

# Automatic Storage Management 11 g release 2



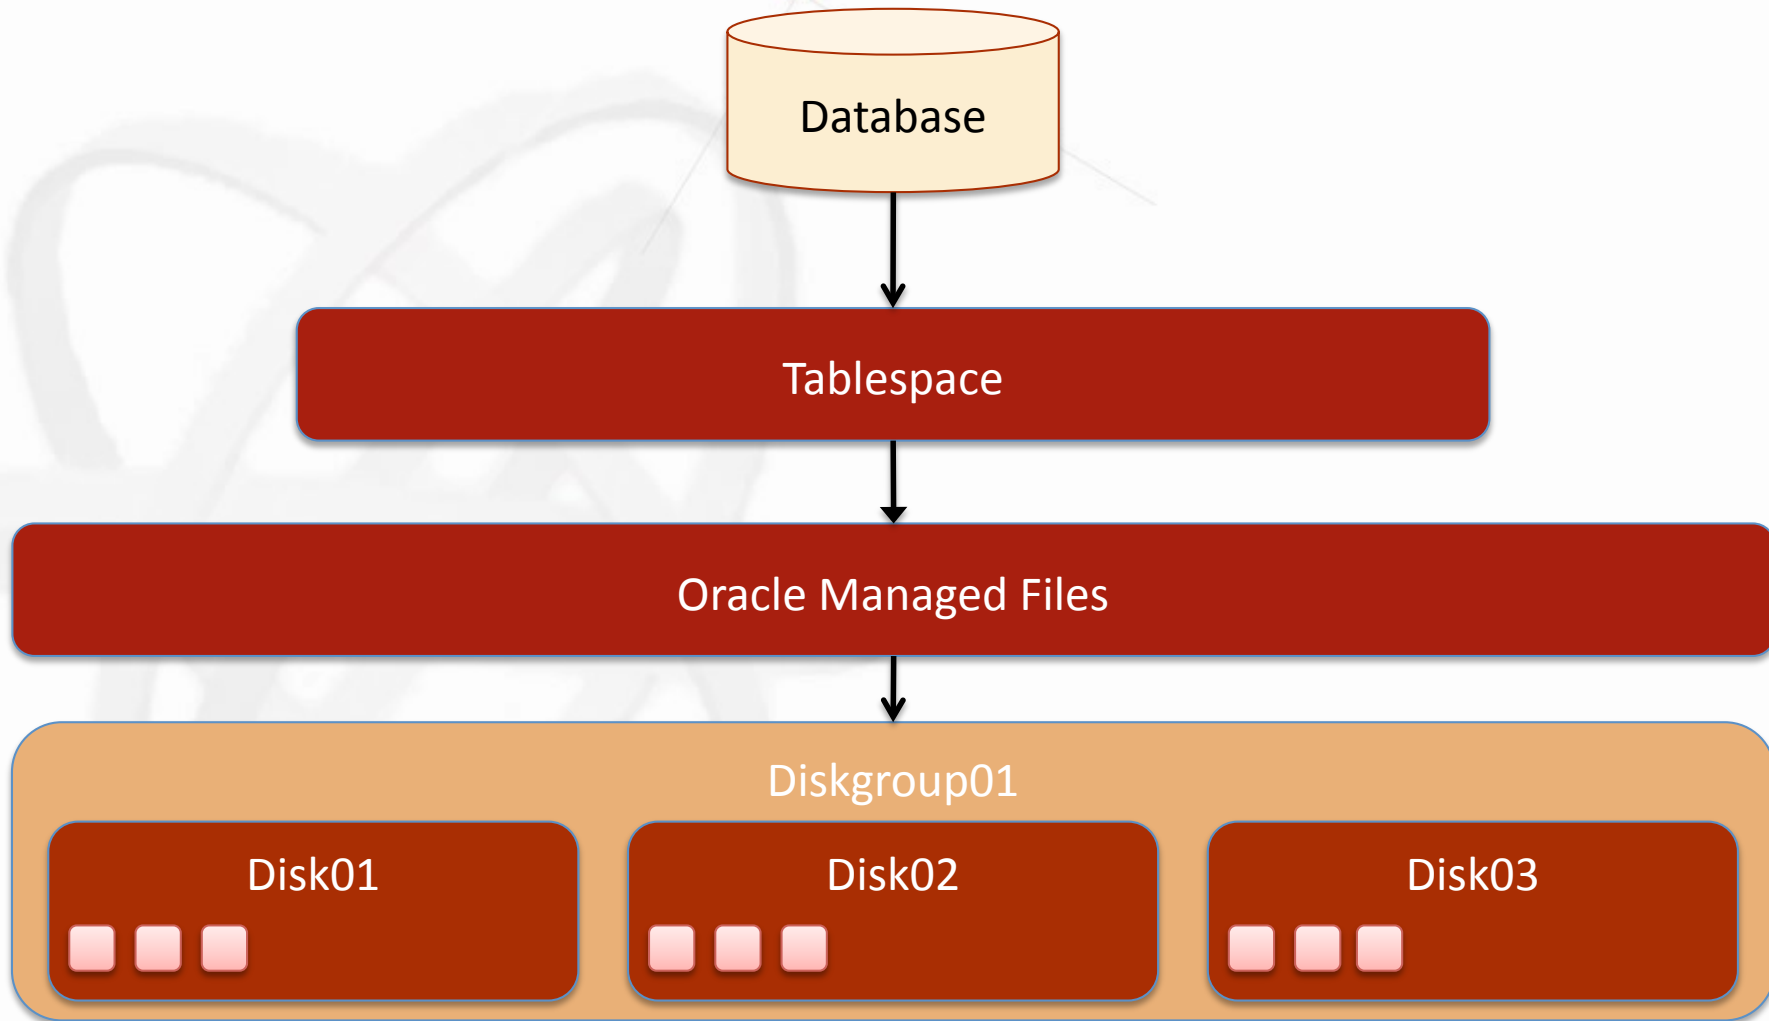
OGH DBA DAG

14 september 2010

Rob den Braber

- Wat is ASM?
- New Features ASM 11g release 2
- Infrastructuur ASM
- Oracle ACFS
- Voor- en nadelen van ASM-gebruik

- Geïntegreerde volume manager voor het opslaan van Oracle database files:
  - Datafiles
  - Redo logs
  - Backups
  - Controlfiles
  - Spfiles
- Data wordt automatisch verdeeld over alle disken in een diskgroup.



```
+ASM1.__oracle_base='/u01/app/oracle'  
+ASM1.asm_diskgroups='DGOGH'  
*.asm_diskstring='/dev/asmdisk*'  
*.asm_power_limit=1  
*.diagnostic_dest='/u01/app/oracle'  
*.instance_type='asm'  
*.large_pool_size=12M  
*.remote_login_passwordfile='EXCLUSIVE'
```

```
SQL> show parameter db_create_file_dest
```

NAME	VALUE
db_create_file_dest	+DGOGH

```
SQL> show parameter spfile
```

NAME	VALUE
spfile	+DGOGH/oghdb/spfileoghdb.ora

```
SQL> select name from v$datafile;
```

```
NAME
```

```
-----  
+DGOGH/oghdb/datafile/system.260.727806171  
+DGOGH/oghdb/datafile/sysaux.261.727806183  
+DGOGH/oghdb/datafile/undotbs1.262.727806197  
+DGOGH/oghdb/datafile/users.264.727806219
```

- Oracle Cluster Registry (OCR) and Voting Files stored in Oracle ASM
- Oracle Restart
- Oracle ASM Configuration Assistant ASMCA
- ASMCMD verbeteringen
- Intelligent Data Placement
- Oracle ASM File Access Control
- Diskgroup rename

- Oracle Automatic Storage Management Cluster File System (ACFS)
- Oracle ASM Dynamic Volume Manager
- Oracle Automatic Storage Management Cluster File System Snapshots
- Oracle Enterprise Manager Support for Oracle ACFS

- ASM is onderdeel van de grid-infrastructuur (Clusterware + ASM)
- Diskgroups worden ook als resource geregistreerd in CRS en zijn dus te beheren met CRSCTL en SRVCTL

```
[oracle@rac01 ~]$ crsctl status resource |grep dg  
NAME=ora.DG_GRIDDB.dg  
NAME=ora.DISKGROUP01.dg  
NAME=ora.GRID_DATA.dg
```

- Oracle Cluster Registry (OCR) in Oracle ASM

```
[root@rac01 ~]# ocrcheck -config
Oracle Cluster Registry configuration is :
      Device/File Name          : +GRID_DATA
ASMCMD [+GRID_DATA/rac-cluster/ocrfile] > ls -l
Type      Redund  Striped  Time                Sys  Name
OCRFILE   UNPROT   COARSE   SEP 28 11:00:00    Y    REGISTRY.255.697565077
```

- Voting Files in Oracle ASM

```
[root@rac01 ~]# crsctl query css votedisk
##  STATE      File Universal Id                File Name                        Disk group
--  -
1.  ONLINE     332b75f479474f29bfd84fc540fda392 (/dev/asmdisk1p1) [GRID_DATA]
Located 1 voting disk(s).
```

- Beheer ASM verwijderd uit de DBCA

ASM Configuration Assistant: Configure ASM: ASM Cluster File Systems

ASM Instances | Disk Groups | Volumes | **ASM Cluster File Systems**

ASM Cluster File System (ACFS) can be used to store files such as Executables, Oracle Diagnostic files, Application configuration files, etc. To use ACFS, you need to create ASM Volume first.

Tip: The table shows both mounted and dismounted filesystems. For dismounted filesystems, the last known mountpoint is shown. To perform operations on an ASM Cluster File System, right mouse click on the row.

ASM Cluster File Systems

Mount Point	State	Registered Mount ...	Volume Device	Size (GB)	Volume
/u01/volume01	MOUNTED(1 of 1)	/u01/volume01	/dev/asm/volume01-380	1.00	VOLUME01

Note: Some ACFS commands can be executed as privileged/root user only. If you choose any of these operations, ASMCA will generate the command that can be executed as privileged/root user manually.

Buttons: Create, Show Mount All Command, Show Dismount All Command

Buttons: Help, Exit

- Compleet beheer vanuit EM mogelijk

ORACLE Enterprise Manager 11g Database Control

Setup Preferences Help Logout  
Cluster Database

### Automatic Storage Management: +ASM1\_rac01

Home Performance Disk Groups Configuration Users ASM Cluster File System

Data Retrieved Nov 9, 2009 9:02:11 AM CET Refresh

#### General

Current Status **Up**  
 Up Since **Nov 9, 2009 8:10:11 AM CET**  
 Availability (%) **7.15**  
(Last 24 hours)  
 Instance Name **+ASM1**  
 Version **11.2.0.1.0**  
 Host **rac01**  
 Oracle Home **/u01/app/11.2.0/grid**

#### Diagnostic Summary

Alert Log **No ORA- errors**  
 Active Incidents [0](#)

#### Serviced Databases

Name	Disk Groups	Failure Groups	Allocated Space (GB)	Availability	Alerts
<a href="#">GRIDDB_GRIDDB_1</a>	DG_GRIDDB	n/a	2.02		<a href="#">1</a>
RAC-CLUSTER	GRID_DATA	n/a	0.26		Not Monitored

#### Serviced ASM Cluster File Systems

Mount Point	Availability	State	Used (%)	Used (GB)	Size (GB)	Allocated Space (GB)	Volume	Disk Group
<a href="#">/u01/volume01</a>		Mounted		9.77	0.0978	1	<a href="#">VOLUME01</a>	<a href="#">DISKGROUP01</a>

#### Disk Group Usage (GB)

Legend:  
■ Unallocated  
■ Internal  
■ GRIDDB\_GRIDDB\_1  
■ RAC-CLUSTER  
■ ASM Cluster File Systems

- Compleet beheer vanuit EM mogelijk

ORACLE Enterprise Manager 11g Database Control

Setup Preferences Help Logout  
Cluster Database

Logged in As SYS / SYSASM




Automatic Storage Management: +ASM1\_rac01

Home Performance **Disk Groups** Configuration Users ASM Cluster File System

Create Mount All Dismount All

Mount Dismount Rebalance Check Delete

Select All | Select None

Select	Name	State	Redundancy	Size (GB)	Used (GB)	Used (%)	Usable Free (GB)	Member Disks
<input type="checkbox"/>	DG_GRIDDB	MOUNTED	EXTERN	4.00	2.12	 53.00	1.88	1
<input type="checkbox"/>	DISKGROUP01	MOUNTED	EXTERN	8.00	1.06	 13.29	6.93	2
<input type="checkbox"/>	GRID_DATA	MOUNTED	EXTERN	4.00	0.39	 9.67	3.61	1

TIP The usable free space specifies the amount of space that can be safely used for data. A value above zero means that redundancy can be properly restored after a disk failure.

TIP Mount All and Dismount All operation will only mount and dismount the disk groups specified in the Auto Mount Disk Groups parameter.

Home Performance **Disk Groups** Configuration Users ASM Cluster File System

**Related Links**

<a href="#">Access</a>	<a href="#">Alert History</a>	<a href="#">Alert Log Contents</a>
<a href="#">All Metrics</a>	<a href="#">Blackouts</a>	<a href="#">Metric and Policy Settings</a>
<a href="#">Metric Collection Errors</a>	<a href="#">Monitoring Configuration</a>	<a href="#">Target Properties</a>

Cluster | Database | Setup | Preferences | Help | Logout

Copyright © 1996, 2009, Oracle. All rights reserved.  
Oracle, JD Edwards, PeopleSoft, and Retek are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.  
[About Oracle Enterprise Manager](#)

- Compleet beheer vanuit EM mogelijk

ORACLE Enterprise Manager 11g Database Control Setup Preferences Help Logout  
Cluster Database

Logged in As SYS / SYSASM

**Automatic Storage Management: +ASM1\_rac01**

Home Performance Disk Groups Configuration Users **ASM Cluster File System**

Create Mount All Dismount All

Mount Dismount Delete Create Snapshot View Content Register Deregister Actions  Go

Select	Mount Point	Availability	State	Snapshots	Used (%)	Used (GB)	Size (GB)	Allocated Space (GB)	Redundancy	Volume	Disk Group
<input checked="" type="checkbox"/>	/u01/volume01		Mounted	0	9.78	0.0978	1	1	1 Unprotected	VOLUME01	DISKGROUP01

TIP Allocated Space is determined by the size and redundancy of the File System

Home Performance Disk Groups Configuration Users **ASM Cluster File System**

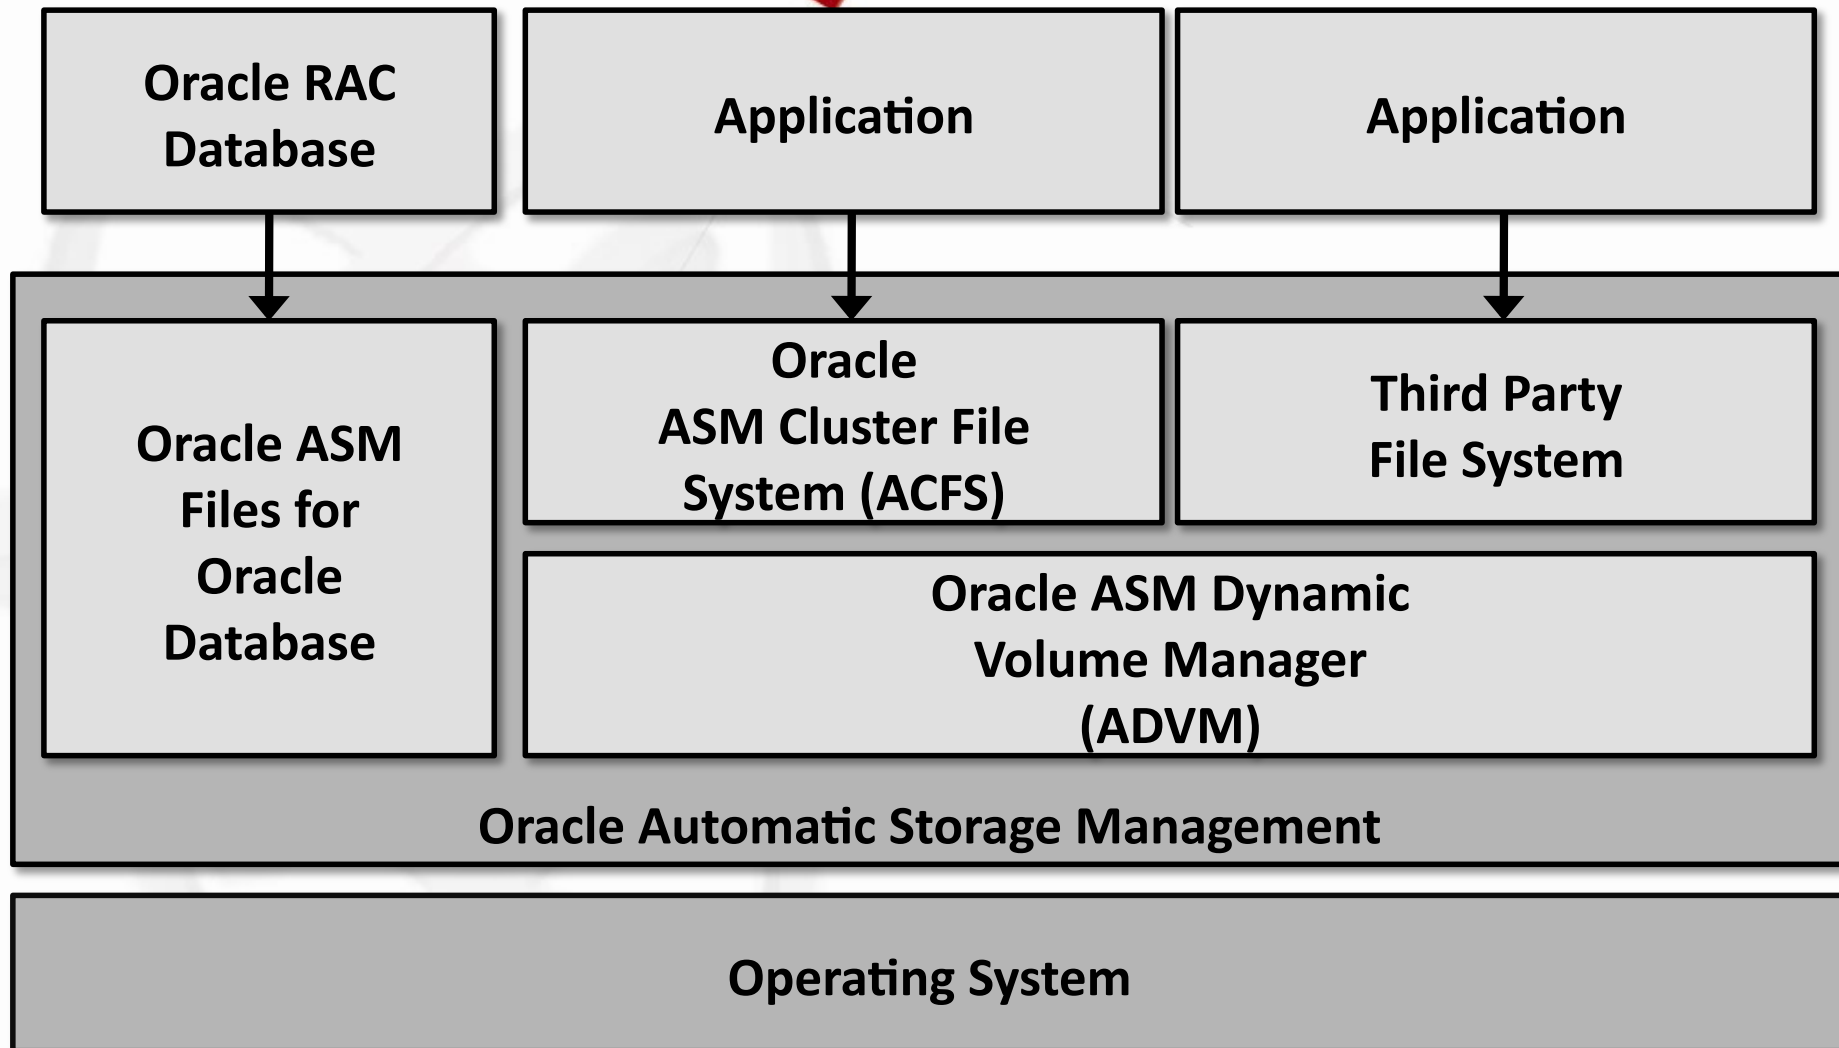
**Related Links**

<a href="#">Access</a>	<a href="#">Alert History</a>	<a href="#">Alert Log Contents</a>
<a href="#">All Metrics</a>	<a href="#">Blackouts</a>	<a href="#">Jobs</a>
<a href="#">Metric and Policy Settings</a>	<a href="#">Metric Collection Errors</a>	<a href="#">Monitoring Configuration</a>
<a href="#">Open Telnet Session</a>	<a href="#">Target Properties</a>	

Cluster | Database | Setup | Preferences | Help | Logout

- Voor het hernoemen van een diskgroup in het geval van een clone.
- Gebeurt in twee fases:
  - Genereren van een configuratie file
  - Hernoemen van de diskgroup aan de hand van de configuratie file

Concreet in Oracle oplossingen



# Files geïnstalleerd door root script

Concreet in Oracle oplossingen

- /sbin directory
  - fsck.acfs.bin
  - advmutil.bin
  - mount.acfs.bin
  - mkfs.acfs.bin
  - acfsdbg.bin
  - acfsutil.bin
  - mount.acfs
  - mkfs.acfs
  - fsck.acfs
  - advmutil
  - acfsutil
  - acfsdbg
- /lib/modules/2.6.18-8.el5/extra/usm/
  - oracleacfs.ko
  - oracleadvn.ko
  - oracleoks.ko

- Oracle ACFS: handelt de Oracle ACFS file en directory operaties af
- Oracle ADVN: Zorgt voor de block devices op het OS aan de hand van de ASM volume files
- Oracle Kernel Services Driver (OKS): zorgt voor de memory allocatie en de distributed locking service.

```
[root@rac01 ~]# lsmod | grep oracle
```

Module	Size	Used by
oracleacfs	787460	2
oracleadvn	177792	6
oracleoks	226656	2
oracleacfs,oracleadvn		

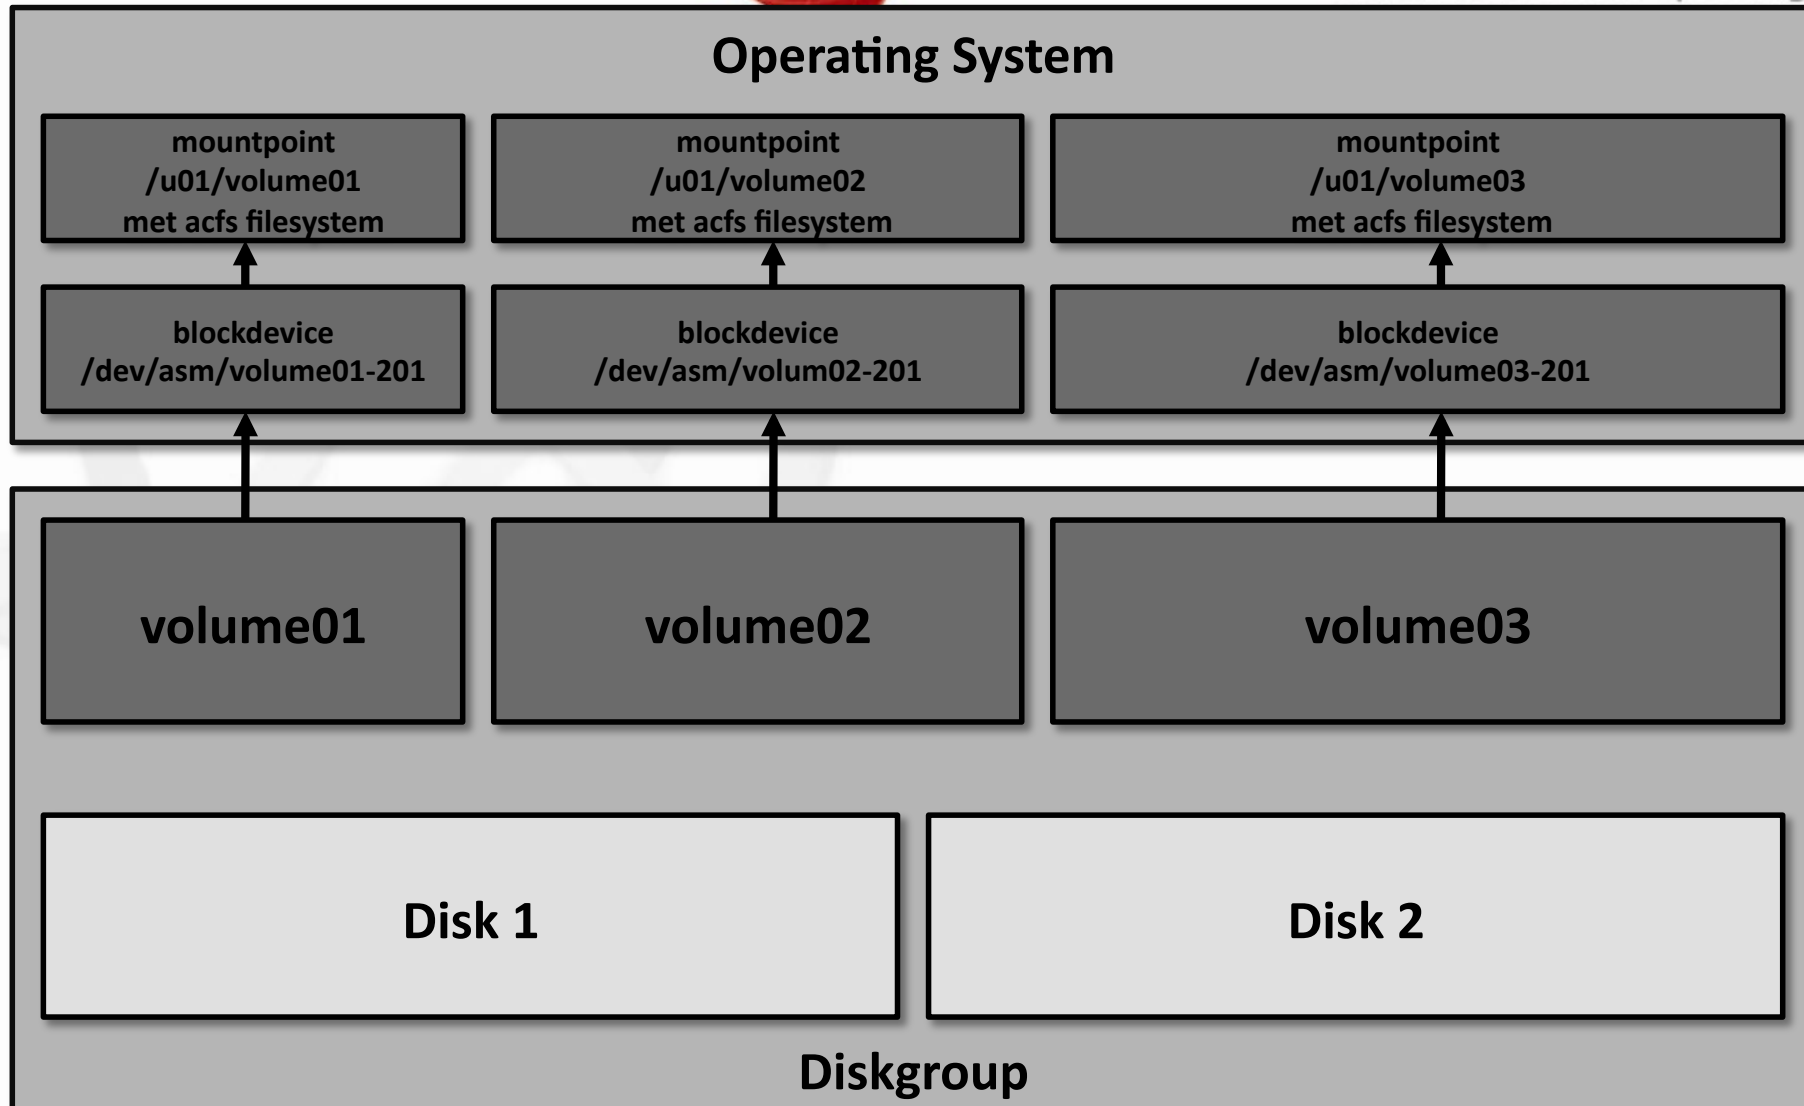
- Resource ora.drivers.acfs wordt gestart door Oracle High Availability Services

```
[root@rac01 ~]# crsctl stat resource ora.drivers.acfs -init  
NAME=ora.drivers.acfs  
TYPE=ora.drivers.acfs.type  
TARGET=ONLINE  
STATE=ONLINE on rac01
```

- Extra processen voor ADVM/ACFS

```
[oracle@rac01 ~]$ ps -ef |egrep "vbg|acfs|oks|vmb" |grep -v grep |sort --key 8
root      3389      1  0 08:08 ?        00:00:00 [acfs_bast0]
root      3390      1  0 08:08 ?        00:00:00 [acfs_bast1]
root      3391      1  0 08:08 ?        00:00:00 [acfs_bast2]
root      3392      1  0 08:08 ?        00:00:00 [acfs_bast3]
root      3393      1  0 08:08 ?        00:00:00 [acfs_bast4]
root      3394      1  0 08:08 ?        00:00:00 [acfs_bast5]
root      3395      1  0 08:08 ?        00:00:00 [acfs_bast6]
root      3396      1  0 08:08 ?        00:00:00 [acfs_bast7]
root      30003     1  0 11:29 ?        00:00:00 [acfsds2]
root      3388      1  0 08:08 ?        00:00:00 [acfsioerrlog]
root      30004     1  0 11:29 ?        00:00:00 [acfssnap2]
root      30001     1  0 11:29 ?        00:00:00 [acfsvol2]
oracle    4380      1  0 08:12 ?        00:00:06 asm_acfs_+ASM1
oracle    4378      1  0 08:12 ?        00:00:00 asm_vbg0_+ASM1
oracle    4406      1  0 08:12 ?        00:00:00 asm_vbg1_+ASM1
oracle    4408      1  0 08:12 ?        00:00:00 asm_vbg2_+ASM1
oracle    4404      1  0 08:12 ?        00:00:00 asm_vmb0_+ASM1
oracle    4397      1  0 08:12 ?        00:00:00 [oks_conn]
oracle    4394      1  0 08:12 ?        00:00:00 [oks_rbld]
oracle    4395      1  0 08:12 ?        00:00:00 [oks_rbld]
oracle    4396      1  0 08:12 ?        00:00:00 [oks_rbld]
root      3296      1  0 08:08 ?        00:00:00 [oks_wkq]
```

Concreet in Oracle oplossingen



- ASMCA
- Enterprise Manager
- ASMCMD

```
ASMCMD> volcreate -G DG01 -s 4G volume1
```

- sqlplus

```
SQL> alter diskgroup DG01 add volume 'volume1' size 4G;
```

- ASMCA
- Enterprise Manager
- mkfs.acfs

```
mkfs.acfs /dev/asm/volume1-201
```

- OCR, indien ACFS gebruikt wordt voor het opslaan een Oracle database home, in verband met afhankelijkheid van de database instance.
- Oracle ACFS Registry. Deze registry bevat alle ACFS filesystems, indien deze geregistreerd worden. Dit zorgt ervoor dat wanneer de resource ora.registry.acfs gestart wordt, de betreffende filesystems gemount worden.

- ASMCA
- Enterprise Manager
- acfsutil

```
[root@rac01 ~]# acfsutil -a /dev/asm/volume1-201 /volume1
```

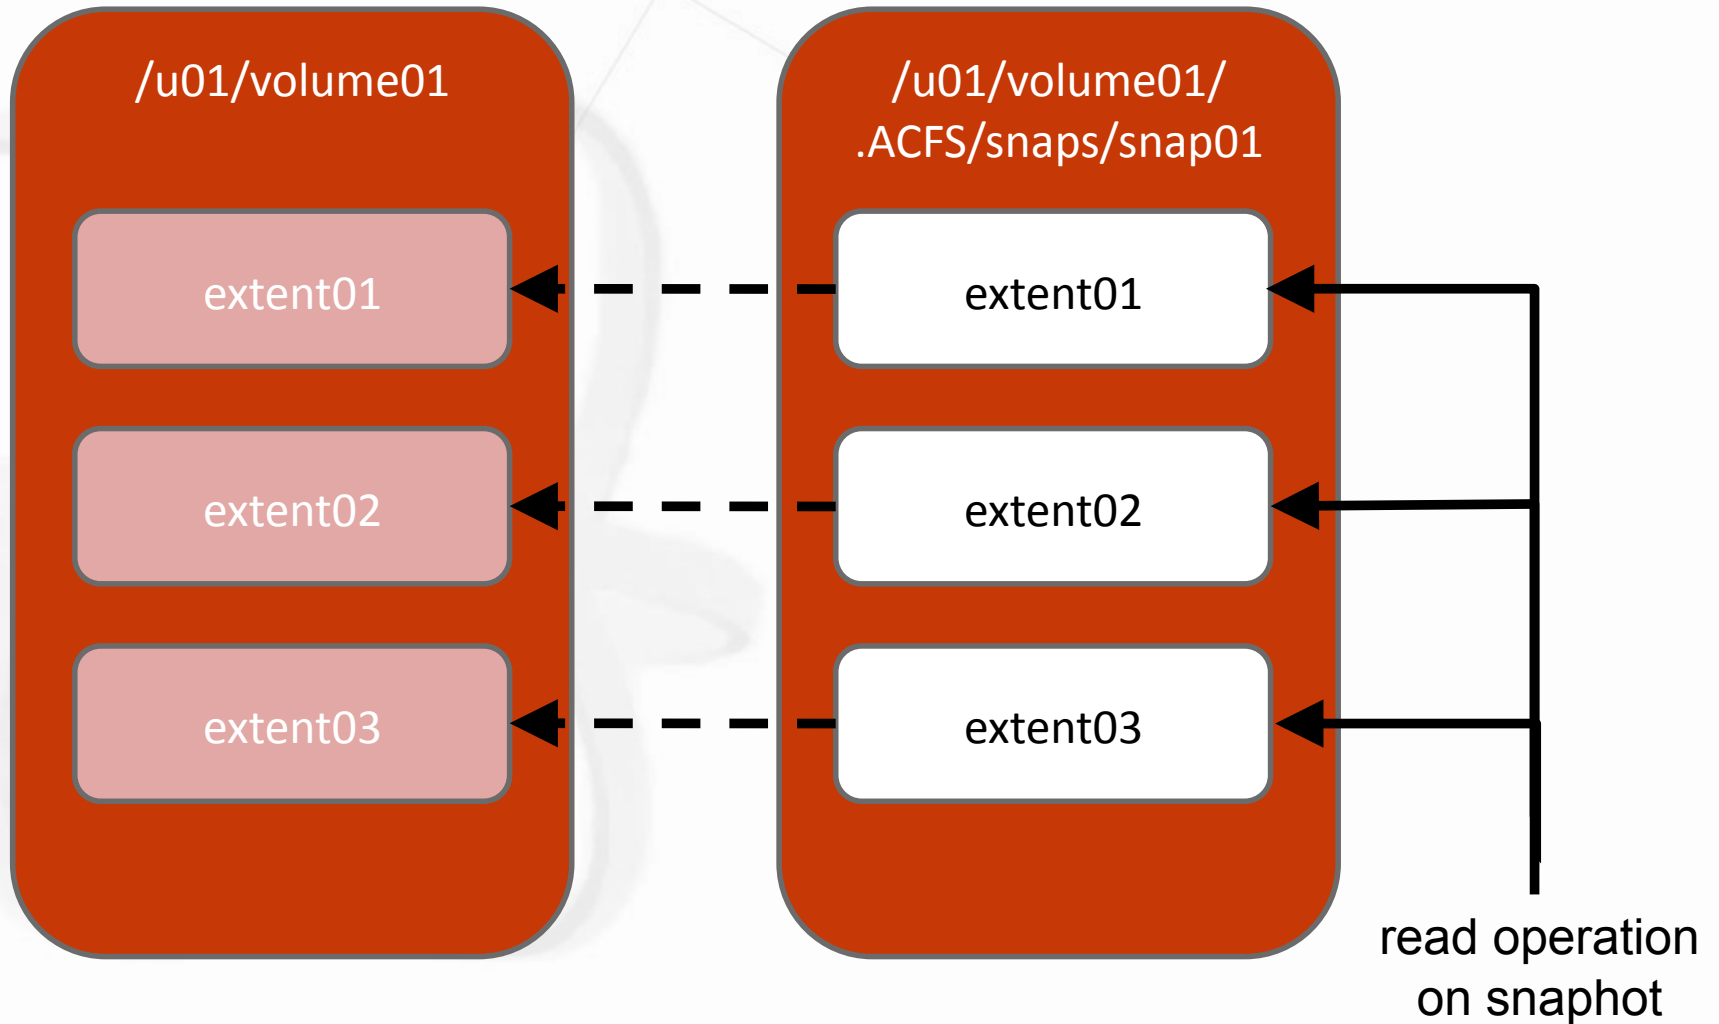
- ASMCA
- Enterprise Manager
- mount.acfs
  - Mounten 1 filesystem  

```
[root@rac01 ~]# mount.acfs /dev/asm/volume1-201 /volume1
```
  - Mounten alle filesystemen in ACFS Registry  

```
[root@rac01 ~]# mount.acfs -o all
```

- Online, read-only, point in time copy
- Copy-on-Write functionality
- Verwijzing naar de storage allocation information
- 63 snapshots per filesystem ondersteund
- Gelijk beschikbaar onder directory `.ACFS/snaps` van het originele filesystem
- Kan gebruikt worden voor het maken van een consistente backup van het filesystem (maak backup van `/u01/volume01/.ACFS/snaps/snap01`)

Concreet in Oracle oplossingen



- Enterprise Manager
- acfsutil

```
[oracle@rac01 ~]acfsutil snap create snap01 /volume01
```

```
[oracle@rac01 ~]acfsutil snap delete snap01 /volume01
```

# Vergroten en verkleinen ACFS volume

Concreet in Oracle oplossingen

- ASMCA
- Enterprise Manager
- acfsutil

```
[oracle@rac01 ~]acfsutil size +2GB -d /dev/asm/  
volum01-201 /u01/volume01
```

```
[oracle@rac01 ~]acfsutil size 3GB -d /dev/asm/  
volume01-201 /u01/volume01
```

```
[oracle@rac01 ~]acfsutil size -2GB -d /dev/asm/  
volume01-201 /u01/volume01
```

- Database logfiles in cluster omgevingen
- Datapump dump files in cluster omgevingen
- Scripts die gebruikt worden op meerdere systemen binnen een cluster omgeving
- Vervanger voor NFS binnen cluster omgevingen
- Shared filesystem voor Weblogic servers

- Eenvoudige manier voor het configureren en beheren van storage voor alle typen files (database, binaries etc.)
- Indien gebruik makend van Oracle Enterprise Linux, ondersteuning gehele stack (OS, ASM, RDBMS, applicatie server) door één leverancier

- In traditionele omgevingen (System Administrators, DBAs) lastig door te voeren
- Extra kennis nodig bij de Oracle DBAs (DBA 2.0)
- Indien zich een probleem voordoet (bug) met ASM, kan dat gevolgen hebben voor de beschikbaarheid van de databases, filesystems in het hele cluster

- Whitepaper ACFS Technical Overview and Deployment Guide
- Oracle Database Storage Administrator's Guide



GRID IT

Concreet in Oracle oplossingen





**[rdbbraber@grid-it.nl](mailto:rdbbraber@grid-it.nl)**