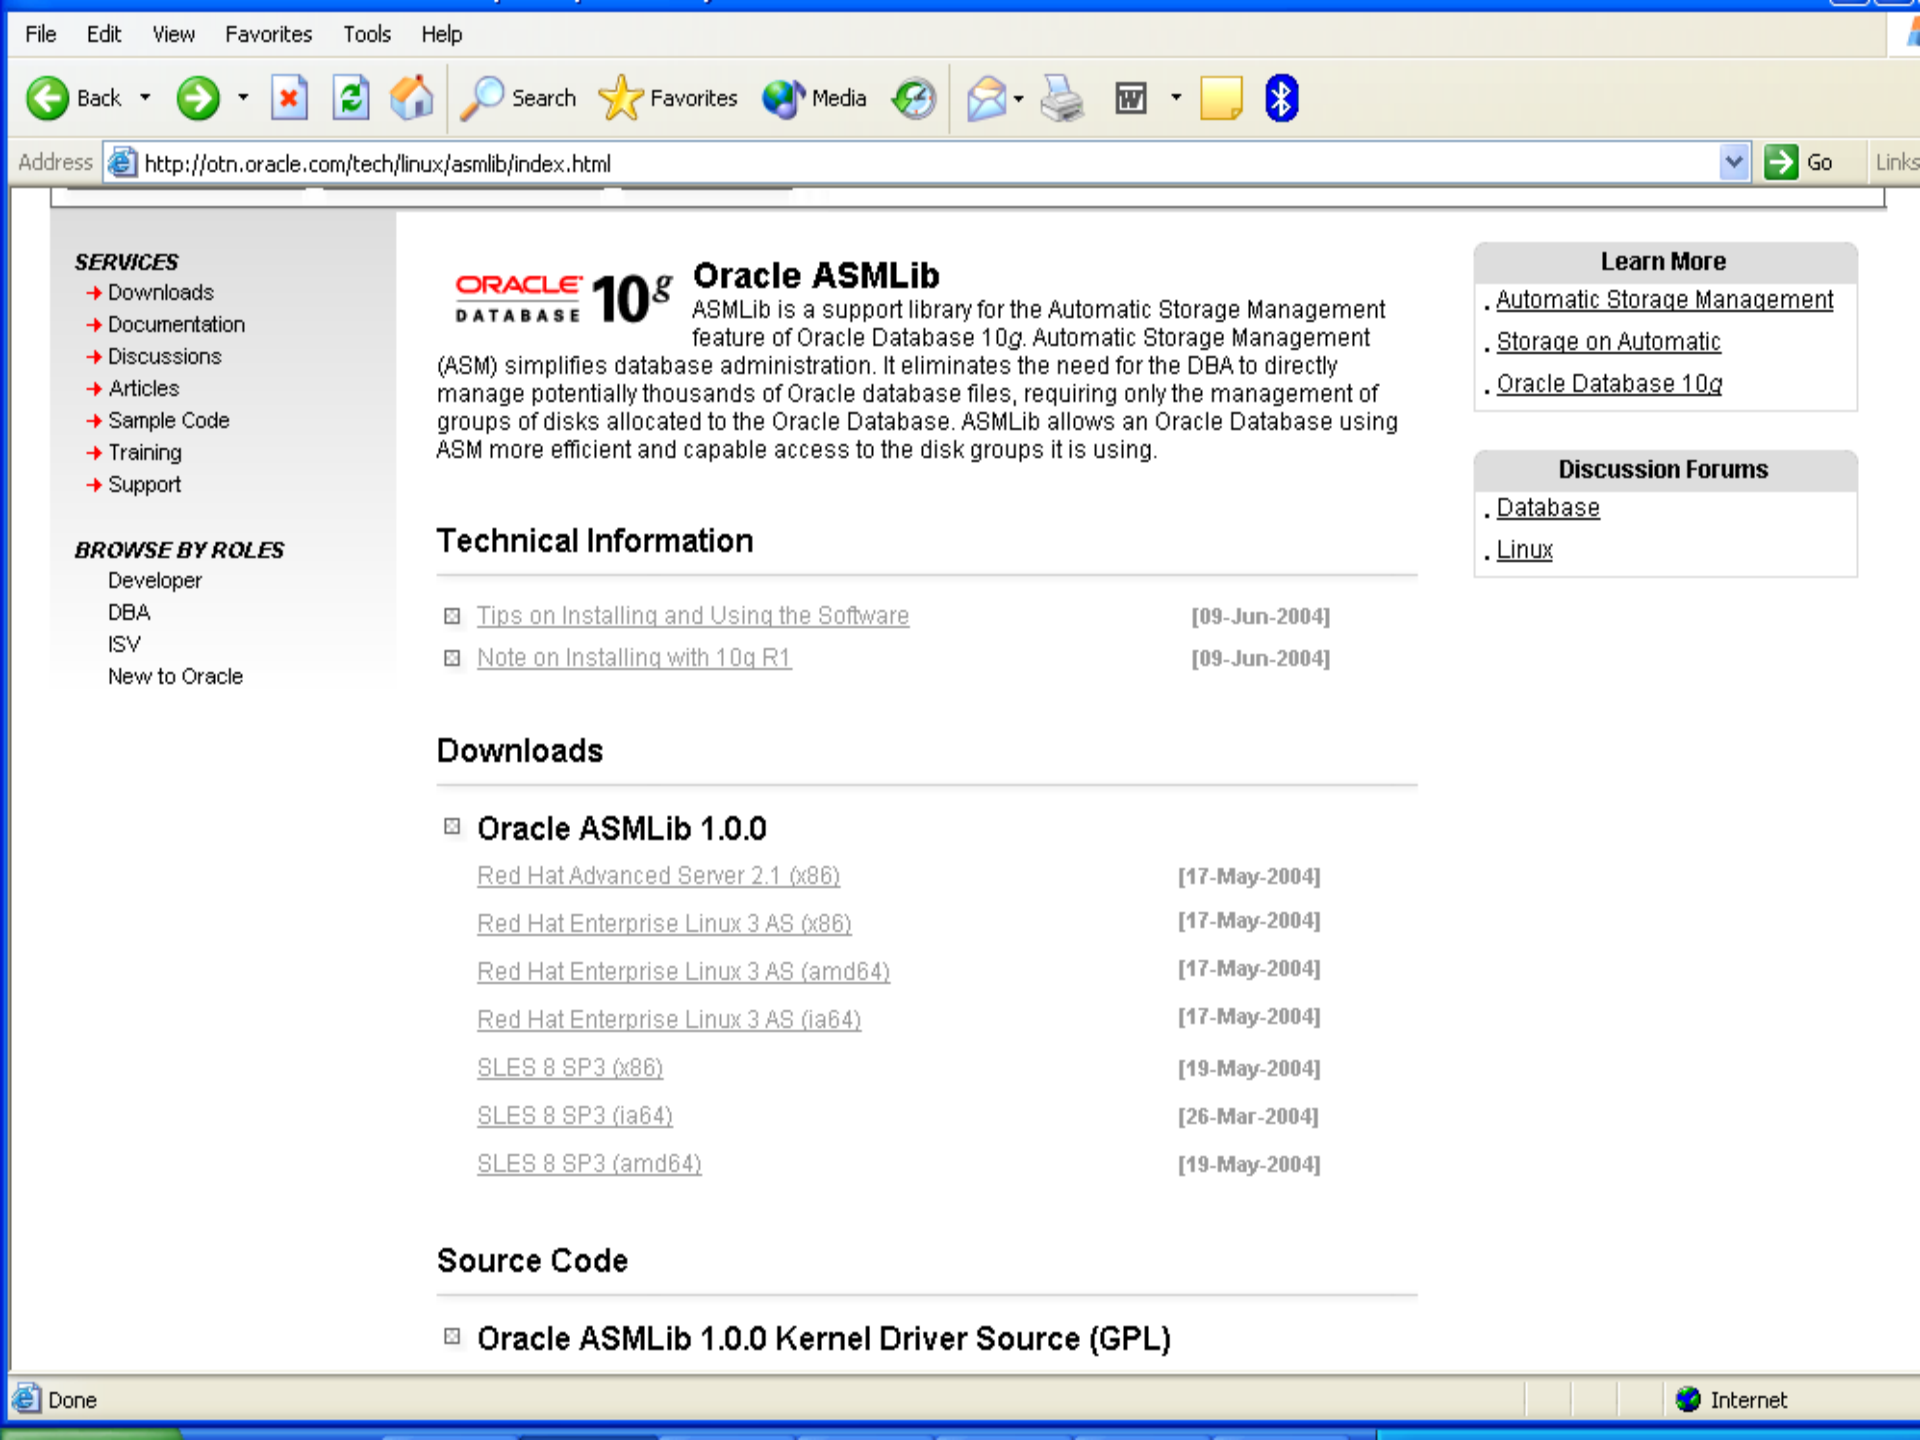
### Occupancy

<b>Total:</b>	963.26 GB
<b>Alarm level:</b>	90 %

### Comments

# ASM and EVA Best Practice

- Oracle recommends to use ASMLib
  - Native async IO
  - Disk discovery
- Install asm packages before installing Oracle
  - [oracleasm-1.0.0-1.i386.rpm](#) – User space library
  - [oracleasm-support-1.0.2-1.ia64.rpm](#) - Driver support files
  - [oracleasm-2.4.21-EL-1.0.0-1.ia64.rpm](#) - Driver for all kernels
- Perform custom Oracle install (R1 restriction)
  - Otherwise DBCA will not discover ASM disks



**SERVICES**

- Downloads
- Documentation
- Discussions
- Articles
- Sample Code
- Training
- Support

**BROWSE BY ROLES**

- Developer
- DBA
- ISV
- New to Oracle



**Oracle ASMLib**

ASMLib is a support library for the Automatic Storage Management feature of Oracle Database 10g. Automatic Storage Management (ASM) simplifies database administration. It eliminates the need for the DBA to directly manage potentially thousands of Oracle database files, requiring only the management of groups of disks allocated to the Oracle Database. ASMLib allows an Oracle Database using ASM more efficient and capable access to the disk groups it is using.

**Technical Information**

- [Tips on Installing and Using the Software](#) [09-Jun-2004]
- [Note on Installing with 10g R1](#) [09-Jun-2004]

**Downloads**

- Oracle ASMLib 1.0.0**
  - [Red Hat Advanced Server 2.1 \(x86\)](#) [17-May-2004]
  - [Red Hat Enterprise Linux 3 AS \(x86\)](#) [17-May-2004]
  - [Red Hat Enterprise Linux 3 AS \(amd64\)](#) [17-May-2004]
  - [Red Hat Enterprise Linux 3 AS \(ia64\)](#) [17-May-2004]
  - [SLES 8 SP3 \(x86\)](#) [19-May-2004]
  - [SLES 8 SP3 \(ia64\)](#) [26-Mar-2004]
  - [SLES 8 SP3 \(amd64\)](#) [19-May-2004]

**Source Code**

- Oracle ASMLib 1.0.0 Kernel Driver Source (GPL)**

**Learn More**

- . [Automatic Storage Management](#)
- . [Storage on Automatic](#)
- . [Oracle Database 10g](#)

**Discussion Forums**

- . [Database](#)
- . [Linux](#)

# ASM and EVA Best Practice

- Decide about number of ASM disk groups
  - Number of databases
  - Size of database
  - RAID requirement
  - Group disks with similar characteristics into same disk group
  - Consider manageability requirements

# ASM and EVA Best Practice

- For Enterprise Virtual Array (EVA)
  - Usually no need to have more than 2 EVA disk groups
    - One EVA disk group is often sufficient
  - Define RAID level based on customer requirements
  - Decide about number of virtual disks based on
    - Single or redundant IO path configuration
    - Online ASM configuration changes
    - Multiple storage systems
  - Define external redundancy for the ASM disk group

# ASM Administration and Monitoring

- ASM views
- Enterprise Manager
- HP OpenView

# ASM Administration and Monitoring

- v\$asm views
  - v\$asm\_disk
  - v\$asm\_client
  - ....
- Available for ASM instance and DB instance
  - Provide different information
  - For example v\$asm\_client
    - **ASM** ,information about databases using disk groups of ASM instance
    - **DB** one row for ASM instance when open ASM files


Automatic Storage Management: [joe\\_+ASM1](#) > Disk Group: DG1

## Disk Group: DG1


**General** [Performance](#) [Templates](#) [Files](#)

Name **DG1**  
 State **MOUNTED**  
 Redundancy **EXTERN**  
 Total (GB) **5.66 GB**  
 Free (GB) **4.50 GB**  
 Pending Operations [Q](#)

### Disk Group Usage (GB)



### Disk Group Usage History (GB)



### Member Disks

View

Select	ASM Disk Name	By Failure Group	Path	Read/Write Errors	State	Size (GB)	Used (GB)	Used (%)
<input type="checkbox"/>	<a href="#">DG1_0000</a>	DG1_0000	/dev/raw/raw3	0	NORMAL	5.66	1.16	20.53

**General** [Performance](#) [Templates](#) [Files](#)



Automatic Storage Management: saptux1.dem.cpqcorp.net\_+ASM1 > Disk Group: DG1


# Disk Group: DG1

General Performance Templates Files


This page shows cluster-wide performance information. Use the links under the graphs to view graphs that show the performance of individual cluster members.

Collected From Target July 2, 2004 10:53:36 AM CEST Refresh Real Time: 15 Second Refresh Refresh


## Disk Group I/O Response Time



## Disk Group I/O Operations



## Disk Group Throughput



## Disk Group I/O Cumulative Statistics

Accumulated Since May 25, 2004 12:00:00 AM

Expand All | Collapse All

Average Response	Average Throughput (MB per second)	Total Read	Total Write	Total Read	Total Write

Telnet saptux1

GlancePlus - IO By Disk

File Reports Configure

System: saptux1 Last Update: 12:56:56

Active Disks: All 7 Selected Total I/O

Device Name	Phys Reads	Phys Read KB Rate
sda	50	286.0
sdb	15096	121123.0
sdc	31	92.0
sdd	0	0.0
sde	0	0.0
sdf	0	0.0

GlancePlus - Main

File Reports Adviser Configure Help

System: saptux1 Graph Points: 15 Int: ?

Range : 12:55:36 - 12:56:56 Dur: 1:2

ALARM

CPU

Memory

Disk

Network

GlancePlus - IO By Disk

File Reports Configure

System: saptux2 Last Update: 13:02:21

Active Disks: All 7 Selected Total I/O

Device Name	Phys Reads	Phys Read KB Rate
sda	27	86.2
sdb	16809	134937.3
sdc	0	0.0
sdd	0	0.0
sde	0	0.0
sdf	0	0.0

GlancePlus - Main

File Reports Adviser Configure Help

System: saptux2 Graph Points: 15 Int: ?

Range : 13:01:01 - 13:02:21 Dur: 1:2

ALARM

CPU

Memory

Disk

Network

# Summary

- ASM and HP storage both provide storage virtualization
- Let the storage system do what it is designed for
  - Stripping and mirroring
  - Snapshots
- Take advantage of additional ASM functionality which is not possible on disk controller level
- ASM and HP storage provide
  - Increased Uptime
  - Reduces Cost of Management
  - Optimized Resource Utilization



# HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:



RECOMMENDED TRAINING VENUE FOR THE  
**HP Certified Professional**

