



**Hewlett Packard**  
Enterprise

# **Best Practices for Oracle Database on HPE 3PAR StoreServ Storage**

Best practices, configuration and settings for deploying Oracle on  
HPE 3PAR StoreServ Storage

# Contents

Executive summary.....	3
Solution overview.....	3
Solution components.....	3
Hardware.....	3
Servers.....	3
Storage array.....	4
Fibre Channel Switch.....	5
Fibre Channel Zoning.....	5
Software.....	5
Server OS.....	5
Application software.....	5
Oracle 12.2.0.1.0 Enterprise Edition.....	5
Load generators (for Oracle database).....	5
Storage reclamation.....	5
Best practices and configuration guidance.....	6
HPE 3PAR configuration and management.....	6
Configuring HPE 3PAR Common Provisioning Groups for Oracle.....	6
Configuring RAID for Oracle.....	7
Configuring Oracle ASM diskgroups with HPE 3PAR in Linux environments.....	7
Using Oracle ASM Storage Reclamation Utility for HPE 3PAR thin storage reclamation.....	16
HPE 3PAR thin provisioning and Adaptive Data Reduction.....	18
Thin provisioning with Oracle.....	18
Adaptive Data Reduction with Oracle.....	19
Oracle Performance with HPE 3PAR.....	23
HPE 3PAR Adaptive Flash Cache for Oracle (hybrid storage array).....	23
HPE 3PAR Priority Optimization for Oracle.....	24
Oracle OLTP performance tuning.....	24
Oracle data protection and high availability with HPE 3PAR.....	24
Oracle and HPE StoreOnce.....	24
HPE Recovery Manager Central for Oracle.....	24
HPE 3PAR Remote Copy Software for Oracle.....	25
Summary.....	26
List of acronyms.....	27
Resources and additional links.....	28
Learn more at HPE 3PAR StoreServ Storage.....	28



## Executive summary

Oracle® database environments sit at the heart of the enterprise, often running demanding and business-critical applications. Whether using Oracle to power online transaction processing (OLTP) workloads or to run decision support systems (DSS) and data warehousing, businesses rely on Oracle to keep their applications up and running. To meet the unique demands of Oracle workloads, Hewlett Packard Enterprise offers HPE 3PAR StoreServ Storage, a comprehensive storage portfolio.

Oracle Database Administrators (DBAs) can use HPE 3PAR StoreServ Storage (HPE 3PAR) to address their most significant Oracle challenges including:

**Performance** – HPE 3PAR delivers high throughput and low latencies in multi-tenant, mixed-workload Oracle environments. With industry-leading performance of up to 3.8 million IOPS and sub-millisecond latencies, HPE 3PAR gives Oracle DBAs high transactions-per-second (TPS) and minimal wait times for OLTP, using features such as Priority Optimization and Adaptive Flash Cache.

**Efficiency and Data Reduction** – In many Oracle environments, overprovisioning the primary database has become a matter of survival. As Oracle databases grow, every added core creates additional license and support fees. HPE 3PAR helps reduce Oracle sprawl and simplifies instance consolidation, driving higher TPS and providing a reduced storage footprint. What's more, HPE 3PAR Adaptive Data Reduction technologies, such as compression, can boost storage efficiency while helping enterprises bypass costly Oracle compression license fees.

**High Availability and Data Protection** – Many Oracle DBAs face a growing primary database with an increasing number of applications writing to that database. Oracle databases are getting bigger while backup windows, recovery point objectives (RPOs), and recovery time objectives (RTOs) are getting harder to meet. HPE 3PAR offers a broad set of solutions that drive high availability (HA) for Oracle, including snapshots and Remote Copy. For direct snapshot copies to HPE StoreOnce Systems as the target backup appliance, HPE 3PAR includes Recovery Manager Central (HPE 3PAR RMC-O) software at no additional cost, delivering fast backup and recovery for Oracle data. The HPE StoreOnce Catalyst Plug-in for Oracle is also offered free of charge and is tightly integrated with Oracle Recovery Manager (RMAN), giving the Oracle DBA complete visibility and control for backup and recovery tasks.

Taken together, these features help Oracle DBAs and storage administrators manage even the most demanding Oracle environments, delivering the performance, data protection, efficiency, and high availability needed to keep critical applications and business processes up and running.

## Solution overview

This paper outlines best practices for configuring and using HPE 3PAR in Oracle database environments, with recommendations based on testing by HPE engineers, and experience with thousands of Oracle customers around the world. Even though each Oracle installation is unique, most require high performance, availability, efficiency and data protection. While these requirements can sometimes be met using host-enabled technologies, the result is often significant additional host overhead that would be better utilized in application processing. This guide lays out a better option—using HPE 3PAR features to offload much of that Oracle host overhead, while still meeting stringent enterprise demands.

This guide is intended for solution architects, database administrators, and system administrators involved in the design and deployment of Oracle database environments. Combined with more detailed information about specific HPE 3PAR features (see “Resources and additional links”), this paper helps Oracle architects and administrators in building and optimizing Oracle environments with HPE 3PAR StoreServ Storage.

## Solution components

An HPE 3PAR StoreServ 20800 Storage array was the storage platform used for testing in support of this document.<sup>1</sup> Red Hat® Enterprise Linux® 7.2 running on HPE ProLiant DL380 Generation9 (Gen9) servers was used to host a three-node Oracle 12c2 RAC installation with Oracle Automatic Storage Management (ASM) managed disk groups.

### Hardware

#### Servers

Three HPE ProLiant DL380 Gen9 servers with 256 GB RAM, as follows:

- Two with 2 x Intel® Xeon® E5-2698v3 processors with a clock speed of 2300 MHz, 16/16 cores; 32 threads
- One with 2 x Intel® Xeon® E5-2690v3 processors with a clock speed of 2600 MHz, 12/12 cores; 24 threads

<sup>1</sup> Note that the content of this paper applies to all HPE 3PAR StoreServ Storage arrays in the current product line.






**Figure 1.** HPE ProLiant DL380 Gen9 server

**Storage array**

An HPE 3PAR StoreServ 20800 Storage array (release version 3.3.1, Patches: P01, P02, P04, and P13) with the following:

- Eight storage controller nodes
- SSDs: 12 @ 1.92 TB (Oracle database)
- FC drives: 128
  - 64@300 GB, 15K rpm
  - 64@900 GB, 10K rpm
- CPGs employed:
  - SSD\_RAID6 – Set size 6 data, 2 parity (Oracle data, REDO and FRA)
  - SSD\_RAID1 – Set size 2 data (Oracle REDO)
  - FC\_RAID1 (Oracle OCR and vote)



**Figure 2.** HPE 3PAR StoreServ 20800 Storage



**Fibre Channel Switch**

Two HPE SN6000B 16 GB switches

- Kernel: 2.6.14.2
- Fabric OS: v7.4.1d

**Fibre Channel Zoning**

- Two FC ports/initiators per Oracle host
- Two FC ports/targets per HPE 3PAR storage node
- One initiator to multiple targets per zone
- Oracle hosts mirrored to storage node pairs

Refer to [HPE 3PAR StoreServ Storage: A reference and best practices guide](#) for detailed recommendations on **FC Hosts Zoning**.



**Figure 3.** HPE SN6000B Fibre Channel Switch

**Software****Server OS**

Red Hat Enterprise Linux Server (release 7.2, kernel 3.10.0-327.el7.x86\_64)

**Application software****Oracle 12.2.0.1.0 Enterprise Edition**

- Oracle Grid Infrastructure 12c
- Oracle Database 12c

**Load generators (for Oracle database)**

- Swingbench 2.5.0.971 (OLTP and DSS)
- HammerDB v2.20 (OLTP)

**Storage reclamation**

Oracle ASM Storage Reclamation Utility (ASRU) (version 1.2)

Figure 4 illustrates the installation that was used to perform testing for this paper. An Oracle 12c RAC cluster was installed on three HPE ProLiant DL380 Gen9 servers. Storage for the solution was provided by Oracle ASM-managed disk groups with external redundancy built on thin-provisioned virtual volumes composed of flash devices from an HPE 3PAR StoreServ 20800 Storage array. DSS and OLTP workloads were run on one or two of the RAC instances at a time, with the other clients serving as load generators.



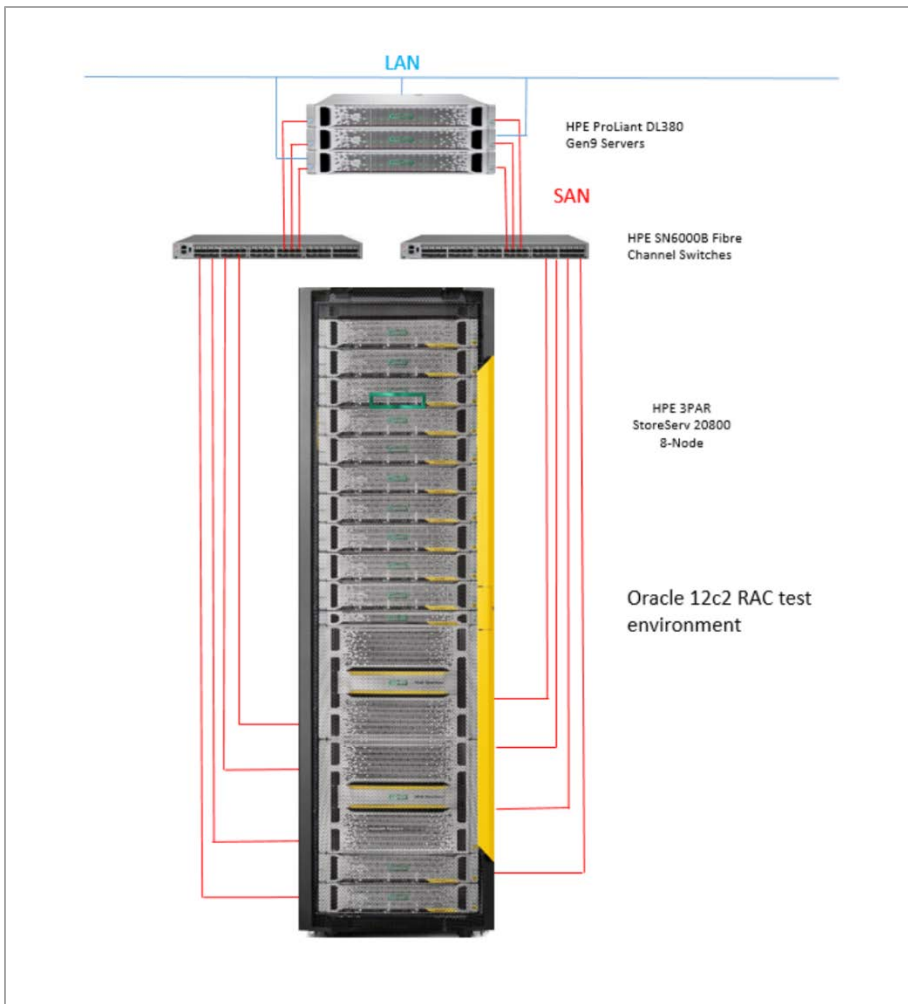


Figure 4. Oracle RAC on HPE 3PAR StoreServ Storage

## Best practices and configuration guidance

### HPE 3PAR configuration and management

#### Configuring HPE 3PAR Common Provisioning Groups for Oracle

In order to maximize performance and get the most efficiency out of an Oracle implementation on HPE 3PAR, follow these guidelines when configuring HPE 3PAR Common Provisioning Groups (HPE 3PAR CPGs):

- Use HPE 3PAR CPGs with SSDs for virtual volumes that require the highest performance.
- All members of an Oracle Automatic Storage Manager (ASM) disk group should be provisioned from the same HPE 3PAR CPG.
- HPE 3PAR CPG availability settings define the level of hardware fault tolerance, including **drive HA** (the ability to tolerate the loss of one or more drives depending on RAID level) and **cage HA** (the ability to tolerate the loss of an entire drive enclosure). When creating HPE 3PAR CPGs in a production environment, it is not necessary to create multiple CPGs with the same characteristics; doing so will not provide any level of separation, could reduce the efficiency of the system, and will have the side effect of increasing management. For management separation, volume sets (VVsets) can be used to implement consistency groups.



Refer to the appropriate HPE 3PAR StoreServ implementation guide for operating system-specific details, multipath software details, array port configuration, and other pertinent information. Implementation guides are available at the [Hewlett Packard Enterprise Information Library \(Storage\)](#).

### Configuring RAID for Oracle

HPE 3PAR storage offers the flexibility of a range of RAID levels to meet Oracle requirements. RAID implementation is abstracted from physical disks and allows multiple RAID levels to coexist on the same physical drives. This allows each volume to be provisioned using the most appropriate RAID level.

As a general rule HPE recommends RAID Multiple Parity (MP) for a good balance between data availability, capacity, and performance on HPE 3PAR storage. However, for some write-intensive Oracle workloads it might be desirable to use RAID 1 or RAID5, especially for redo logs.

Exponential growth in HDD and SSD capacity without commensurate improvements in reliability or performance has led to increased risk of data loss. For example, consider the 600 GB FC disks and 6 TB Nearline disks available on HPE 3PAR StoreServ Storage systems. The capacity difference alone means that reconstruction of a failed disk can be expected to take more than six times as long with the 6 TB disk. The Nearline disks are also slower, which further increases the amount of time to rebuild degraded drives relative to smaller, faster SSD or FC drives. A longer mean time to repair creates a larger window during which a second disk failure could cause data loss when using RAID 1 or RAID 5. RAID MP was created to address this problem. Like RAID 5, HPE 3PAR RAID MP uses distributed parity, but it stores two parity values in a manner that allows the data to be reconstructed, even in the event of two drive failures. High I/O loads make traditional RAID 6 deployments problematic; the implementation of RAID MP on HPE 3PAR arrays provides the extra level of data protection without drastically compromising performance.

- HPE 3PAR StoreServ allows selection of the number of data blocks per parity block (n+2) to suit different needs. For RAID MP, the default is 6+2; however, it is possible to select from any value from 4+2, 6+2, 8+2, 10+2 and 14+2. Higher values of “n” result in higher storage efficiency but can reduce the performance of random writes.
- When considering the use of RAID MP for higher reliability, it is important to understand RAID MP performance relative to RAID 1 and RAID 5. On HPE 3PAR, read performance is comparable for all the RAID types. If the workload is read intensive, then RAID MP will likely have little or no impact compared to other RAID levels. However, as the write ratio increases, overhead can be expected in terms of total IOPS capability.
- RAID 1 or RAID 5 may be considered for Oracle installations, when it is necessary to meet higher write performance requirements. For example, Oracle redo logs entail sequential reads and writes. RAID 1 ensures the highest performance for sequential writes due to the high write rate typical of Oracle redo logs. For a thorough discussion of performance and data security tradeoffs between different HPE 3PAR RAID levels, refer to [HPE 3PAR RAID 6: Securing your data investment](#).

### Configuring Oracle ASM diskgroups with HPE 3PAR in Linux environments

For ease of implementation in a Linux environment, HPE recommends the use of multipathing and `udev` rules to map whole raw devices without a partition table, which removes the need to impose partition alignment. For Oracle on Linux environments where disk partitioning is required or preferred, refer to “Partitioning Device Mapper nodes” in the [HPE 3PAR Red Hat Enterprise Linux, CentOS Linux and Oracle Linux Implementation Guide](#). The following is an example of how to create HPE 3PAR virtual volumes and to build Oracle ASM diskgroups from the mapped raw devices.



### Make sure multipathing is active on the Oracle host

1. Device mapper multipathing must be active on the Oracle host. Figure 5 shows how to verify the presence of the device mapper package on a Linux host using the `rpm` command. Install the device mapper package if it is not already present.

```
[root@solddb01 ~]# rpm -qa | grep device-mapper
device-mapper-libs-1.02.107-5.el7.x86_64
device-mapper-event-1.02.107-5.el7.x86_64
device-mapper-event-libs-1.02.107-5.el7.x86_64
device-mapper-persistent-data-0.5.5-1.el7.x86_64
device-mapper-multipath-libs-0.4.9-85.el7.x86_64
device-mapper-1.02.107-5.el7.x86_64
device-mapper-multipath-0.4.9-85.el7.x86_64
[root@solddb01 ~]# 3. █
```

Figure 5. Verifying device mapper is installed

2. Use the Linux `systemctl` command to verify the `multipathd` service status is active (running) on the Oracle host, as shown in Figure 6.

```
[root@solddb01 ~]# systemctl status multipathd
● multipathd.service - Device-Mapper Multipath Device Controller
   Loaded: loaded (/usr/lib/systemd/system/multipathd.service; enabled; vendor preset: enabled)
   Active: active (running) since Fri 2017-11-03 11:23:53 MDT; 4h 13min ago
     Main PID: 2555 (multipathd)
    CGroup: /system.slice/multipathd.service
           └─2555 /sbin/multipathd

Nov 03 11:24:38 solddb01 multipathd[2555]: ofsctl: spurious uevent, path already in pathvec
Nov 03 11:24:38 solddb01 multipathd[2555]: ofsctl: HDIO_GETGEO failed with 25
Nov 03 11:24:38 solddb01 multipathd[2555]: ofsctl: failed to get path uid
Nov 03 11:24:38 solddb01 multipathd[2555]: uevent trigger error
Nov 03 11:24:48 solddb01 multipathd[2555]: ofsctl: add path (uevent)
Nov 03 11:24:48 solddb01 multipathd[2555]: ofsctl: spurious uevent, path already in pathvec
Nov 03 11:24:48 solddb01 multipathd[2555]: ofsctl: HDIO_GETGEO failed with 25
Nov 03 11:24:48 solddb01 multipathd[2555]: ofsctl: failed to get path uid
Nov 03 11:24:48 solddb01 multipathd[2555]: uevent trigger error
Nov 03 15:32:31 solddb01 systemd[1]: Started Device-Mapper Multipath Device Controller.
[root@solddb01 ~]# █
```

Figure 6. `systemctl` status shows `multipathd` is active

3. If `multipathd` is not active, use `systemctl` to start it, and enable `chkconfig` or `systemctl` to make sure it will be restarted at boot of the Oracle host, as shown in Figure 7.



```
[root@solldb04 ~]# systemctl status multipathd
● multipathd.service - Device-Mapper Multipath Device Controller
   Loaded: loaded (/usr/lib/systemd/system/multipathd.service; enabled; vendor preset: enabled)
   Active: inactive (dead) since Fri 2017-11-03 16:09:34 MDT; 12s ago
     Process: 28665 ExecStart=/sbin/multipathd (code=exited, status=0/SUCCESS)
     Process: 28661 ExecStartPre=/sbin/multipath -A (code=exited, status=0/SUCCESS)
     Process: 28653 ExecStartPre=/sbin/modprobe dm-multipath (code=exited, status=0/SUCCESS)
    Main PID: 28669 (code=exited, status=0/SUCCESS)

Nov 03 16:09:33 solldb04 multipathd[28669]: ocr_votedisk3: stop event checker thread (140022275835648)
Nov 03 16:09:33 solldb04 multipathd[28669]: ocr_votedisk4: stop event checker thread (140022275802880)
Nov 03 16:09:33 solldb04 multipathd[28669]: compress_fra2: stop event checker thread (140022275770112)
Nov 03 16:09:34 solldb04 multipathd[28669]: compress_red01: stop event checker thread (140022275737344)
Nov 03 16:09:34 solldb04 multipathd[28669]: ocr20800_1: stop event checker thread (140022275704576)
Nov 03 16:09:34 solldb04 multipathd[28669]: ocr20800_2: stop event checker thread (140022275671808)
Nov 03 16:09:34 solldb04 multipathd[28669]: ocr20800_3: stop event checker thread (140022275639040)
Nov 03 16:09:34 solldb04 multipathd[28669]: ocr20800_4: stop event checker thread (140022275606272)
Nov 03 16:09:34 solldb04 multipathd[28669]: -----shut down-----
Nov 03 16:09:34 solldb04 systemd[1]: Stopped Device-Mapper Multipath Device Controller.
[root@solldb04 ~]# systemctl start multipathd
[root@solldb04 ~]# chkconfig multipathd on
Note: Forwarding request to 'systemctl enable multipathd.service'.
[root@solldb04 ~]# █
```

Figure 7. Starting multipathd and enabling daemon restart on boot.

For further information on how to setup multipathing software for Oracle ASM on Linux, refer to the [HPE 3PAR Red Hat Enterprise Linux, CentOS Linux and Oracle Linux Implementation Guide](#).

### Create the HPE 3PAR volumes

1. To create HPE 3PAR volumes to include in the ASM diskgroup, follow these steps:
  - a. Create the HPE 3PAR volumes, as shown below.
  - b. Export the volumes to all clients that will be mounting the volumes.

Figure 8 shows an example using the HPE 3PAR StoreServ Management Console (HPE 3PAR SSMC) to create a 100 GB thin-provisioned virtual volume on a RAID 6 (RAID MP) all-SSD common provisioning group. Note that all actions on virtual volumes can also be performed on the HPE 3PAR command line interface (CLI) or through the HPE 3PAR REST client.<sup>2</sup>

<sup>2</sup> Go to [python-3parclient](#) at GitHub.



Create Virtual Volume ?

General  Advanced options

Name

System

Domain

Provisioning

Dedup

Compression

CPG  ✕ 🔍

RAID 6 SSD

Size

**Additional Settings**

▶ Volume

▼ Export

Export To

Host Name	Host Set	Host OS	
soldb01	db_20800	RHE Linux 7.x	✕
soldb02	db_20800	RHE Linux 7.x	✕
soldb03	db_20800	RHE Linux 7.x	✕

3 added

Export With

LUN  ✕ 🔍  Auto

Changed: LUN to "17"

Figure 8. Creating a virtual volume, exported to three RAC nodes

- After the HPE 3PAR virtual volume has been created, the worldwide name (WWN) of the volume can be matched with the device discovered on the client. Figure 9 shows the *disk overview* of the virtual volume on HPE 3PAR SSMC. The WWN for the device is provided.



The screenshot shows the 'DB\_DATA1' disk overview. The 'General' tab is active, displaying various attributes. The 'WWN' attribute is circled in red and has the value '60002AC0000000005012B6F0007E1C3'. Other attributes include Name (DB\_DATA1), ID (76655), Volume set (DB\_DATA\_VOLS), System (20-14a1\_2080Q), Domain (—), Virtual size (100.00 GiB), Compaction (2.8:1), Type (Base), Provisioning (Thin), Dedup (Yes), Compression (No), Mode (Read/Write), RAID (RAID 6), Adaptive optimization (—), CPG (SSD r6\_2080Q), and Copy CPG (SSD r6\_2080Q). The 'Health' section shows State as Normal and State description as Normal. The 'Copies' section shows Clones (0), Snapshots (0), Snapshot policy (none), and Remote copy group (none). The 'Exports' section shows it is exported to 3 hosts, with a table listing exports to soldb01, soldb02, and soldb03.

Figure 9. HPE 3PAR disk overview provides the WWN of the device

### Configure device mapping on the Oracle host

1. On the Linux client, the `lsscsi -i` command (Figure 10) can be used to list the `udev` IDs of the multiple paths to all SCSI devices. The WWN for the HPE 3PAR virtual volume is found in the output. The numeral **3** type identifier is prepended to the WWN and alpha characters are lower case. Use this form of the WWN for mapping the device with device manager multipathing.

```
[root@soldb01 trace]# lsscsi -i > lsscsi.out
[root@soldb01 trace]# grep -i 60002AC0000000005012B6F0007E1C3 lsscsi.out
[2:0:1:17] disk 3PARdata VV 3310 /dev/sdbh 360002ac0000000005012b6f0007e1c3
[2:0:3:17] disk 3PARdata VV 3310 /dev/sdeo 360002ac0000000005012b6f0007e1c3
[2:0:5:17] disk 3PARdata VV 3310 /dev/sdxg 360002ac0000000005012b6f0007e1c3
[2:0:6:17] disk 3PARdata VV 3310 /dev/sdyh 360002ac0000000005012b6f0007e1c3
[2:0:7:17] disk 3PARdata VV 3310 /dev/sdaar 360002ac0000000005012b6f0007e1c3
[2:0:8:17] disk 3PARdata VV 3310 /dev/sdack 360002ac0000000005012b6f0007e1c3
[2:0:9:17] disk 3PARdata VV 3310 /dev/sdaed 360002ac0000000005012b6f0007e1c3
[3:0:2:17] disk 3PARdata VV 3310 /dev/sdjj 360002ac0000000005012b6f0007e1c3
[3:0:3:17] disk 3PARdata VV 3310 /dev/sdlc 360002ac0000000005012b6f0007e1c3
[3:0:4:17] disk 3PARdata VV 3310 /dev/sdmv 360002ac0000000005012b6f0007e1c3
[3:0:5:17] disk 3PARdata VV 3310 /dev/sdoo 360002ac0000000005012b6f0007e1c3
[3:0:6:17] disk 3PARdata VV 3310 /dev/sdqh 360002ac0000000005012b6f0007e1c3
[3:0:7:17] disk 3PARdata VV 3310 /dev/sdsa 360002ac0000000005012b6f0007e1c3
[3:0:8:17] disk 3PARdata VV 3310 /dev/sdtt 360002ac0000000005012b6f0007e1c3
[3:0:9:17] disk 3PARdata VV 3310 /dev/sdvm 360002ac0000000005012b6f0007e1c3
[root@soldb01 trace]#
```

Figure 10. Finding virtual volume WWN on the Linux client



2. Modify the `/etc/multipath.conf` file to include entries for newly added volumes, associating the device WWN with a descriptive alias. See Figure 11 for an example showing how to add four new devices, including the volume created in Figure 8.

```

}
multipath {
    wwid          360002ac00000000005012b6f0007e1c3
    alias         db_data1
    path_grouping_policy  multibus
    path_selector "round-robin 0"
    failback      manual
    rr_weight     priorities
    no_path_retry 5
}
multipath {
    wwid          360002ac00000000005012b700007e1c3
    alias         db_data2
    path_grouping_policy  multibus
    path_selector "round-robin 0"
    failback      manual
    rr_weight     priorities
    no_path_retry 5
}
multipath {
    wwid          360002ac00000000005012b710007e1c3
    alias         db_data3
    path_grouping_policy  multibus
    path_selector "round-robin 0"
    failback      manual
    rr_weight     priorities
    no_path_retry 5
}
multipath {
    wwid          360002ac00000000005012b720007e1c3
    alias         db_data4
    path_grouping_policy  multibus
    path_selector "round-robin 0"
    failback      manual
    rr_weight     priorities
    no_path_retry 5
}
}
blacklist {
}

```

Figure 11. Entries for new volumes added to `/etc/multipath.conf` on the Oracle host node



3. Add the new devices to the device map rules in `/etc/udev/rules.d`, as shown in Figure 12. For an Oracle RAC environment, copy the `/etc/multipath.conf` file and the device map rules in `/etc/udev/rules.d/` to all of the cluster peers.

```
[root@solddb01 rules.d]# pwd
/etc/udev/rules.d
[root@solddb01 rules.d]# tail 12-dm-permissions.rules
ENV{DM_NAME}=="fra20800_3", OWNER=="oracle", GROUP=="sysasm", MODE=="660"
ENV{DM_NAME}=="fra20800_4", OWNER=="oracle", GROUP=="sysasm", MODE=="660"
ENV{DM_NAME}=="redo20800_1", OWNER=="oracle", GROUP=="sysasm", MODE=="660"
ENV{DM_NAME}=="redo20800_2", OWNER=="oracle", GROUP=="sysasm", MODE=="660"
ENV{DM_NAME}=="redo20800_3", OWNER=="oracle", GROUP=="sysasm", MODE=="660"
ENV{DM_NAME}=="redo20800_4", OWNER=="oracle", GROUP=="sysasm", MODE=="660"
ENV{DM_NAME}=="db_data1", OWNER=="oracle", GROUP=="sysasm", MODE=="660"
ENV{DM_NAME}=="db_data2", OWNER=="oracle", GROUP=="sysasm", MODE=="660"
ENV{DM_NAME}=="db_data3", OWNER=="oracle", GROUP=="sysasm", MODE=="660"
ENV{DM_NAME}=="db_data4", OWNER=="oracle", GROUP=="sysasm", MODE=="660"
[root@solddb01 rules.d]#
```

**Figure 12.** Entries for the mapped devices are added to the device map rules in `/etc/udev/rules.d`. Ownership is assigned to the ASM admin owner and group

4. Enable the new mappings on each Oracle host. Either reboot all clients to force a rescan or execute the following sequence of commands as root on each client:

Sequence	Command	Action
1	<code>rescan-scsi-bus.sh</code>	Rescan the SCSI bus to add the new devices without rebooting. This may take a few minutes.
2	<code>udevadm control reload-rules</code>	Reload the <code>udev</code> rules without a reboot.
3	<code>udevadm trigger</code>	Manually trigger device events.
4	<code>multipath -r</code>	Force device map reload.



5. On each client, verify correct entries in the `/dev/mapper` directory and ASM admin ownership of the referenced devices, as shown in Figure 13.

```
[root@soldb01 mapper]# pwd
/dev/mapper
[root@soldb01 mapper]# ls -l db_data*
lrwxrwxrwx 1 root root 8 Nov  2 13:07 db_data1 -> ../dm-56
lrwxrwxrwx 1 root root 8 Nov  2 13:07 db_data2 -> ../dm-57
lrwxrwxrwx 1 root root 8 Nov  2 13:07 db_data3 -> ../dm-58
lrwxrwxrwx 1 root root 8 Nov  2 13:07 db_data4 -> ../dm-60
[root@soldb01 mapper]# ls -l ../dm-56
brw-rw---- 1 oracle sysasm 253, 56 Nov  2 13:07 ../dm-56
[root@soldb01 mapper]# ls -l ../dm-57
brw-rw---- 1 oracle sysasm 253, 57 Nov  2 13:07 ../dm-57
[root@soldb01 mapper]# ls -l ../dm-58
brw-rw---- 1 oracle sysasm 253, 58 Nov  2 13:07 ../dm-58
[root@soldb01 mapper]# ls -l ../dm-60
brw-rw---- 1 oracle sysasm 253, 60 Nov  2 13:07 ../dm-60
[root@soldb01 mapper]#
[root@soldb01 mapper]#
```

**Figure 13.** Following reboot of the Oracle host, verify the `/dev/mapper` entry was created for the device and that the mapped device is owned by the ASM admin owner and group

### Create the ASM disk group

1. Create a new ASM disk group.
  - a. Select the devices and choose **External Redundancy** since RAID protection is provided by the HPE 3PAR storage array.
  - b. Keep the logical sector size for the disk group at **512 bytes** in alignment with the 512 byte physical sector size of the underlying devices.
  - c. The mapped HPE 3PAR volumes now appear as eligible disks in the Oracle ASM Configuration Assistant, as shown in Figure 14.



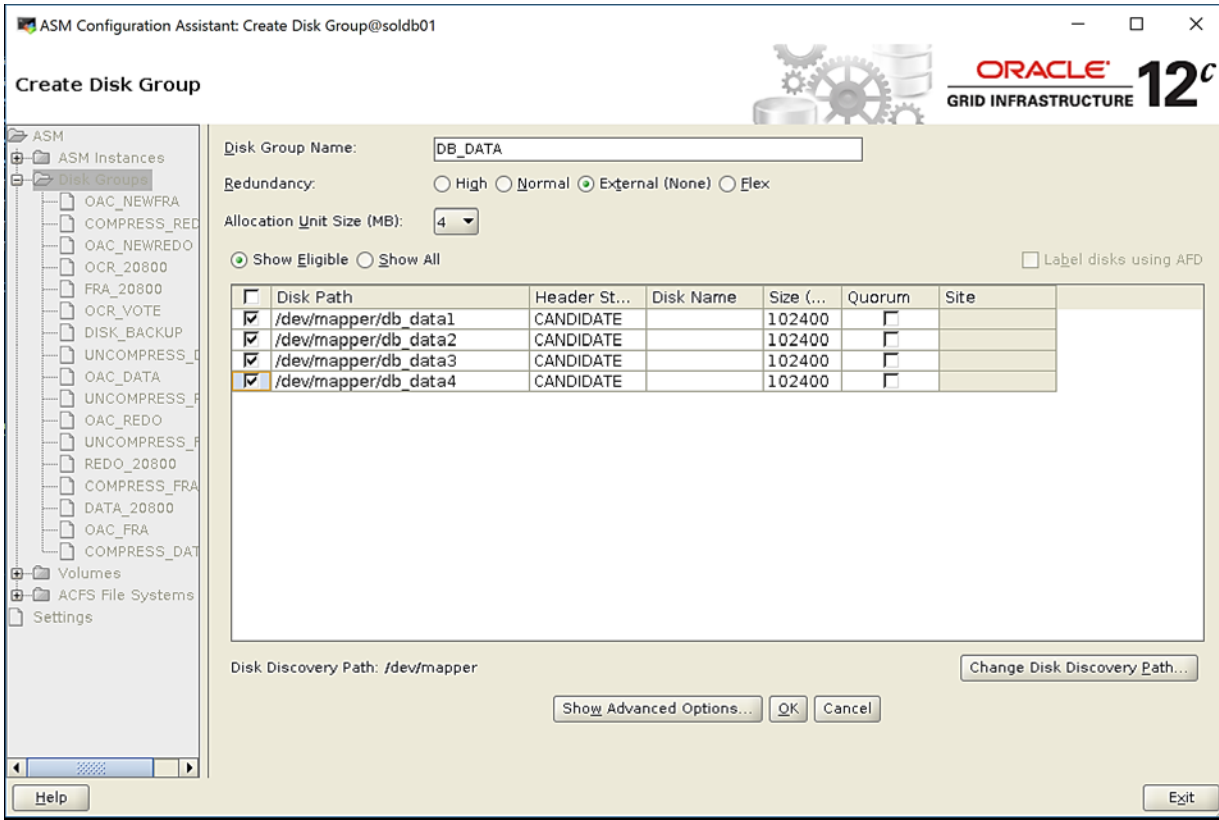


Figure 14. The asmca tool can be used to create a new ASM diskgroup from the mapped devices



After the ASM disk group has been created, verify that it is mounted on all nodes, as shown in Figure 15.

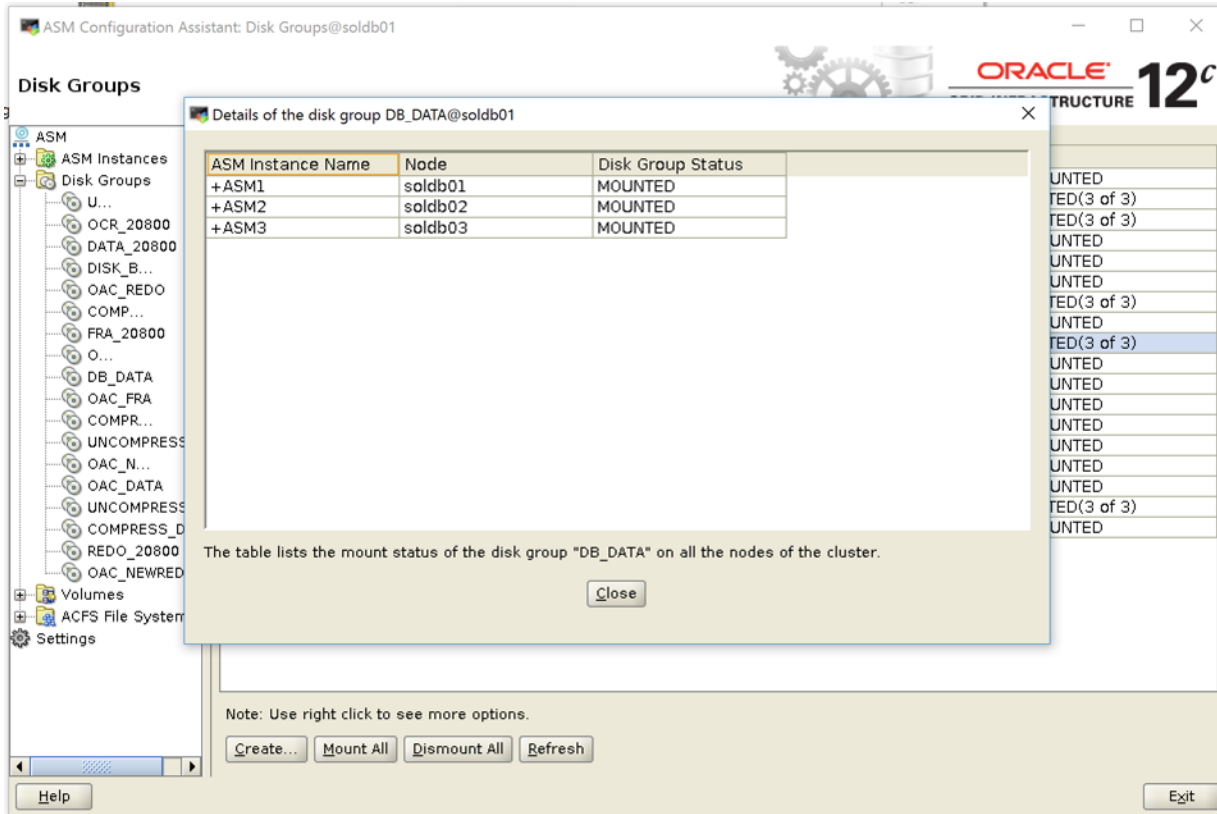


Figure 15. Following creation of the ASM diskgroup, check diskgroup status to verify that it is mounted on all ASM instances

### Using Oracle ASM Storage Reclamation Utility for HPE 3PAR thin storage reclamation

In order to exploit the full potential of thin-provisioned storage, use the Oracle ASM Storage Reclamation Utility (ASRU) to reclaim space allocated for thin provisioning and then deleted by Oracle ASM. ASRU zeroes out ASM-deleted space, so that the HPE 3PAR controller can detect that the space is no longer in use and can be re-allocated as needed. ASRU can be downloaded from the [Oracle Technology Network](#).

### Deleted space within ASM diskgroups still appears “used” to the HPE 3PAR controller

Figure 16 shows an ASM diskgroup that is 99% free, according to ASM; however, in Figure 17, the HPE 3PAR CLI shows the underlying virtual volumes to be 46% used.



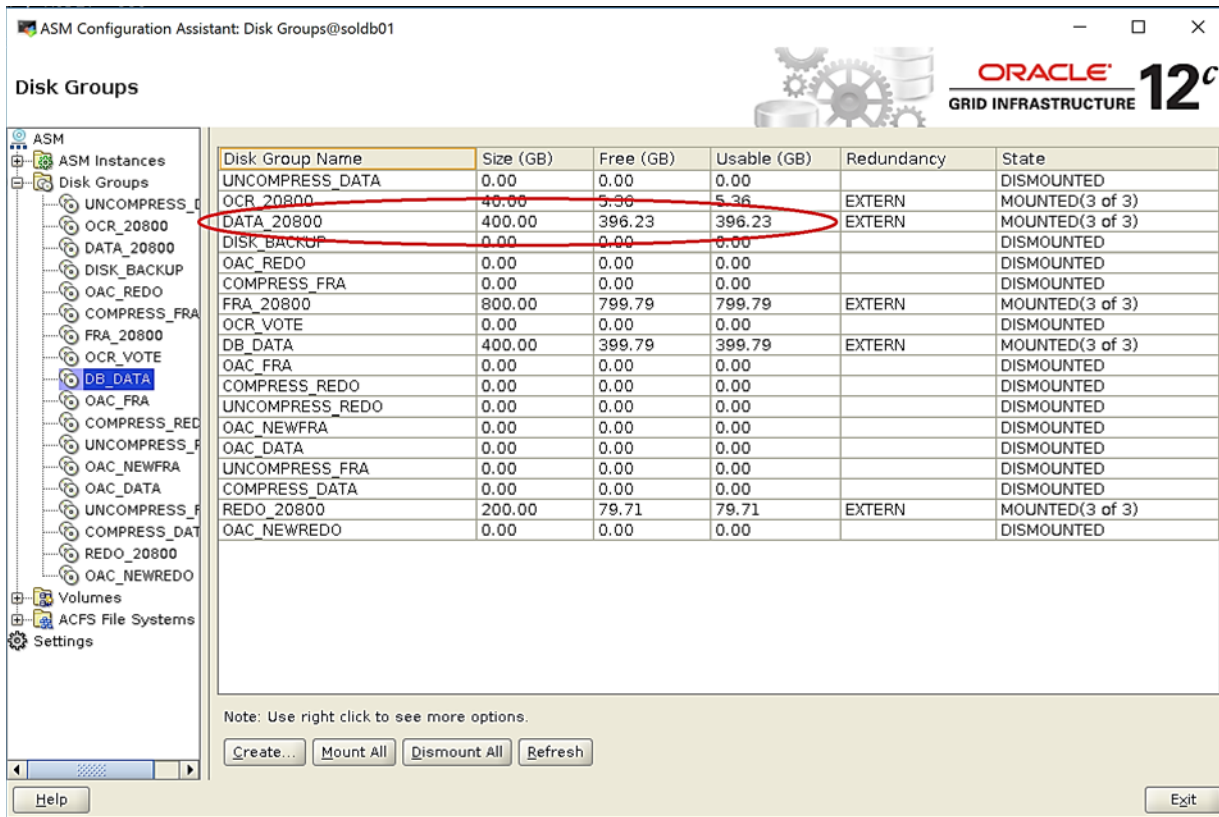


Figure 16. Disk Group DATA\_20800 is 99% free according to ASM




Figure 17. The HPE 3PAR CLI shows the underlying volumes to be 46% used

**ASM Storage Reclamation Utility corrects the discrepancy**

Follow these steps to correct the discrepancy between HPE 3PAR and ASM storage capacity numbers:

1. Execute the ASM Storage Reclamation Utility (ASRU) as ASM admin, as shown in Figure 18.

**Notice**

Interruption of ASRU during the process of disk resizing can result in **ASM diskgroup corruption**. HPE recommends performing this operation using a VNC session or on the system console to avoid ASRU being interrupted due to session disconnect.<sup>3</sup>

<sup>3</sup> Go to Oracle support [Doc ID 1668673.1: How To Validate ASM Diskgroup Consistency/State After ASM Reclamation Utility \(ASRU\) Execution Aborted](#).



```
[oracle@soldb01 ASRUv12]$ env | grep SID
ORACLE_SID=+ASM1
[oracle@soldb01 ASRUv12]$ ./ASRU DATA_20800
Useless use of \E at /home/oracle/ASRUv12/ASRU.pl line 544.
Checking the system ...done
Calculating the sizes of the disks ...done
Writing the data to a file ...done
Resizing the disks...done
Calculating the sizes of the disks ...done

/u01/app/12.2.0/grid/perl/bin/perl -I /u01/app/12.2.0/grid/perl/lib/5.22.0 /home/oracle/ASRUv12/zerofill 1
/dev/mapper/data20800_3 1197 102400 /dev/mapper/data20800_2 1213 102400 /dev/mapper/data20800_1 1217 102400i
 /dev/mapper/data20800_4 1213 102400
101203+0 records in
101203+0 records out
106119036928 bytes (106 GB) copied, 56.3473 s, 1.9 GB/s
101187+0 records in
101187+0 records out
106102259712 bytes (106 GB) copied, 59.4367 s, 1.8 GB/s
101183+0 records in
101183+0 records out
106098065408 bytes (106 GB) copied, 52.1092 s, 2.0 GB/s
101187+0 records in
101187+0 records out
106102259712 bytes (106 GB) copied, 55.3662 s, 1.9 GB/s

Calculating the sizes of the disks ...done
Resizing the disks...done
Calculating the sizes of the disks ...done
Dropping the file ...done
[oracle@soldb01 ASRUv12]$
```

Figure 18. ASRU is run against the DATA\_20800 diskgroup to reclaim the deleted space.

2. Figure 19 shows that the discrepancy in free storage capacity between ASM and HPE 3PAR has been eliminated by the ASRU run.

```
20-14a1_20800 cli% showvv -s DATA_20800*
```

Id	Name	Prov	Compr	Dedup	Type	-----Snp-----				-----Usr-----				-----Total-----				---Efficiency---			
						Rsvd	Used	Used	Wrn	Lim	Rsvd	Used	Used	Wrn	Lim	Rsvd	Used	HostWr	VSize	Compact	Compress
76638	DATA_20800_1	tpvv	No	No	base	512	0	0.0	0	0	49664	1005	1.0	0	0	50176	1005	1004	102400	>25	--
76639	DATA_20800_2	tpvv	No	No	base	512	0	0.0	0	0	49664	998	1.0	0	0	50176	998	997	102400	>25	--
76640	DATA_20800_3	tpvv	No	No	base	512	0	0.0	0	0	49664	986	1.0	0	0	50176	986	985	102400	>25	--
76641	DATA_20800_4	tpvv	No	No	base	512	0	0.0	0	0	49664	980	1.0	0	0	50176	980	979	102400	>25	--
4 total						2048	0				198656	3969				200704	3969	3965	409600		

```
20-14a1_20800 cli%
```

Figure 19. After the ASRU run, the HPE 3PAR CLI now shows the virtual volumes as only 1% used, in agreement with ASM

For additional information on the ASM Storage Reclamation Utility, refer to the [Oracle Automatic Storage Management and Thin Reclamation white paper](#).

## HPE 3PAR thin provisioning and Adaptive Data Reduction

### Thin provisioning with Oracle

Thin provisioning allows a volume to be created and made available to a host as a LUN without the need to dedicate physical storage until it is actually needed, reducing the amount of capacity required to store the data. HPE 3PAR thin provisioning leverages the **dedicate-on-write** feature of HPE 3PAR, allowing storage administrators to allocate only the disk capacity actually needed.



HPE 3PAR thin provisioning enhances storage space utilization. In an Oracle environment, this feature provides the highest value in the following cases:

- **Autoextend-enabled table spaces.** Oracle ASM-managed data files with the **autoextend** feature grow as the need for table space grows. Using such auto-extendable data files on thin provisioned volumes allows the system to allocate disk space to a database as it grows.

The process of extending a table space is I/O intensive on any kind of storage device and can affect the performance of the database during the file extension.

**Best practice:** Lessen the frequency of the performance impact by increasing the increment of space that Oracle adds using the **autoextend** on **next** parameter of the **create tablespace** command. Set **autoextend** to a larger value for rapidly growing data files and to a smaller value if the data is growing at a slower rate.

- **ASM disk groups as archive log destinations.** The combination of Oracle ASM and HPE 3PAR thin provisioning enables an expanding ASM archive log destination. Oracle databases become inaccessible when the archive log destinations fill up.

**Best practice:** Put Oracle archive logs on thin provisioned ASM disk groups, which allows the underlying storage to self-tune, accommodating unexpected increases in log switch activity. After a **level 0** backup, remove the archive logs and use the Oracle ASM Storage Reclamation Utility to free up the storage that was allocated to the old logs.<sup>4</sup> (For more information, see “Using Oracle ASM Storage Reclamation Utility for HPE 3PAR thin storage reclamation” on page 16.)

### Adaptive Data Reduction with Oracle

Adaptive Data Reduction is a collection of technologies that come standard with HPE StoreServ and are designed to reduce the data footprint on flash storage. When used alone or in combination, these technologies maximize HPE 3PAR’s flash capacity and reduce total cost of storage, while improving flash media endurance. Among these technologies are deduplication and compression.<sup>5</sup>

#### HPE 3PAR deduplication

Oracle databases usually do not have enough duplication of data blocks to make deduplication of the underlying volumes worthwhile. HPE 3PAR provides the ability through the HPE 3PAR StoreServ Management Console (HPE 3PAR SSMC) or the HPE 3PAR CLI to preview estimated deduplication savings on a virtual volume. This makes it easy to evaluate the effectiveness of HPE 3PAR deduplication on a given volume before applying it to the volume. Figure 20 shows an example of how to estimate deduplication savings with HPE 3PAR SSMC by selecting **Estimate dedup savings** from the Actions dropdown menu.

<sup>4</sup> In some Oracle environments, it may be advisable to put archive logs on fully provisioned volumes for performance reasons, as noted in “Oracle OLTP performance tuning” on page 20.

<sup>5</sup> For general information on HPE 3PAR Adaptive Data Reduction, see [Adaptive Data Reduction - HPE 3PAR StoreServ Storage with Zero Detect, Deduplication, Compression, and Data Packing](#)



**DB\_DATA2** Overview

**General**

Name	DB_DATA2
ID	76656
Volume set	
System	20-14a1_20800
Domain	—
WWN	60002AC00000000005012B700007E1C3
Virtual size	100.00 GiB
Compaction	2.8:1
Type	Base
Provisioning	Thin
Dedup	No
Compression	No
Mode	Read/Write
RAID	RAID 6
Adaptive optimization	—
CPG	SSD_r6_20800
Copy CPG	SSD_r6_20800
Comments	—

**Health**

State: ● Normal  
State description: Normal

**Copies**

Clones	0
Snapshots	0

**Exports**

Exported to hosts: 3 hosts  
Exported to host sets: 0

Exported To	Export Type	Exported
<a href="#">sldb01</a>	Volume	DB_DATA2
<a href="#">sldb02</a>	Volume	DB_DATA2
<a href="#">sldb03</a>	Volume	DB_DATA2

**Actions**

- Create
- Create similar
- Add to virtual volume set
- Save as template
- Edit
- Delete
- Export
- Unexport
- Convert
- Tune
- Restart tune
- Rollback tune
- Create snapshot
- Promote snapshot
- Create clone
- Stop clone
- Resync clone
- Promote clone
- Start Peer Motion
- Estimate dedup savings**
- Estimate compression savings

Figure 20. Estimating deduplication savings with HPE 3PAR SSMC

Results of the deduplication estimation task can be viewed in the task details in the HPE 3PAR CLI or on the HPE 3PAR SSMC. Figure 21 shows the deduplication preview estimate task detail on HPE 3PAR SSMC for virtual volume DB\_DATA2. In this case, with a volume containing Oracle datafiles, the deduplication space saving would be virtually nil, with a ratio of **1**. Unlike some storage products, HPE 3PAR provides the ability to disable deduplication, improving performance by using a more streamlined data path when there is little benefit to be gained by deduplication.



**General**

System 20-14a1\_20800  
 Serial number 4UW0001450  
 ID 1557  
 Type Preview dedup ratio  
 Start time Nov 10, 2017 10:52:27 AM MST  
 End time Nov 10, 2017 10:54:19 AM MST  
 Duration 1 minute 52 seconds  
 Priority —  
 Origin System

▼ Task Detail

2017-11-10 10:52:27 MST Created task.  
 2017-11-10 10:52:27 MST Started checkvV space estimation started with option -dedup\_dryrun  
 2017-11-10 10:52:27 MST Updated checkvV process for VV DB\_DATA2  
 2017-11-10 10:52:27 MST Updated VV DB\_DATA2 stats updated: 0 MB analyzed.  
 2017-11-10 10:52:29 MST Updated checkvV process for VV DB\_DATA2  
 2017-11-10 10:52:29 MST Updated VV DB\_DATA2 stats updated: 465 MB analyzed.  
 2017-11-10 10:52:30 MST Updated checkvV process for VV DB\_DATA2  
 2017-11-10 10:52:30 MST Updated VV DB\_DATA2 stats updated: 930 MB analyzed.  
 2017-11-10 10:52:31 MST Updated checkvV process for VV DB\_DATA2  
 . . .  
 2017-11-10 10:54:19 MST Updated VV DB\_DATA2 stats updated: 37501 MB analyzed.  
 2017-11-10 10:54:19 MST Updated checkvV process for VV DB\_DATA2  
 2017-11-10 10:54:19 MST Updated VV DB\_DATA2 stats updated: 37501 MB analyzed.  
 2017-11-10 10:54:19 MST Finished checkvV space estimation process finished

---

Dedup Estimate Results (MiB)

ID	Name	User Data	Estimated	Ratio
76656	DB_DATA2	36924	—	—
	<b>Total</b>	<b>36924</b>	<b>36910</b>	<b>1</b>

Figure 21. HPE 3PAR task detail for deduplication estimation

**HPE 3PAR compression**

In many production environments, compression can be applied to Oracle data for significant storage-space savings. Some compression technologies come with a high cost; for example, Oracle Advanced Compression requires an Oracle license key priced on a per-core or per-user basis. The HPE 3PAR compression feature offers advanced compression capabilities on flash storage at no additional license cost. HPE 3PAR compression is easy to apply, and, as with deduplication, HPE 3PAR provides the ability through the HPE 3PAR SSMC or the CLI to preview estimated compression space savings on a virtual volume. Figure 22 shows an example of using the checkvV command to obtain a compression ratio estimate for the UNCOMPRESS\_DATA\_1 volume.



```

soldb_3Par cli% checkvv -compr_dryrun UNCOMPRESS_DATAv1
The -compr_dryrun command relies on HPE 3PAR Thin Compression
technology to emulate the benefit of converting VVs to compressed volumes. The
information provided is applicable to HPE 3PAR Thin Compression only and it
should be considered for estimation purposes only; results of a conversion may
differ slightly.

The command can be run against live production volumes and will generate a non-intrusive
background task. The task may take some time to run, depending on the size and number of volumes,
and can be monitored via the showtask commands.

Do you want to continue with this operation?
select q=quit y=yes n=no: y
Task 8935 has been started to validate administration information for VVs: UNCOMPRESS_DATAv1
soldb_3Par cli% showtask -d 8935
  Id Type          Name      Status Phase Step -----StartTime----- -----FinishTime----- -Priority- -User--
8935 compr_dryrun checkvv done    ---  --- 2017-11-14 10:29:18 MST 2017-11-14 10:29:23 MST n/a      3paradm

Detailed status:
2017-11-14 10:29:18 MST Created      task.
2017-11-14 10:29:18 MST Started      checkvv space estimation started with option -compr_dryrun
2017-11-14 10:29:23 MST Finished     checkvv space estimation process finished

      User Data      (Compressed)
  Id Name          Size(MB) Size(MB) Ratio
102 UNCOMPRESS_DATAv1 275661 135128 2.04
-----
  1 total          275661 135128 2.04
soldb_3Par cli% █

```

Figure 22. Estimating compression on a thin provisioned virtual volume using the HPE 3PAR checkvv command

HPE 3PAR compression can be enabled when a virtual volume is created, or it can easily be applied to an existing volume. Figure 23 shows an example of applying compression to an existing volume from the CLI with the tunevv command.

```

soldb_3Par cli% tunevv usr_cpg SSD_r6_1920 -compr UNCOMPRESS_DATAv1
Are you sure you want to tune VV "UNCOMPRESS_DATAv1"?
select y=yes n=no: y
Task 8955 started

soldb_3Par cli% █

```

Figure 23. Applying compression to volume with the tunevv command



In this example, the compression task completes in 5 minutes and 52 seconds, resulting in an actual compression ratio of 2.23:1, as shown in Figure 24.

```

sldb_3Par cli% showtask -done
  Id Type      Name      Status Phase Step -----StartTime----- -----FinishTime----- -Priority- -User--
8935 compr_dryrun checkvv done    ---   --- 2017-11-14 10:29:18 MST 2017-11-14 10:29:23 MST n/a      3paradm
8939 compr_dryrun checkvv done    ---   --- 2017-11-14 11:28:46 MST 2017-11-14 11:28:50 MST n/a      3paradm
8953 compr_dryrun checkvv done    ---   --- 2017-11-14 16:08:47 MST 2017-11-14 16:08:51 MST n/a      3paradm
8954 compr_dryrun checkvv done    ---   --- 2017-11-14 16:10:00 MST 2017-11-14 16:10:03 MST n/a      3paradm
8955 convert_vv UNCOMPRESS DATAv1 done    ---   --- 2017-11-14 16:11:40 MST 2017-11-14 16:17:32 MST n/a      3paradm
sldb_3Par cli%
sldb_3Par cli% showvv -s UNCOMPRESS_DATAv1
-----Snp-----Usr-----Total-----
--(MiB)--  -(% VSize)--  ---(MiB)---  -(% VSize)--  -----(MiB)-----  ---Efficiency---
 Id Name      Prov Compr Dedup Type Rsvd Used Used Wrn Lim Rsvd Used Used Wrn Lim Rsvd Used HostWr VSize Compact Compress
-----
 102 UNCOMPRESS_DATAv1 tpvv Yes No base 512 0 0.0 0 0 139776 123461 16.1 0 0 140288 123461 275789 768000 6.22 2.23
-----
 1 total 512 0 139776 123461 140288 123461 275789 768000
sldb_3Par cli% █
    
```

**Figure 24.** The showtask command displays the completed task and the resulting compression ratio exceeds the prediction

HPE 3PAR compression is an excellent data reduction technology and is well-suited for administrators seeking to reduce flash-storage space and to drive efficiency. The benefits of compression should be weighed against other considerations, such as performance and resource allocation. Here are key points to keep in mind when using HPE 3PAR compression with Oracle databases.

- Before applying HPE 3PAR compression, make sure that storage controllers are not CPU-bound. To run without impacting performance, HPE 3PAR compression requires some spare CPU cycles on the storage side.
- Avoid making assumptions about the compressibility of any particular dataset. The actual rate of HPE 3PAR compression may vary, depending on unique workloads and environments.
- Note that pre-initialized tablespaces are highly compressible. Unused allocated space compresses to almost nothing.

**Types of Oracle data for compression on HPE 3PAR**

- HPE 3PAR compression is an economical and effective space-saving option for Oracle datafiles in both DSS and OLTP environments.
- HPE 3PAR compression may also be considered for Oracle online redo logs, temporary data, and fast recovery area where additional space saving is desired.

For additional information about Oracle and HPE 3PAR compression, refer to [HPE Best Practices for Oracle 12c1 on HPE 3PAR with Compression](#).

**Oracle Performance with HPE 3PAR**

**HPE 3PAR Adaptive Flash Cache for Oracle (hybrid storage array)**

Random-read workloads, such as Oracle workloads, are some of the most difficult for a hybrid storage array to handle. HPE 3PAR Adaptive Flash Cache (HPE 3PAR AFC) functionality built into HPE 3PAR StoreServ allows SSDs in a hybrid array to act as an extended read cache to hold random read data that has aged out of DRAM cache. By increasing the size of the read cache, the chances are increased for random read data to be serviced out of a much faster tier than the spinning media where it resides.

Adaptive Flash Cache increases the overall random read IOPS and reduces the average latency an array can deliver by reducing load on the back end of the array and serving these requests from flash, which results in higher performance and lower latency. To simulate the benefits of using HPE 3PAR AFC in Oracle environments, create Flash Cache in simulation mode. Using simulation mode, helps assess whether augmenting on-board DRAM benefits current workloads. Based on the results, decide whether to add SSDs or use a portion of existing SSD capacity for AFC<sup>6</sup>.

<sup>6</sup> See the [HPE 3PAR Adaptive Flash Cache](#) white paper for a full description of the Adaptive Flash Cache simulator.



### **HPE 3PAR Priority Optimization for Oracle**

Priority Optimization is a quality of service (QoS) offering that provides the ability to set maximum and minimum limits for IOPS or bandwidth available to a VVset. Priority Optimization also ensures low latency through the unique ability to set latency goals. QoS is key for Oracle environments where online transaction processing (OLTP) and online analytical processing (OLAP) environments often reside side-by-side, or production and non-production environments may need to co-exist. When setting maximum limits, minimum goals, and latency goals, it is important to carefully set QoS rules, understand system performance, and consider historical performance data. A good understanding of the applications that reside on the volumes is equally important.

When configuring QoS rules in Oracle environments, follow the HPE 3PAR Priority Optimization guidelines. For more information and best practices for Priority Optimization, refer to the [HPE 3PAR Priority Optimization](#) product documentation.

### **Oracle OLTP performance tuning**

HPE recommends that Oracle environments be tuned to meet the performance and scalability needs of the particular environment. No configuration setting can be guaranteed to enhance performance in all situations. The following are some general suggestions for Oracle OLTP workloads.

- As explained in “Configuring Oracle ASM diskgroups with HPE 3PAR in Linux environments” on page 7, it is advisable to either build ASM disk groups on raw (unpartitioned) devices or, where partitioned devices are required, ensure proper partition alignment. Misaligned partitions can negatively impact performance.
- OLTP performance can be hampered by slow redo log response times, which cause commits to be a write bottleneck. For example, when an Oracle AWR report shows `log file sync` to be a top foreground wait event, this indicates that redo latency may be causing a performance issue.

Consider the following measures to enhance redo performance:

- Dedicate an ASM disk group to redo log files built on HPE 3PAR virtual volumes comprised of flash devices.
- Employ a faster HPE 3PAR RAID level for the redo disk group: RAID1 or RAID5.
- Utilize fully provisioned (as opposed to thin) virtual volumes for the redo log disk group.
- Increase the size of redo logs to reduce the frequency of log file switches.

## **Oracle data protection and high availability with HPE 3PAR**

### **Oracle and HPE StoreOnce**

For data protection in Oracle environments HPE recommends HPE StoreOnce, the industry-leading disk-based backup and deduplication solution offering a single, federated deduplication technology across the entire enterprise—from small, remote offices to enterprise data centers. HPE StoreOnce deduplication can be deployed at the application source, such as on Oracle database servers, with backup applications (for example, Veeam, Veritas or Micro Focus Data Protector), and on target HPE StoreOnce systems. This means that an IT organization can deploy a deduplication solution where it makes the most business sense, not based on vendor limitations. This capability is made possible with HPE StoreOnce Catalyst, a backup-optimized, high-performance protocol.

HPE StoreOnce Catalyst Plug-in for Oracle provides Oracle DBAs with control and visibility of backups and restores from HPE 3PAR StoreServ to HPE StoreOnce, free of charge. HPE StoreOnce is integrated with Oracle RMAN through HPE StoreOnce Catalyst, providing DBAs faster backup and accelerated recovery of Oracle databases. The HPE StoreOnce Catalyst Plug-in is installed directly onto Oracle database servers and enables direct backup and restore of single or clustered Oracle databases to HPE StoreOnce Catalyst stores. In the case of deploying third-party backup/restore software, integration of Oracle RMAN and backup software such as Veeam, Veritas and Micro Focus Data Protector is made possible through HPE StoreOnce Catalyst.

### **HPE Recovery Manager Central for Oracle**

HPE Recovery Manager Central for Oracle (RMC-O) allows Oracle snapshot data to be automatically backed up directly from HPE 3PAR to HPE StoreOnce systems. Backups can then be replicated to the cloud through HPE Cloud Bank Storage. Combining the performance of local and remote snapshots with the protection of backups, HPE RMC for Oracle enables fast, efficient, reliable, and simple protection of Oracle databases and instances. It is included with every HPE 3PAR, and no additional licenses need to be purchased.

For more information, refer to [HPE 3PAR Recovery Manager Central for Oracle configuration and best practices guide](#).



### **HPE 3PAR Remote Copy Software for Oracle**

The ability to replicate data to remote locations is critical for Oracle database environments. HPE 3PAR Remote Copy is a unique replication technology that allows for efficient, affordable protection and sharing of Oracle data.

HPE recommends using HPE 3PAR Remote Copy in Oracle environments to reduce the cost of remote data replication and disaster recovery (DR) through Asynchronous and Synchronous Replication. These thin-copy technologies help meet low-recovery-time objectives (RTOs) and zero-data-loss recovery point objectives (RPOs). HPE 3PAR Remote Copy also makes it economical for customers to implement DR solutions by allowing replication across high-end, mid-range, and flash platforms within the HPE 3PAR StoreServ family.

HPE 3PAR also offers autonomic DR configuration that facilitates setup and test of an entire Oracle DR deployment in just minutes, including multisite replication, using multiple models of HPE 3PAR systems, and multiple replication modes.

With native IP and Fibre Channel (FC) support, HPE 3PAR Remote Copy delivers out-of-the-box connectivity over an IP network. Oracle database administrators can replicate data more cost effectively by leveraging existing infrastructure. This enables robust DR in Oracle environments without expensive converters and the management burden of a separate, special-purpose infrastructure.

Since HPE 3PAR Remote Copy is thin-aware, target volumes provide the same cost and ease-of-use as thin-source volumes. With the combination of HPE 3PAR thin provisioning and HPE 3PAR Remote Copy in Oracle environments, both primary and remote sites can share in the benefits of allocating volumes just once, while consuming only necessary physical capacity.

#### **Asynchronous Replication**

- With Asynchronous Replication solutions, the primary storage array does not wait for acknowledgment from the secondary storage array when processing I/O.
- When network bandwidth is at a premium, HPE 3PAR Remote Copy offers the Asynchronous Periodic replication mode.
- Changed data within a volume or a Remote Copy volume group is transferred in bulk and only once per period, no matter how many times the data may have been changed by hosts between synchronization periods.
- For Oracle applications that require a low recovery-point objective (RPO), HPE 3PAR offers the Asynchronous Streaming replication mode, which ensures a data protection RPO measured in seconds.

#### **Synchronous Replication**

- With Synchronous Replication, write I/O initiated by the host requires the primary storage array to wait for I/O acknowledgement from the secondary storage array before sending an acknowledgement to the host. Thus, the two arrays hold fully synchronized copies of the same data. This is essential for instant failover of transactional Oracle applications. Synchronous mode solutions are suitable for Oracle data centers within metro distances and have an RPO of zero seconds.
- Synchronous Long Distance (SLD) mode delivers a disaster recovery solution across long distances with a potential for a zero data loss RPO and an RTO of minutes. This is achieved with a replication configuration that uses three sites to simultaneously replicate a virtual volume from the primary array (in synchronous mode) to an HPE 3PAR array located at a synchronous site. SLD allows for resilient data access in the event of an entire site failure and can provide solutions with multiple restore locations.



**Best practice:** HPE 3PAR arrays allow volumes to be placed in consistency groups. The use of volume sets ensures that coordinated snapshots are taken across the volumes that make up the Oracle ASM disk group to provide consistent data in the snapshots.

**Best practice:** All the volumes comprising an ASM disk group should belong to the same HPE 3PAR virtual volume set when using HPE snapshot or Remote Copy. This makes sure that the volumes are copied with the same point-in-time image.

Refer to the [Oracle Automatic Storage Management Administrator's Guide](#) for more information about ASM. For more information about 3PAR Remote Copy, refer to [HPE 3PAR Replication Software Suite](#).

## Summary

HPE 3PAR offers a unique set of features to meet the challenge of implementing Oracle database applications with outstanding workload performance, efficient storage utilization, high availability, and data protection. With models ranging from a few terabytes to many petabytes, the HPE 3PAR product family includes primary storage to fit the business requirements of every Oracle installation. The recommendations and guidelines in this document will assist the solution architect, storage administrator, or Oracle DBA in implementing Oracle in an HPE 3PAR environment, scaled to suit the specific requirements of their organization.

HPE recommends implementing a proof-of-concept using an Oracle test environment that matches as closely as possible the planned production environment. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept, contact an HPE Services representative ([hpe.com/us/en/services/consulting.html](http://hpe.com/us/en/services/consulting.html)) or your HPE partner.



## List of acronyms

<b>ASM</b>	Oracle Automatic Storage Management.	<b>OLAP</b>	Online analytical processing.
<b>ASRU</b>	Oracle ASM Storage Reclamation Utility.	<b>OLTP</b>	Online transaction processing.
<b>AWR</b>	Oracle Automatic Workload Repository.	<b>PIT</b>	Point-in-time.
<b>CLI</b>	Command line interface.	<b>QoS</b>	Quality of Service.
<b>DBA</b>	Oracle Database Administrator.	<b>RAC</b>	Oracle Real Application Cluster.
<b>DR</b>	Disaster recovery.	<b>RAID MP</b>	RAID Multiple Parity (RAID 6).
<b>DSS</b>	Decision Support System.	<b>RMAN</b>	Oracle Recovery Manager.
<b>FC</b>	Fibre Channel.	<b>RPO</b>	Recovery point objective.
<b>GUI</b>	Graphical user interface.	<b>RTO</b>	Recovery time objective.
<b>HA</b>	High availability.	<b>SLD</b>	Synchronous Long Distance.
<b>HPE 3PAR AFC</b>	HPE 3PAR StoreServ Adaptive Flash Cache.	<b>TPS</b>	Transactions-per-second.
<b>HPE 3PAR CPG</b>	HPE 3PAR StoreServ Common Provisioning Group.	<b>VVset</b>	Virtual volume set
<b>HPE 3PAR RMC-O</b>	HPE 3PAR StoreServ Recovery Manager for Oracle.	<b>WWN</b>	Worldwide name.
<b>HPE 3PAR SSMC</b>	HPE 3PAR StoreServ Management Console.		



## Resources and additional links

HPE 3PAR RAID 6: Securing your data investment

[hpe.com/h20195/v2/Getdocument.aspx?docname=a00000244enw](http://hpe.com/h20195/v2/Getdocument.aspx?docname=a00000244enw)

HPE Best Practices for Oracle 12c1 on HPE 3PAR StoreServ Storage with Compression

[h20565.www2.hpe.com/hpsc/doc/public/display?docId=a00022246en\\_us](http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=a00022246en_us)

HPE 3PAR Red Hat Enterprise Linux, CentOS Linux, and Oracle Linux Implementation Guide

[h20565.www2.hpe.com/hpsc/doc/public/display?docId=c04448818](http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=c04448818)

HPE 3PAR StoreServ Storage Concepts Guide

[h20564.www2.hpe.com/hpsc/doc/public/display?docId=c04204225](http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=c04204225)

HPE 3PAR StoreServ Storage: A reference and best practices guide for HPE 3PAR StoreServ Storage

[hpe.com/h20195/v2/GetPDF.aspx/4AA4-4524ENW.pdf](http://hpe.com/h20195/v2/GetPDF.aspx/4AA4-4524ENW.pdf)

HPE 3PAR StoreServ Storage software: Family data sheet

[hpe.com/h20195/v2/GetPDF.aspx/4AA2-8396ENN.pdf](http://hpe.com/h20195/v2/GetPDF.aspx/4AA2-8396ENN.pdf)

HPE 3PAR OS Command Line Interface Reference

[support.hpe.com/hpsc/doc/public/display?docId=c04204279](http://support.hpe.com/hpsc/doc/public/display?docId=c04204279)

HPE 3PAR StoreServ Architecture technical white paper

[h20195.www2.hpe.com/v2/Getdocument.aspx?docname=4AA3-3516ENW](http://h20195.www2.hpe.com/v2/Getdocument.aspx?docname=4AA3-3516ENW)

Oracle Automatic Storage Management and Thin Reclamation: An Oracle White Paper

[oracle.com/technetwork/database/oracle-automatic-storage-management-132797.pdf](http://oracle.com/technetwork/database/oracle-automatic-storage-management-132797.pdf)

Oracle Database Performance Tuning Guide 12c Release 2

[docs.oracle.com/database/122/TGDBA/TGDBA.pdf](http://docs.oracle.com/database/122/TGDBA/TGDBA.pdf)

Maintaining High Storage Utilization with Oracle ASM Storage Reclamation Utility and 3PAR Thin Persistence

[oracle.com/technetwork/database/oracle-asru-3par.pdf](http://oracle.com/technetwork/database/oracle-asru-3par.pdf)


Swingbench – [dominicgiles.com](http://dominicgiles.com) [dominicgiles.com/swingbench.html](http://dominicgiles.com/swingbench.html)

HammerDB [hammerdb.com](http://hammerdb.com)

HPE Advisory and Transformation Services [hpe.com/us/en/services/consulting.html](http://hpe.com/us/en/services/consulting.html)

## Learn more at HPE 3PAR StoreServ Storage

[hpe.com/us/en/storage/3par.html](http://hpe.com/us/en/storage/3par.html)



**Sign up for updates**

---

© Copyright 2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Oracle is a registered trademarks of Oracle and/or its affiliates. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries

a00038978en\_us, January 2018, Rev 2

