



# Oracle Database 10g Automatic Storage Management with HP StorageWorks Arrays on HP-UX

Providing best practices and customer guidance for HP and Oracle environments

## Table of contents

<b>Introduction</b> .....	<b>2</b>
<b>Overview of Oracle Database 10g ASM</b> .....	<b>2</b>
Using HP StorageWorks Secure Path with Oracle Database 10g ASM on HP-UX .....	4
<b>Test environment</b> .....	<b>4</b>
Configuring ASM with HP StorageWorks EVA and XP Arrays on HP-UX .....	5
Creating an Oracle RDBMS utilizing ASM for the storage .....	9
Configuring archive logs to utilize ASM .....	10
Adding a LUN (disk) to an ASM disk group .....	10
Path fail-over tests for StorageWorks EVA and XP Arrays on HP-UX .....	10
<b>Best practices for using Oracle ASM with HP StorageWorks on HP-UX</b> .....	<b>11</b>
<b>LVM or ASM striping in a single EVA disk group vs. multiple EVA disk groups</b> .....	<b>11</b>
<b>Oracle ASM and HP Logical Volume Manager (LVM)</b> .....	<b>11</b>
<b>Conclusion</b> .....	<b>12</b>

# Introduction

HP and Oracle® have shared a long and productive relationship for over 25 years, working together to help our customers maximize the business value of their database solutions and support ever-increasing volumes of data. Today, 37 percent of Oracle database customers use HP systems. The reason: These customers know they can count on HP to deliver the performance, availability, and scalability they need to meet the demands of a dynamic and highly competitive marketplace.

HP StorageWorks arrays provide the ideal platform for Oracle 10g databases. As a leader in the storage arena, HP has a keen understanding of the critical issues Oracle database administrators (DBAs) face today, including disk optimization, cloning of production volumes for testing, and shrinking backup windows. HP StorageWorks arrays address these issues by offering:

- Extensive scalability with very attractive price performance
- Intelligent management of multiple arrays and hundreds of terabytes from a single, centralized location
- High throughput and zero downtime, with no single point of failure and rapid disaster recovery
- A compact footprint to conserve valuable floor space

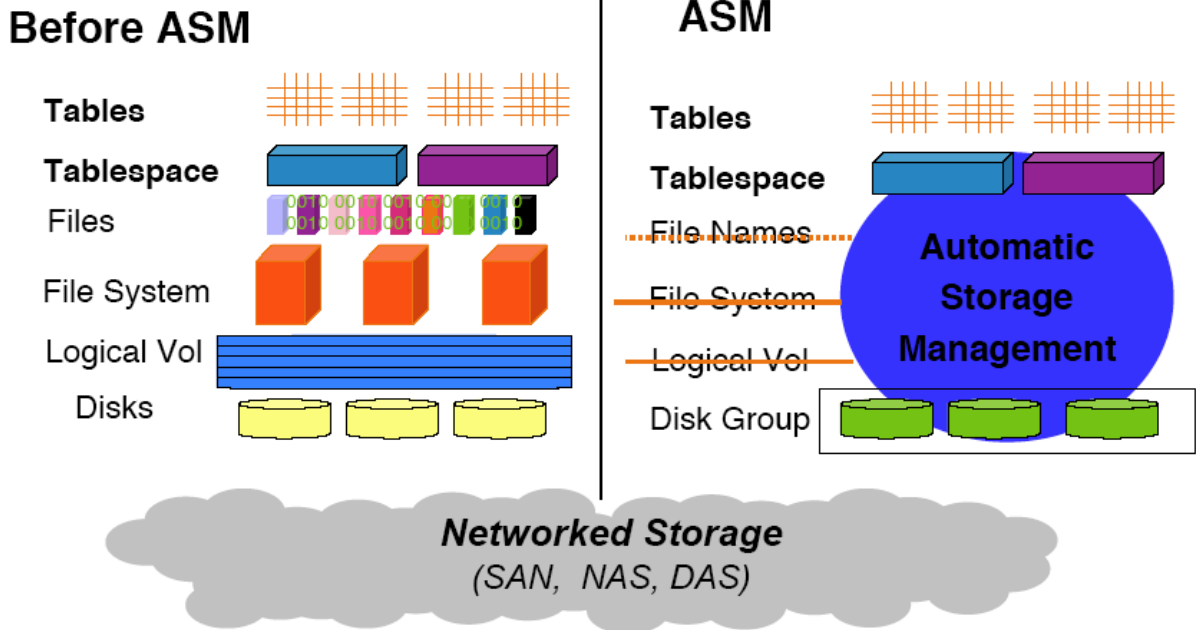
HP and Oracle have conducted extensive testing using the HP StorageWorks Enterprise Virtual Array (EVA) and XP Arrays with Oracle Database 10g Automatic Storage Management (ASM) on HP-UX demonstrating that Oracle ASM can optimize workloads and provide best performance. ASM implemented internally the Stripe and Mirror Everything (SAME) storage database layout methodology from Oracle. In this paper, HP and Oracle discuss how HP StorageWorks arrays are used with ASM. We then provide details of the lab tests, including system configurations, test descriptions, results, and best practices recommendations.

## Overview of Oracle Database 10g ASM

The ASM feature in Oracle Database 10g is a database file system that provides cluster file system and volume manager capabilities integrated into the Oracle database 10g kernel. It is applicable only to Oracle database files.

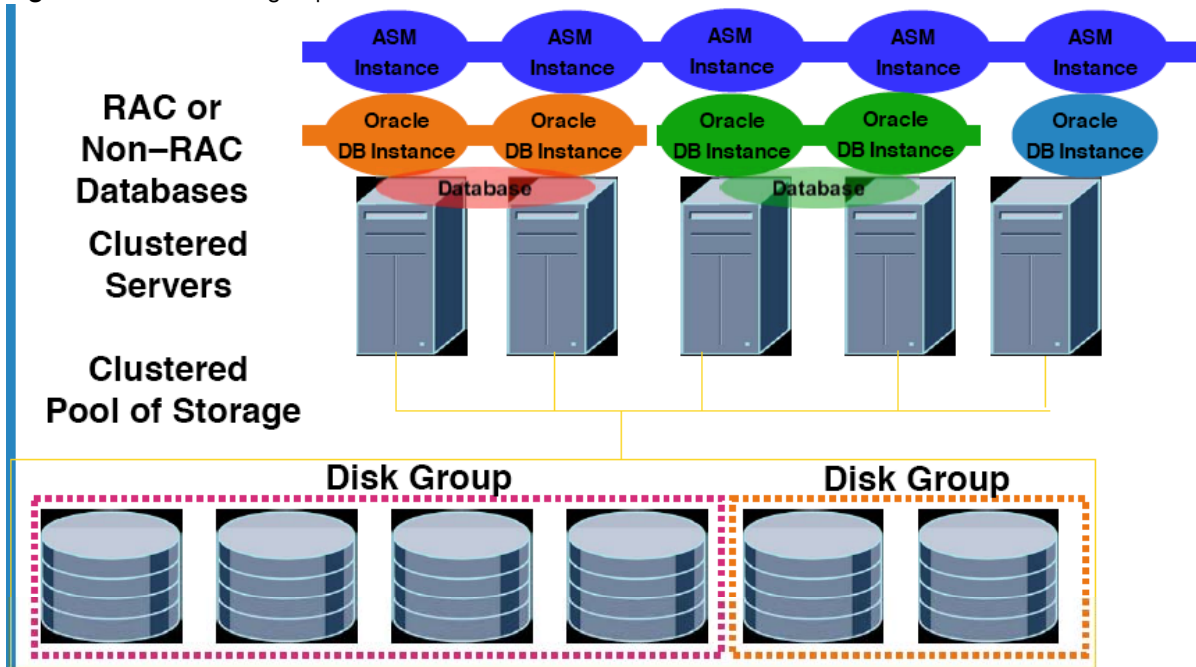
HP with Oracle Database 10g ASM leverages HP storage and Oracle database technologies along with blended best practices and ongoing joint development to provide an efficient IT environment.

Figure 1 illustrates the efficiencies gained with this solution.



ASM disk groups, as shown above, provide a pool of disks that can be managed as a logical unit. By dividing total disk space into uniform-sized megabyte units, ASM spreads each Oracle file evenly across all disks in a disk group. ASM can then automatically rebalance the disk group online whenever storage configuration changes. Oracle only moves the necessary amount of data to guarantee equal distribution for any given disk group.

Figure 2 illustrates disk groups in the ASM architecture.



ASM also provides data mirroring. Choices for disk group redundancy include:

- External: defers to hardware protection
- Normal: 2-way mirroring
- High: 3-way mirroring

In addition, ASM failure groups can be created, providing a set of disks sharing a common resource whose failure needs to be tolerated. Redundant copies of an extent are stored in separate failure groups. ASM failure groups can be assigned by DBAs or automatically by ASM.

### **Using HP StorageWorks Secure Path with Oracle Database 10g ASM on HP-UX**

HP Secure Path is server-based, multipathing software that enhances the HP StorageWorks disk arrays. HP Secure Path is supported by several operating systems but is not part of the operating system.

HP Secure Path bundle contains two drivers, one for Active/Passive model based arrays such as the EVA3000 and EVA5000 and one for Active/Active arrays like the EVA3000/4000/5000/6000/8000 and XP family. The EVA3000 and EVA5000 come in two flavors, Active/Active and Active/Passive. This is purely firmware based.

Secure Path provides the following capabilities:

- Presents a single virtual path to the host for the Active/Passive arrays and expose all the I/O paths for the Active/Active arrays
- Provides a high-availability solution for HP StorageWorks arrays, preventing any single point of failure
- Provides automatic fail-over and fail-back
- Increases I/O throughput by providing different load-balancing policies
- Detects path failures and seamlessly reroutes the I/O without affecting the application

Secure Path compliments ASM by providing a single virtual path for Active/Passive arrays or the preferred path for Active/Active arrays for ASM while still providing transparent path fail-over. The user can choose any active path as a preferred path in the case of Active/Active arrays.

## **Test environment**

A test environment was created with DSS-like workload to simulate load on the database server, running eight different queries in parallel. The objective was to run the server at ~80 percent capacity (CPU / I/O, whichever is higher). Please note load balancing was only activated for the EVA array.

The test environment was configured with the following:

- One HP Integrity rx4640 server with 2 CPUs
- Oracle 10gR1
- HP Secure Path A.3.0D.00F.00F
- HP StorageWorks XP12000 disk array
- 3 HBAs connected to an FC 2/16 switch (XP)
- HP-UX 11.23
- Oracle 10gR1 ASM
- HP StorageWorks EVA5000
- 2 HBAs connected to an FC 2/16 switch (EVA)

### Configuring ASM with HP StorageWorks EVA and XP arrays on HP-UX

For StorageWorks EVA and XP arrays, install and configure HP Secure Path using the supplied installation script (*install.sh*). Alternate paths for path failover is provided through Secure Path to compliment ASM, which is designed to address a single path. The recommendation for ASM is to enable path fail-over via SecurePath. There is no configuration file to edit with SecurePath.

Determine the "access device files" for ASM under control by Secure Path for Active/Passive arrays like the EVA3000/5000 as follows:

```
# spmgr display
Server:          demo.hp.com          Report Created: Fri, Apr 15 15:28:32 2005
Command:        spmgr display
=====
Storage:                5000-1FE1-5004-1730
Load                   On             Auto-restore: Off   Balance Policy: Least Service Time
Balance:
Path                   On             Verify Interval: 30
Verify:
HBAs:                 td1             td3
Controller:          P5849E1AAQ6024,
                    Operational
                    P5849E1AAQ603Y,
                    Operational
Devices:              c58t0d0 c58t0d1   c58t0d2   c58t0d3   c58t0d4   c58t0d5   c58t0d6
TGT/LUN               Device      WWLUN_ID      H/W_Path      #_Paths
0/0                   c58t0d0    6005-08B4-0010-0924-0000-5000-01DC-
                    0000
                    255/255/0/0.0
Controller            Path_Instance  HBA          Preferred?    Path_Status
P5849E1AAQ6024
                    c44t0d1      td1          no           Standby
                    c48t0d1      td1          no           Standby
                    c36t0d1      td3          no           Standby
                    c40t0d1      td3          no           Standby
Controller            Path_Instance  HBA          Preferred?    Path_Status
P5849E1AAQ603Y
                    c46t0d1      td1          YES          Active
                    c51t0d1      td1          YES          Active
                    c38t0d1      td3          YES          Active
                    c42t0d1      td3          YES          Active
```

You can notice that SecurePath generated a virtual path `"/dev/rdisk/c58t0d0"` and associated it with eight physical paths `-- /dev/rdisk/c44t0d1, /dev/rdisk/c48t0d1, /dev/rdisk/c36t0d1, /dev/rdisk/c40t0d1, /dev/rdisk/c46t0d1, /dev/rdisk/c51t0d1, /dev/rdisk/c38t0d1, /dev/rdisk/c42t0d1 --`

The disk "/dev/rdisk/c58t0d0" is an ASM disk candidate. In case of a path failure, the I/O will be re-routed to the remaining paths.

Determine the "access device files" for ASM under control by Secure Path for Active/Active arrays like the EVA3000/4000/5000/6000/8000 and the XP family as follows:

```
# autopath display all
HPswsp Version      :      A.3.0D.00F.00F
Details for path    :      /dev/dsk/c8t0d0
Load Balance Policy :      No Load Balance

=====
Device Paths        Device Status
=====
/dev/dsk/c8t0d0     Active
/dev/dsk/c10t0d0    Active
/dev/dsk/c14t0d0    Active
/dev/dsk/c49t0d0    Active
/dev/dsk/c54t0d0    Active
/dev/dsk/c56t0d0    Active
=====
Details for path    :      /dev/dsk/c8t0d1
...
```

In this example, -- /dev/dsk/c8t0d0, /dev/dsk/c10t0d0, /dev/dsk/c14t0d0, /dev/dsk/c49t0d0, /dev/dsk/c54t0d0, /dev/dsk/c56t0d0-- refer to the same physical device. Any of these devices can be picked as an ASM disk, therefore will become the preferred path. You can observe that all paths are disclosed. If the chosen path failed, then SecurePath will re-route I/O to the other available paths.

Note that RAW / character device files must be used for Oracle ASM, e.g., /dev/rdisk/.... For Oracle to use these devices, they should be owned by the Oracle user:

```
Chown oracle /dev/rdisk/c58t*
```

Create an Oracle parameter file for ASM. For example:

```
# Oracle ASM init.ora parameters
*.background_dump_dest=/oracle/RDBMS/10.1.0/Server/admin/+ASM/bdump'
*.core_dump_dest=/oracle/RDBMS/10.1.0/Server/admin/+ASM/cdump'
*.user_dump_dest=/oracle/RDBMS/10.1.0/Server/admin/+ASM/udump'
*.instance_type='asm'
*.asm_diskgroups=EVDG,XPDG
*.large_pool_size=12M
*.asm_diskstring='/dev/rdsk/c58t0d[012345] ',
'/dev/rdsk/c8t0d*'
*.remote_login_passwordfile='SHARED'
```

Verify that Oracle ASM detects all the devices:

```
oracle > export ORACLE_SID=+ASM
oracle > sqlplus "/ as sysdba"
SQL*Plus: Release 10.1.0.3.0 -. . .
. . .
SQL > STARTUP NOMOUNT;
ASM instance started
Total System Global Area 130023424 bytes
. . .
SQL > SELECT PATH FROM V$ASM_DISK;
PATH
-----
/dev/rds/c58t0d0 /dev/rds/c58t0d1 /dev/rds/c58t0d2 /dev/rds/c58t0d3 /dev/rds/c58t0d4
```

. . .

Create the ASM disk group for the StorageWorks XP or EVA array as follows:

### StorageWorks EVA

```
SQL> CREATE DISKGROUP EVDG EXTERNAL REDUNDANCY DISK '/dev/rds/c58t0d[012345]';
```

Diskgroup created.

```
SQL> SELECT * FROM V$ASM_DISKGROUP WHERE NAME = 'EVDG';
```

<u>GROUP_</u>		<u>SECTOR_</u>		<u>ALLOCATION</u>					
<u>NUMBER</u>	<u>NAME</u>	<u>SIZE</u>	<u>BLOCK_SIZE</u>	<u>UNIT_SIZE</u>	<u>STATE</u>	<u>TYPE</u>	<u>TOTAL_MB</u>	<u>FREE_MB</u>	
1	EVDG	1024	4096	1048576	Mounted	Extern	61440	61380	

### StorageWorks XP

```
SQL> CREATE DISKGROUP XPDG EXTERNAL REDUNDANCY DISK '/dev/rds/c8t0d*';
```

Diskgroup created.

```
SQL> SELECT * FROM V$ASM_DISKGROUP WHERE NAME = 'XPDG';
```

<u>GROUP_</u>		<u>SECTOR_</u>		<u>ALLOCATION</u>					
<u>NUMBER</u>	<u>NAME</u>	<u>SIZE</u>	<u>BLOCK_SIZE</u>	<u>UNIT_SIZE</u>	<u>STATE</u>	<u>TYPE</u>	<u>TOTAL_MB</u>	<u>FREE_MB</u>	
2	XPDG	1024	4096	1048576	Mounted	Extern	56336	56272	

Note that external redundancy is used because the disk groups are created on fault-tolerant arrays.

Verify that all the ASM disks were used and initialized by ASM for the disk group as follows:

#### StorageWorks EVA

```
SQL> SELECT NAME,DISK_NUMBER,FREE_MB,TOTAL_MB FROM V$ASM_DISK WHERE
2 GROUP_NUMBER =
3 (SELECT GROUP_NUMBER FROM V$ASM_DISKGROUP WHERE NAME = 'EVDG')
4 ORDER BY DISK_NUMBER;
```

<u>NAME</u>	<u>DISK_NUMBER</u>	<u>FREE_MB</u>	<u>TOTAL_MB</u>
EVDG_0000	0	10231	10240
EVDG_0001	1	10225	10240
EVDG_0002	2	10233	10240
EVDG_0003	3	10229	10240
EVDG_0004	4	10229	10240
EVDG_0005	5	10233	10240

#### StorageWorks XP

```
SQL> SELECT NAME,DISK_NUMBER,FREE_MB,TOTAL_MB FROM V$ASM_DISK WHERE
2 GROUP_NUMBER =
3 (SELECT GROUP_NUMBER FROM V$ASM_DISKGROUP WHERE NAME = 'XPDG')
4 ORDER BY DISK_NUMBER;
```

<u>NAME</u>	<u>DISK_NUMBER</u>	<u>FREE_MB</u>	<u>TOTAL_MB</u>
XPDG_0000	0	7034	7042
XPDG_0001	1	7030	7042
XPDG_0002	2	7035	7042
XPDG_0003	3	7036	7042
XPDG_0004	4	7033	7042
XPDG_0005	5	7034	7042
XPDG_0006	6	7034	7042
XPDG_0007	7	7036	7042



If the disk group is listed in the init+ASM.ora parameter file, ASM will automatically mount the disk group on startup:

```
SQL> SHUTDOWN;  
ASM diskgroups dismounted  
ASM instance shutdown  
SQL> STARTUP;  
ASM instance started
```

```
Total System Global Area 130023424 bytes  
Fixed Size 1296616 bytes  
Variable Size 128726808 bytes  
Database Buffers 0 bytes  
Redo Buffers 0 bytes
```

ASM diskgroups mounted

```
SQL> SELECT * FROM V$ASM_DISKGROUP ORDER BY GROUP_NUMBER;
```

<u>GROUP_</u> <u>NUMBER</u>	<u>NAME</u>	<u>SECTOR_</u> <u>SIZE</u>	<u>BLOCK_</u> <u>SIZE</u>	<u>ALLOCATION</u> <u>UNIT_SIZE</u>	<u>STATE</u>	<u>TYPE</u>	<u>TOTAL_MB</u>	<u>FREE_MB</u>
1	EVDG	1024	4096	1048576	Mounted	Extern	61440	29243
2	XPDG	1024	4096	1048576	Mounted	Extern	56336	56272

### Creating an Oracle RDBMS utilizing ASM for the storage

The special destination of '+<diskgroup>' is used to create files inside ASM from a database perspective. REDO log files will also be created on the target destination. For example, the *init.ora* might be as follows:

```
...  
db_domain=cup.hp.com  
db_name=ASMEV  
db_create_file_dest=+EVDG  
  
compatible=10.1.0.2.0 control_files=+EVDG/asmev/controlfile/current.257.3,  
  
+EVDG/asmev/controlfile/current.256.3 . . .
```

Note that if the DBA did not supply the control file names during database creation, they can be obtained from the V\$PARAMETER view as in the following example:

```
SELECT CONCAT('CONTROL_FILES= ',  
CONCAT(REPLACE(VALUE, ' ', ' '), ' '))  
CTL_FILES FROM V$PARAMETER WHERE NAME  
='CONTROL_FILES';
```

A much simpler way is to use the following command:

```
SQL> show parameter control_files
```

To create tablespaces, the same special destination is used:

```
SQL> CREATE TABLESPACE TS_ODS_FACT
2 DATAFILE '+EVDG'
3 EXTENT MANAGEMENT LOCAL UNIFORM NOLOGGING;
```

The V\$ASM\_ALIAS and V\$ASM\_FILE views can be used to get a list of files and sizes stored inside ASM. Oracle will create ASM aliases for the files listed in the Oracle file map if supplied.

```
SQL> SELECT A.NAME, F.BYTES, F.SPACE, F.TYPE FROM
V$ASM_ALIAS A,V$ASM_FILE F WHERE A.GROUP_NUMBER =
F.GROUP_NUMBER AND A.FILE_NUMBER = F.FILE_NUMBER;
```

NAME	BYTES	SPACE	TYPE
Current.256.3	2899968	8388608	CONTROLFILE
Current.257.3	2899968	8388608	CONTROLFILE
SYSTEM.258.3	520101888	522190848	DATAFILE
group_1.263.3	33555456	41943040	ONLINELOG
USERS.264.3	67117056	68157440	DATAFILE
TEMP.265.3	268443648	269484032	TEMPFILE
SYSAUX.266.3	268443648	269484032	DATAFILE
group_1.267.3	33555456	41943040	ONLINELOG
UNDOTBS1.268.3	520101888	522190848	DATAFILE
spfile.269.1	3072	1048576	PARAMETERFILE
spfileASMEV.ora	3072	1048576	PARAMETERFILE

### Configuring archive logs to utilize ASM

Oracle allows archive logs also to be placed inside ASM. This can be accomplished by setting the LOG\_ARCHIVE\_DEST *init.ora* parameter to point to the desired ASM disk group. For example:  
\*.log\_archive\_dest='+EVDG'

### Adding a LUN (disk) to an ASM disk group

ASM allows the DBA to add and remove disks while the database is open. When storage is added or removed, a REBALANCE operation will typically be triggered unless asm\_power\_limit=0 or the storage is not used. The asm\_power\_limit parameter indicates how aggressive ASM is when REBALANCING with respect to I/O and CPU usage.

The value can range from 1 to 11 with 11 being the most aggressive setting. However, it should be noted that even at the highest setting, the database would still have a higher priority than the rebalancing. Lab testing showed that rebalancing has very low or no impact on database performance.

## Path fail-over tests for StorageWorks EVA and XP arrays on HP-UX

### Environment

- Two physical paths from server to switch and two from switch to StorageWorks array
- To force path fail-over, one switch port is forced down, eliminating the physical path to the StorageWorks EVA and XP arrays.

### Findings

- Fail-over was completely transparent to ASM.
- Database performance was not affected. Our tests were not disk performance oriented. For heavy workload there will be a performance impact.
- Path restoration is also transparent to ASM.

## Best practices for using Oracle ASM with HP StorageWorks on HP-UX

As a result of the testing described above, HP has developed a set of best practices for using Oracle ASM with HP StorageWorks.

- Configure ASM disk groups to use external redundancy.
- When building a disk group or adding to an existing one, use disks of similar capacity and performance characteristics in the same disk group.
- To leverage I/O distribution across as many resources as possible, it is best to present more than one LUN to a disk group (allowing ASM to do the striping).
- Use HP Secure Path with ASM because it complements the high availability and performance of the entire stack.
- Each device (LUN) should be managed either by Oracle ASM or by HP-UX LVM, but not both.
- Care should be taken not to attempt inadvertently to manage an ASM disk by a traditional volume manager or vice versa.
- Configure async I/O (Please consult the Oracle Administration Guide documentation).
- Use "insf" instead of "insf -e" to create the devices associated with new hardware because "insf -e" will reset the ownership to user "bin"

### LVM or ASM striping in a single EVA disk group vs. multiple EVA disk groups

LVM double striping was evaluated in a comprehensive series of tests and a white paper [http://www.oracle.com/technology/deploy/performance/pdf/EVA\\_ORACLE\\_paper.pdf](http://www.oracle.com/technology/deploy/performance/pdf/EVA_ORACLE_paper.pdf) was published on OTN in 2005. The purpose of this exercise was to determine the impact of LVM striping in a single disk group compared to LVM striping across multiple disk groups. From the test results we determined that the number of EVA disk groups does not affect performance when all of the data and logs are striped across all the EVA disk groups and that there is no penalty for double striping. This same principle applies to ASM and EVA striping.

## Oracle ASM and HP Volume Managers (Logical Volume Manager /Shared Logical Volume Manager)

When using Oracle ASM, there is no need for a separate volume manager. However, ASM can be used with Oracle data only. Therefore, a volume manager and file system is still needed for non-Oracle data.

HP Logical Volume Managers (LVM/SLVM), the default volume managers on HP-UX 11i, are mature and growing products that are well integrated with HP hardware and software, including all HP storage, HP Serviceguard (uses LVM), and SGeRAC (uses SLVM). New online features include disable/enable path and deactivate/activate physical volume.

HP has recently announced ASM support on top of SLVM for Oracle Database 10g RAC with SGeRAC 11.17 and HP-UX 11.23. This means that a raw logical volume (SLVM) can be used as an ASM disk, thus taking advantage of the multi-pathing software builds into HP volume manager.

By providing additional volume managers in addition to ASM support, HP is able to meet the broadest possible range of customer needs.

## Conclusion

When following the configuration recommendations provided in this paper, Oracle DBAs can take advantage of Oracle Database 10g ASM automatically to balance Oracle I/O workloads without the need for a separate volume manager.

HP tests with Oracle show that using ASM with HP StorageWorks EVA and XP arrays enable high availability while maintaining optimal performance levels. The ASM feature is an efficient way to manage storage for Oracle Database 10g databases.

Oracle DBAs can take advantage of Oracle 10g ASM automatically to balance Oracle workloads without the need for a separate volume manager. HP tests with Oracle show that using ASM with HP StorageWorks EVA and XP Arrays enable high availability while maintaining optimal performance levels.

For more information on using Oracle Database 10g ASM with HP StorageWorks EVA and XP Disk Arrays, please contact your HP reseller or visit [www.hp.com/go/oracle](http://www.hp.com/go/oracle).

© Copyright 2005 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein. Oracle is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

4AA0-0687ENW, February 2006