



Introduction to Automatic Storage Management

An overview of ASM, the Oracle 10g
file and volume management
system.



Topics

- Introduction to ASM
- Benefits
- Architecture
- Installation tips
- Starting ASM
- ASM Disks, Diskgroups and Failgroups
- Managing ASM with asmcmd
- ASM from the Database Instance
- ASM and SAN Tips
- Miscellaneous Release 2 enhancements



What is ASM?...

- ASM = *Automatic* Storage Management
 - It does lots of storage-related things for you
- Built-in “database file system”
- Storage for database-related things only
 - Tables, indexes, tempfiles, backup sets, control files, parameter files, redo logs, archive logs, dump sets



...What is ASM?

- Mirrors and stripes data (RAID 1+0)
 - Mirrored stripes
 - Good R/W performance
 - Good availability
 - Loss of disk causes background rebalance operation to remirror stripes on available space
- Works with
 - Single server or RAC environments
 - NAS / SAN or DAS devices
 - Can turn off ASM mirror – use NAS
- Bundled with database – in kernel



Benefits of ASM...

- Simpler administration
 - Manage small number of disk groups
 - Not hundreds / thousands of datafiles
 - ASM balances IO load
 - Dynamic parallel load balancing
 - Helps prevent hot spots
 - IO spread evenly across available disks
 - Control for intrusiveness of rebalance operations
 - See `ASM_POWER_LIMIT`



...Benefits of ASM...

- Performance
 - Raw device performance
 - Without raw device management
 - Oracle TPC tests using ASM
- Redundancy
 - 2 or 3 way mirroring
 - Or use hardware mirroring



...Benefits of ASM

- High Availability
 - Add / remove drives online
 - Move files online
 - Automatic, controllable, online rebalancing
- Bundled with database
 - SE1, SE and EE



Automated S.A.M.E.

- S.A.M.E.
 - “Stripe and Mirror Everything”
- Striping to balance I/O load
 - Good performance
 - Uses optimal stripe size for Oracle datafiles
 - Typically 1MB
 - 128k stripe for control files
- Mirroring for redundancy, availability



Other Points...

- ASM built on top of raw devices
 - Don't use on top of other file systems
 - Except maybe for testing
 - Undoc'd parameter `_asm_allow_only_raw_disks`
- Raw performance w/o raw management
- Can coexist with other file management techniques
 - Can migrate incrementally
 - 10g `DBMS_FILE_TRANSFER` can be used
 - RMAN commands



...Other Points

- ASM files are *not* visible to OS
 - Use RMAN to backup
- 10g “BIGFILES” are designed for ASM
 - Or another striping solution
- OEM provides nice management interface
- ASM complements Oracle9i Managed Files

```
LINUX> show parameter db_create
```

NAME	TYPE	VALUE
db_create_file_dest	string	+GROUP1

OMF points to
disk group



The ASM Instance...

- ***One ASM instance required on server***
 - ***One ASM instance per node in RAC***
- ***A separate instance from DB instance***
 - ***See parameter `INSTANCE_TYPE=+ASM`***
 - ***ASM instance services one or more database instances on a node***
- ***Requires approximately 100MB memory***
 - ***Memory requirements said to be static***



...The ASM Instance...

- Manages diskgroup metadata
 - Diskgroup is group of physical disks / LUNs
- Each physical Oracle server has one ASM instance
 - Used by all database instances on server
 - Each node in RAC cluster has 1 ASM instance



...The ASM Instance

- ***Starts background processes to manage ASM disk metadata***
 - ***RBAL***
 - ***ARBn***
 - ***GMON***
- ***ASM instance has no database***
 - ***No datafiles, control file, log files***
 - ***ASM disk metadata triple-mirrored on ASM disks***
 - **Metadata backed up with database backups**
 - ***No need for separate ASM backups***



ASM Install Tips

- DBCA can
 - Create ASM instance only
 - Even if server already contains DB
 - Create ASM and database instance
 - Create database using existing ASM instance
- Consider putting in ASM separate home
 - Patch separately
- Give ASM raw disks to discover
 - See the ASM_DISKSTRING parameter
- CSS required





ASM Parameters

```
instance_type='asm'  
+asm.asm_diskgroups='DATAFILE_GROUP','FLASHBACK_GROUP'  
asm_diskstring='c:\asmdisks\_file*'  
asm_power_limit=1 # 1 (low), 11 (high)  
processes=55 #set to handle 2 database instances  
large_pool_size=12M  
shared_pool_size=40m  
db_cache_size=24m  
remote_login_passwordfile='exclusive'  
user_dump_dest='C:\oracle\product\10.2.0\admin\+ASM\udump'  
background_dump_dest='C:\oracle\product\10.2.0\admin\+ASM\bdump'  
core_dump_dest='C:\oracle\product\10.2.0\admin\+ASM\cdump'  
compatible='10.2.0.1.0'
```



Starting ASM

- Start ASM first
 - Mounts all diskgroups in ASM_DISKGROUPS parameter
- Shutdown ASM last

Default ASM instance name

```
[oracle@springsteen oracle]$ ORACLE_SID=+ASM
[oracle@springsteen oracle]$ sqlplus / as sysdba
. . .
SQL> startup
ASM instance started

Total System Global Area  100663296 bytes
Fixed Size                 777616 bytes
Variable Size             99885680 bytes
Database Buffers          0 bytes
Redo Buffers               0 bytes
ASM diskgroups mounted
SQL>
```

Must (always) connect as SYSDBA

Mounts disks; no database to open!



ASM Disks

- ASM disks must be discovered
 - Give raw devices to ASM
- Parameter `ASM_DISKSTRING` to discover disks
 - At ASM startup
 - Can limit to specific disks

```
SQL> select name, path, total_mb, free_mb,
2      reads, writes
3      from v$asm_disk;
```

NAME	PATH	TOTAL_MB	FREE_MB
GROUP1_0003	/dev/raw/raw4	34710	29852
GROUP1_0002	/dev/raw/raw3	34710	29852
GROUP1_0001	/dev/raw/raw2	34710	29213
GROUP1_0000	/dev/raw/raw1	32718	28416

ASM instance shows all available disks

Database shows disks in use by that database



ASM Diskgroups

- Like a volume or storage group
- Contain ASM files
- Used by one or more database instances
- `ASM_DISKGROUP` parameter identifies diskgroups mounted
- Striping for balanced IO load across disks
 - File's extents “distributed equally across all ... disks in the diskgroup”*

See the `CREATE DISKGROUP` example later in this lesson



ASM Failgroups

- Failgroups provide ASM (internal) mirroring
- Mirroring options
 - External
 - Use EMC or other storage product
 - Normal
 - 2-way
 - High
 - 3-way



Creating Diskgroup

```
+ASM> create diskgroup group1 normal redundancy
 2 failgroup fgroup1 disk '/dev/raw/raw1', '/dev/raw/raw2'
 3 failgroup fgroup2 disk '/dev/raw/raw3', '/dev/raw/raw4';
```

Diskgroup created.

```
SQL> select * from v$asm_diskgroup;
```

GROUP_NUMBER	NAME	SECTOR_SIZE	BLOCK_SIZE
1	GROUP1	512	4096

ALLOCATION_UNIT_SIZE	STATE	TYPE	TOTAL_MB	FREE_MB
1048576	MOUNTED	NORMAL	136848	136744

```
+ASM> select g.name as group_name, d.name as disk_name, d.path, d.failgro
 2 from v$asm_diskgroup g, v$asm_disk d
 3 where g.group_number = d.group_number
```

GROUP_NAME	DISK_NAME	PATH	FAILGROUP
GROUP1	GROUP1_0003	/dev/raw/raw4	FGROUP2
GROUP1	GROUP1_0002	/dev/raw/raw3	FGROUP2
GROUP1	GROUP1_0001	/dev/raw/raw2	FGROUP1
GROUP1	GROUP1_0000	/dev/raw/raw1	FGROUP1



Altering Diskgroup

- Easily add or remove disks while database open
- Automatic rebalancing starts

```
+ASM> alter diskgroup group1 drop disk GROUP1_0002;
```

```
Diskgroup altered.
```

```
+ASM> select * from v$asm_operation
```

```
2 /
```

GROUP_NUMBER	OPERA	STAT	POWER	ACTUAL	SOFAR	EST_WORK	EST_RATE
1	REBAL	RUN	1	1	181	1275	338
EST_MINUTES							

3							



Referencing Disk Group...

- Creating datafiles follow 9i OMF conventions
 - Location
 - Size
- New file name convention

```
LINUX> create tablespace class datafile size 5m;
Tablespace created.
LINUX> select file_name, bytes from dba_data_files
2      where tablespace_name = 'CLASS';
```

FILE_NAME	BYTES
+GROUP1/linux3/datafile/class.1083.1	5242880

DROP removes
file from ASM
diskgroup

```
LINUX> drop tablespace class;
Tablespace dropped.
```



...Referencing Disk Group

➤ Easy to override OMF diskgroup

Just specify diskgroup name; ASM figures out the rest

```
LINUX> create tablespace class datafile '+group1' size 1m;
```

Tablespace created.

```
LINUX> select file_name, bytes from dba_data_files
  2  where tablespace_name = 'CLASS';
```

FILE_NAME	BYTES
+GROUP1/linux3/datafile/class.1083.3	1048576

```
LINUX> alter tablespace class add datafile '+group1' size 1m;
```

Tablespace altered.

Use diskgroup name in ALTER



ASMCMD

- Command line interface
- Alternative to managing with SQL or OEM
- Supports scripting

Commands similar to Unix

```
[oracle@springsteen oracle]$ ORACLE_SID=+ASM
[oracle@springsteen oracle]$ asmcmd
ASMCMD> ls -l group1/linux3/datafile
```

Type	Redund	Striped	Time	Sys	Name
DATAFILE	MIRROR	COARSE	JAN 05 10:00:00	Y	DAVES_TS.1162.9
DATAFILE	MIRROR	COARSE	JAN 05 10:00:00	Y	DEFAULT_USERS.13
DATAFILE	MIRROR	COARSE	JAN 05 10:00:00	Y	EXAMPLE.269.1
DATAFILE	MIRROR	COARSE	JAN 05 10:00:00	Y	SKILL.1221.7
DATAFILE	MIRROR	COARSE	JAN 05 10:00:00	Y	SYSAUX.257.1
DATAFILE	MIRROR	COARSE	JAN 05 10:00:00	Y	SYSTEM.256.1
DATAFILE	MIRROR	COARSE	JAN 05 10:00:00	Y	UNDOTBS1.258.1
DATAFILE	MIRROR	COARSE	JAN 05 10:00:00	Y	USERS.957.7

```
ASMCMD> du
Used_MB      Mirror_used_MB
6790         13591
```



R2 Enhancements...

- OEM provides ASM migration utility
- DBMS_FILE_TRANSFER
 - Migrate to or from ASM / OS file system
- Backward / forward compatibility
 - 10.2 ASM instance supports 10.1 database
 - 10.1 ASM instance supports 10.2 database
- FTP tool to copy files in / out
- XML DB virtual folder support



...R2 Enhancements...

- New views
 - V\$ASM_DISK_STAT, V\$ASM_DISKGROUP_STAT
 - More efficient than V\$DISK, V\$ASM_DISKGROUPUsable free space
 - Space that's safe to use considering mirroring
- RAC and single-instance DB's can share an ASM instance
- STARTUP displays ASM Cache size



...R2 Enhancements

- SHUTDOWN returns error if database connected
 - ORA-15097
- Wait for rebalance op to complete
 - Request WAIT on ALTER DISKGROUP
 - Can help in scripts
- `ASM_POWER_LIMIT=0`
 - Disables automatic rebalance ops
 - Suspends in-progress ops



ASM and SAN Best Practices

- Use ASM external redundancy
 - Let NAS create mirrors
- Use like disks in one disk group
- Give more than 1 LUN to a disk group
 - Let ASM do the striping
 - No performance hit for double striping (stripe-on-stripe)



ASM and SAN Best Practices

- Use a multi-pathing solution to prevent SPOF on SAN
 - e.g. HP Secure Path, EMC Powerpath
 - Manage device (LUN) with ASM or LVM, not both
- Configure async I/O
- Many white papers on ASM home page
 - oracle.com/technology/products/database/asm



Summary...

- ASM is a built-in file system and volume manager
 - Provides automated S.A.M.E.
 - Easier disk management
 - Use instead of other file systems
- ASM instance manages diskgroup metadata
 - One per physical server
- Install ASM with DBCA
 - Put in separate home
- At startup:
 - ASM_DISKSTRING controls disk discovery
 - ASM_DISKGROUP names diskgroups to be mounted



...Summary

- Diskgroup is the main object
 - Define disks in the group
 - Define failgroups for redundancy
- Manage ASM with SQL or asmcmd
 - asmcmd provides Unix-like commands
- Backup datafiles with RMAN
 - Datafiles not visible to the OS
- Can mix ASM files with OS or raw
- R2 provides DBMS_FILE_TRANSFER for migration
- Works with NAS/SAN, DAS
- Works for single-instance or RAC databases
- Nice feature!



Helpful Resources

- Class from SkillBuilders, free tutorials
 - See www.skillbuilders.com
- ASM Home Page
 - oracle.com/technology/products/database/asm
 - Arup Nanda's articles particularly helpful
- www.idevelopment.info
 - Jeff Hunter's home page
 - Cookbook how-to examples
- Oracle10g Administrator's Guide
 - Chapter 12
- Oracle10g Installation Guide for your platform
 - Recommended reading to get up and running



Q & A



Thank You

- ***Thank you for coming to the presentation!***
- ***Thanks to SEOUC for inviting me to speak!***