



Confidence in the connected world.



INTEGRATING ORACLE
AUTOMATIC STORAGE
MANAGEMENT WITH
VERITAS STORAGE
FOUNDATION FOR ORACLE
RAC

A Joint Paper by Oracle and
Symantec

Contents

Content

- Introduction 3
- Veritas Storage Foundation 3
- Oracle Automatic Storage Management 3
- Supported Configurations 4
- Support Statements for Veritas Storage Foundation for Oracle RAC 5
- How to configure initial ASM diskgroup with CVM 5
- How to add additional storage to an ASM diskgroup with CVM 7
- Recommended Dos and Don'ts 8
- Conclusion 9

Introduction

With the release of Oracle Database 10g, Oracle introduced Automatic Storage Management (ASM), which enables database or systems administrators to manage a pool of storage and automate Oracle database file management. Veritas Storage Foundation provides storage management and dynamic multi-pathing for Oracle databases and other critical enterprise workloads. Oracle and Symantec believe that customers can combine ASM and Veritas Storage Foundation to take advantage of the benefits of both storage management offerings. Symantec and Oracle have worked jointly to present our customers with best practices for storage configurations, layout guidelines and interoperability of the joint solution.

Veritas Storage Foundation

Veritas Storage Foundation is a 5th generation suite of software that provides enterprise-class storage and cluster management for large mission critical applications and databases, including Oracle Database software and Oracle RAC with ASM. Fundamental components of Veritas Storage Foundation include: Veritas Volume Manager (VxVM) and Cluster Volume Manager (CVM); Veritas File System (VxFS) and Cluster File System (CFS); and Veritas Cluster Server (VCS).

A key element of Veritas Storage Foundation is Dynamic Multi-pathing (DMP), which provides multiple parallel IO paths between servers and storage, to improve storage availability and scalable performance. Customers who wish to run ASM with Veritas Storage Foundation can benefit from the DMP capabilities built into Veritas Storage Foundation. DMP is supported with all major storage hardware vendors, which allows users to confidently standardize on Storage Foundation to take full advantage of common tools across storage and server infrastructure and reduce their dependency on point storage management tools from hardware vendors.

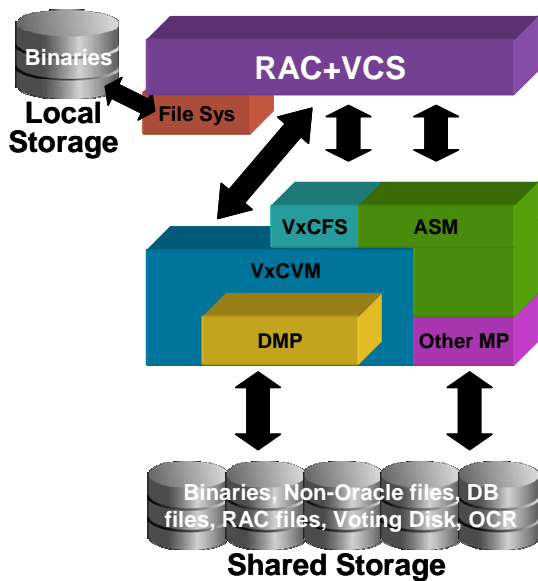
Oracle Automatic Storage Management

ASM is a feature of Oracle Database, first introduced in 10g, that provides integrated cluster file system and volume management capabilities for Oracle database files. ASM lowers the Oracle database storage total cost of ownership and increases storage utilization while improving performance and availability for the Oracle database. ASM may be configured with VxVM or CVM to take advantage of the volume management and DMP Multi-pathing capabilities in Veritas Storage Foundation.

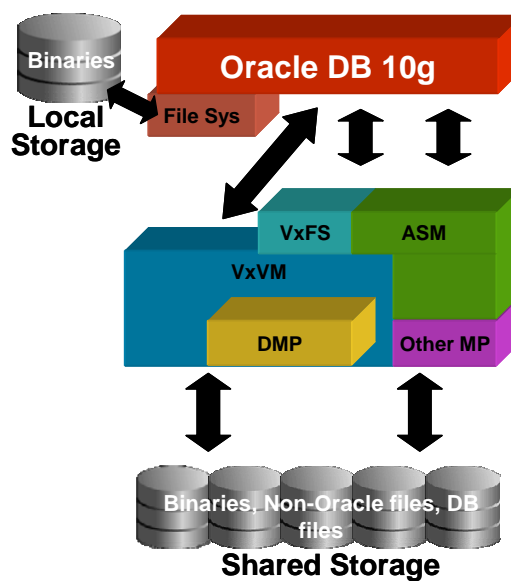
As the foundation of Oracle storage grid, ASM brings significant key values to Oracle Database platforms. ASM reduces management complexity, simplifies and automates database file management, and increases storage utilization through database consolidation. . ASM also supports dynamic and on-line storage reconfigurations in both single instance and Oracle RAC database configurations..

Supported Configurations

Clustered (RAC) Configuration



Non-clustered Configuration



Two Options for Storage Management

Shown above are supported configurations for Oracle ASM with Veritas Storage Foundation

The diagrams above show ASM and Veritas Storage Foundation working together in both clustered (RAC) and non-clustered (single instance) configurations. On the left, clustered versions of the Veritas Volume Manager (VxCVM) and Veritas File System (VxCFS) are shown in support of Oracle RAC. On the right, non-clustered versions (VxVM & VxFS) are shown in support of non-RAC databases.

In both configurations above, binaries, clusterware files and non-Oracle files can be stored in a Veritas File System (Cluster File System for RAC), or on a locally mounted file system on each node. Database files are stored and managed by ASM, which can then leverage Storage Foundation Dynamic Multi-pathing (DMP) through Veritas Volume Manager (clustered for RAC, left). The diagrams also show the possibility of using ASM with other multi-pathing solutions, which are supported by Oracle. However, for any given host or RAC cluster only one multi-pathing solution should be active and managing the underlying devices.

The Oracle RAC Voting Disk and the Oracle Cluster Registry (OCR) (shown at the bottom of the left diagram) are not managed by ASM in Oracle RAC 10g, but can be stored in Veritas Volume Manager and/or Veritas File System (VxCVM or VxCFS), or using raw devices.

Technical Best Practices

Support Statements for Veritas Storage Foundation for Oracle RAC

Symantec will support ASM configurations running with the Storage Foundation under the following conditions:

- Support will be provided for ASM on any Storage Foundation products that are supported with Oracle 10.2.0.3 and later. *Check the Symantec Support site for a listing of supported configurations - http://ftp.support.veritas.com/pub/support/products/DBE_Advanced_Cluster_for_Oracle_RAC/280186.pdf*
- ASM support is provided for configurations that are running over raw logical volumes (VxVM and CVM)
- Platforms supported/certified with SFRAC include AIX, HP-UX and Solaris (Sparc & x64)

How to configure initial ASM diskgroup with CVM

- 1) The example below will illustrate how to add four disks (which is the Oracle recommended best practice) to the diskgroup. This needs to be done only for one node in the case of RAC. As root configure a CVM diskgroup using the following commands :

```
racnode1@root# vxdiskunsetup -C c2t1d13
racnode1@root# vxdisksetup -i c2t1d13
racnode1@root# vxdiskunsetup -C c2t1d14
racnode1@root# vxdisksetup -i c2t1d14
racnode1@root# vxdiskunsetup -C c2t1d15
racnode1@root# vxdisksetup -i c2t1d15
racnode1@root# vxdiskunsetup -C c2t1d16
racnode1@root# vxdisksetup -i c2t1d16
racnode1@root# vxdbg -s init ora_asm_dg c2t1d13 c2t1d14 c2t1d15 c2t1d16
racnode1@root# vxassist -g ora_asm_dg make ora_asm_vol 2000M
racnode1@root# vxedit -g ora_asm_dg \
    set group=dba user=oracle mode=660 ora_asm_vol
```

- 2) As oracle user configure ASM using either Database Configuration Assistant (DBCA) or Enterprise Manager (EM). If you do not use either GUI, you will need to manually create the ASM instances and disk groups. Please check Section 8 (Recommended Do's and don'ts) before starting this process:

- a. DBCA Procedure
 - i. Choose Configure ASM Instance option
 - ii. Create a default ASM diskgroup (Eg: DATADG)
 - iii. Choose appropriate redundancy, for example, choose external if you plan to leverage hardware RAID
 - iv. Change the Discovery String to “/dev/vx/rdisk/ora_asm_dg/ora_asm*”

Technical Best Practices

- v. This should list all the VxVM volumes created and available for ASM.
- vi. Select the disks

At the completion of this step, DBCA starts the ASM instances and mounts the diskgroup on all nodes of the cluster.

Verify the diskgroup status using the following:

```
SQL>select name, state, type, total_mb, free_mb from v$asm_diskgroup.
```

When using DBCA to create a database, select ASM for storage

- b. Manually create the ASM instances and diskgroups :
 - i. Update the init.ora file parameter asm_diskstring with the appropriate diskstring locator:
"/dev/vx/rdisk/ora_asm_dg/ora_asm*"
 - ii. Update init.ora to set up instance_type to ASM
 - iii. Startup the ASM instance specifying the init.ora file
 - iv. SQL> create diskgroup DATADG external redundancy disk

```
"/dev/vx/rdisk/ora_asm_dg/ora_asm_vol1",/dev/vx/rdisk/ora_asm_dg/ora_asm_vol2',/dev/vx/rdisk/o  
ra_asm_dg/ora_asm_vol3',/dev/vx/rdisk/ora_asm_dg/ora_asm_vol4';
```

- v. Verify the diskgroup creation using the following:

```
SQL> select name, state, type, total_mb, free_mb from v$asm_diskgroup;
```

NAME	STATE	TYPE	TOTAL_MB	FREE_MB
----	-----	----	-----	-----
DATADG	MOUNTED	EXTERN	34512	34101

```
SQL> select name, path, mode_status, state, disk_number from v$asm_disk
```

NAME	PATH	MODE_ST	STATE	DISK_NUMBER
-----	-----	-----	-----	-----
DATA_0000	/dev/vx/rdisk/ora_asm_dg/ora_asm_vol1	ONLINE	NORMAL	0

The above manual steps are required to be done only for one node in the case of Oracle RAC database.

Also the diskgroups need to be mounted manually on the remote nodes (and update ASM_DISKGROUPS if using a pfile instead of an spfile)

Technical Best Practices

3) The following are other useful views:

- a. v\$asm_disk
- b. v\$asm_diskgroup
- c. v\$asm_operation
- d. v\$asm_file

4) As user oracle use the following commands for normal startup/shutdown of the ASM instance.

This can be run from any node of the cluster.

For startup

```
racnode1@oracle> srvctl start asm -n node1
```

```
racnode1@oracle> srvctl start asm -n node2
```

For shutdown

```
racnode1@oracle> srvctl stop asm -n node1
```

Note: The ASM instance cannot be shut down when RDBMS instances that are using ASM are up and running. Shut down all RDBMS instances that are accessing ASM files before shutting down the ASM instance.

How to add additional storage to an ASM diskgroup with CVM

1. As root user to add additional storage use the following commands:

```
racnode1@root# vxdisksetup -C c2t1d17
```

```
racnode1@root# vxdisksetup -i c2t1d17
```

```
racnode1@root# vxdg -g ora_asm_dg adddisk c2t1d17
```

```
racnode1@root# vxassist -g ora_asm_dg make ora_asm_vol5 2000M
```

```
c2t1d17
```

```
racnode1@root# vxedit -g ora_asm_dg set group=dba user=oracle mode=660 ora_asm_vol5
```

2. To add a disk to an existing ASM diskgroup

```
SQL>alter diskgroup ora_asm_dg add disk '/dev/vx/rdisk/ora_asm_dg/ora_asm_vol5';
```

This topic should have a process flow visual aid.

Recommended Dos and Don'ts

1. Do not create any file systems on VxVM volumes which will be used with ASM. Doing so may cause data corruption and/or result in headers getting overwritten.
2. Do not include volumes with file systems or raw volumes in the same diskgroup being assigned to ASM. Veritas raw volumes for ASM should be used exclusively by ASM.
3. Do not enable ODM when implementing ASM. ODM is only appropriate for Datafiles residing on the Veritas Filesystem (VxFS and/or CFS). If the user chooses to use Veritas Cluster File System for any other purpose, such as archive logs, ODM is not applicable.
4. Do not stripe or mirror using Veritas Volume Manager. ASM implements an additional layer of striping above hardware RAID configuration.
5. Create two ASM diskgroups (DGs).

Diskgroup1: Datafiles, one set of Control Files, one set of Redo Logs

Diskgroup2: Flash Recovery Area: Flashback rewind logs, Archive Logs, set two of Redo Logs, and one of Control Files, RMAN backups.

Please refer to ASM best practices paper below for more details:

http://www.oracle.com/technology/products/database/asm/pdf/asm_10gr2_bestpractices%2005-07.pdf

6. A VxVM volume should consist of a single disk/LUN as presented to the operating system. Do not stripe, mirror, or concatenate the volumes.
7. Do not change ownership or permissions on the DMP meta nodes , as this could prevent access to storage devices under DMP. (example: /dev/vx/dmp/devicename and /dev/vx/rdmp/devicename) .
8. Do not attempt to have ASM directly address the DMP meta nodes. All access to DMP meta nodes should be through the appropriate Veritas Volume Manager (clustered or non-clustered)
9. Do not combine ASM and CVM/CFS volumes within the same CVM diskgroup. Create separate diskgroups for ASM and CFS/CVM volumes.
10. Combining DMP with multipathing software from another vendor on the same system is not allowed and not supported.
11. Coordinator disks for I/O fencing cannot be configured under ASM. The coordinator disks can be configured only under the Veritas SFRAC stack.
12. Do not attempt VxVM/CVM virtualization operations, i.e. volume resize, concatenation with LUNs that are under ASM control.

13. Volumes presented to a particular ASM diskgroup should be of the same size and performance characteristics.
14. Veritas Storage Foundation commands vxresize or vxassist are not required in this configuration. [Optional] Lower level, under-the-covers description of how things work. Content focus is on bits the user would not touch.

Conclusion

Storage Foundation Volume Manager operations will co-exist with Oracle ASM without any performance penalties. Veritas Storage Foundation Dynamic Multi-Pathing provides seamless multi-path failover capabilities and is complementary to ASM for both high availability and performance. Customers who wish to use Oracle ASM and Veritas Storage Foundation together should be careful to follow the procedures outlined in the paper to ensure a successful installation.

i [Dynamic MultiPathing \(DMP\): http://eval.symantec.com/mktginfo/enterprise/white_papers/ent-whitepaper_vsf_5.0_dynamic_multi-pathing_05-2007.en-us.pdf](http://eval.symantec.com/mktginfo/enterprise/white_papers/ent-whitepaper_vsf_5.0_dynamic_multi-pathing_05-2007.en-us.pdf)

ii A DMP metanode represents a metadvice, a Veritas Volume Manager abstraction that corresponds to a disk or LUN and all the I/O paths on which it can be accessed.

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world. The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

Copyright © 2007 Symantec Corporation. All rights reserved. Symantec and the Symantec logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

10/07 xxxxxxxx