



Technical Report

Element Software for Oracle Configuration Guide

For SolidFire and NetApp HCI

Bobby Oommen, NetApp
August 2020 | TR-4768

TABLE OF CONTENTS

| | | |
|----------|--|-----------|
| 1 | Introduction | 4 |
| 1.1 | Oracle Considerations | 4 |
| 1.2 | CPU Optimization with NetApp HCI | 4 |
| 1.3 | Core Right-Sizing | 4 |
| 2 | Storage Configuration | 5 |
| 2.1 | Create an Account | 5 |
| 2.2 | Create a QoS Policy | 6 |
| 2.3 | Create a Volume | 7 |
| 2.4 | Create Volume Access Groups | 8 |
| 3 | Operating System Configuration | 9 |
| 3.1 | Add Packages | 10 |
| 3.2 | Update Kernel Parameters | 10 |
| 3.3 | Optimize Network Performance | 10 |
| 3.4 | Set User Limits | 10 |
| 3.5 | Set Permissions | 10 |
| 4 | Oracle Configuration | 11 |
| 4.1 | Install Prerequisite Packages | 11 |
| 4.2 | Configure Oracle ASM | 11 |
| 4.3 | Configure ASM Library | 11 |
| 5 | Microsoft Windows Configuration | 11 |
| 6 | Backup and Recovery Using Storage Snapshot Copies | 11 |
| 6.1 | Create Snapshots (Backup) | 12 |
| 6.2 | Restore Snapshots (Recovery) | 13 |
| 7 | RMAN Merge | 14 |
| 7.1 | Enable Block Change Tracking | 14 |
| 7.2 | Create RMAN Incremental | 14 |
| 7.3 | Use Cases for RMAN Merge | 15 |
| 8 | Database Cloning | 15 |
| 8.1 | Clone a Volume | 15 |

| | |
|---|-----------|
| Appendix A: RMAN Merge Output | 17 |
| Where to Find Additional Information | 17 |
| Version History | 18 |

1 Introduction

1.1 Oracle Considerations

This document describes how to configure Oracle and NetApp Element software on NetApp SolidFire storage arrays and NetApp HCI in accordance with NetApp recommended best practices. Element Software on NetApp HCI and SolidFire arrays can support a wide range of database application use cases, such as database consolidation, dynamic resource allocation for development and testing environments, and integration of cloud infrastructure.

By leveraging the per-volume QoS of Element, you can be assured that individual databases will maintain sufficient I/O throughput without being affected by other workloads on the storage cluster. For more in depth information on Element Software QoS, refer to [TR-4644: NetApp SolidFire Quality of Service \(QoS\)](#) and [TR-4632: NetApp HCI QoS and Mixed Workloads](#).

The snapshot capabilities of Element allow you to have point-in-time backups of the active file system or storage volume. You can use application-consistent snapshot copies for rapid recovery of corrupted datasets and to create space-efficient copies of datasets for development and testing. By using group snapshot copies, the multiple volumes for the database will be fully consistent. For more information on Element Software data protection refer to [TR-4641: NetApp HCI Data Protection](#).

There is no restriction on running single Oracle or Oracle Real Application Clusters (RAC) instances on Element OS. However, you should also research other NetApp [products](#) if you want to achieve consistent submillisecond and high-bandwidth I/O performance. Oracle Licensing with NetApp HCI

1.2 CPU Optimization with NetApp HCI

The use of dedicated compute and dedicated storage servers is a key differentiator with NetApp HCI because it means that core-licensed products do not run on storage server CPUs. As a result, you can reduce the number of cores that must be licensed for products such as Oracle Database and can simplify software management and auditing.

The NetApp HCI architecture is also more efficient from a licensing and performance point of view because cores that are licensed for Oracle are not given the task of storage work. The CPU work that is required to present and to maintain a LUN, including features such as compression and deduplication, is isolated to a different system. It therefore does not create a “NetApp HCI tax” on the compute hypervisor, including its cores that were licensed for Oracle work.

1.3 Core Right-Sizing

The processing capabilities of CPUs have increased much faster than the processing demands of most database workloads. Sometimes databases are limited by CPU work, but it is generally a result of the processing limits of a single core and is not a limitation of the CPU. The result is an increasing number of idle cores on database servers that still must be licensed for the Oracle Database software. This underutilization of CPU resources is a waste of capital expenditure, not only in terms of licensing costs, but also in terms of the cost of the server itself, heat output, and so on.

For information NetApp HCI platform specifications, refer to the [DS-3881: NetApp HCI Datasheet](#).

Oracle licensing requirements can be complicated, but as a rule, all cores in which the software product can run without deliberate user intervention must be properly licensed. Virtualization adds requirements because of the native clustering capabilities. There is no requirement to license all ESX clusters under a single VMware vCenter Server, although the ability to migrate workloads between clusters does increase the chance of user error that leads to a license violation. A more common approach is to fully license a particular ESX cluster for core-licensed Oracle products. This approach reduces the risk of user error in VM placement while avoiding the licensing of unnecessary cores. All nodes are then available to provide high-availability (HA) services.

When it is not feasible to license all the cores, such as a large single ESX cluster, VMware affinity rules are often used to segment the cluster. For example, two out of six ESX nodes might be dedicated for Oracle work, and VMware affinity rules can prevent the software from running elsewhere. You must verify that new databases are properly restricted to the desired nodes. Finally, you can use [VMware policies and affinity rules](#) to limit the way that running database instances are tied to specific CPU cores. No controlling Oracle documentation currently exists for Oracle virtualization license policies. The recommended option is to reduce the number of cores that you deploy in a single server so that servers can be fully licensed wherever possible.

Next, by using the previously outlined guidelines, you explain your intended deployment of Oracle products on the NetApp HCI system to your Oracle account team, including the number of cores you are expected to license. If the Oracle account team expresses any concerns, ask them to clarify the policies with documented, binding authority that covers your scenario.

2 Storage Configuration

The configuration that is described in this document is for an Oracle database that uses an Oracle Automatic Storage Management (ASM) layout. You can also use a logical volume manager (LVM) configuration on Linux to configure the Oracle application. Element supports presentation of the storage in a 4K sector size (native mode) and in a traditional 512-byte sector size (512e). Provided that there is no partition misalignment at the host level, there is no performance effect from choosing emulation mode.

For the database layout, follow one of the storage configurations:

- iSCSI LUNs that are managed by the iSCSI initiator on the VM, not on the hypervisor
- Separate datastore and storage LUNs for Oracle data files, redo logs, and archival logs

To get the following benefits, you should avoid the use of datastores for Oracle files:


- **Transparency.** When a guest VM owns the file systems, it is easier for you to identify the source of the file systems for your data.
- **Manageability.** When a guest VM owns the file systems, manageability of provisioning, monitoring, and data protection across the environment becomes easier.
- **Portability.** When a VM owns its file systems, the process of moving an Oracle environment becomes simpler. File systems can easily be moved between virtualized and nonvirtualized guests and to other hypervisors.

Note: NetApp strongly recommends that you use QoS policies to set the IOPS characteristics when you provision the volumes for the Oracle database. The use of QoS policies enables better manageability of IOPS settings at the volume level.

2.1 Create an Account

To create an account, complete the following steps:

1. Log in to the NetApp Element software UI.
2. Select Management > Accounts. The Account List window opens.



3. Click Create Account. The Create a New Account window opens.

4. Enter a user name.
5. In the CHAP Settings section, fill in the following fields:
 - Initiator secret for CHAP node session authentication
 - Target secret for CHAP node session authentication

Note: Leave the credentials field blank if you want the passwords to be generated automatically.
6. Click Create Account.

Note: If an account with the same name exists, an error message will appear.

2.2 Create a QoS Policy

To create a QoS policy, complete the following steps:

1. Log in to the NetApp SolidFire Element UI.
2. Select Management > QoS Policies. QoS policy window opens.
3. Click Create QoS Policy.
4. Enter the QoS policies as required.

For example: minimum 5K, maximum 100K, burst 100K. A proactive monitoring of the systems is performed through NetApp active [IQ analytics](#); there are QoS recommendations for each volume.

Create a New QoS Policy
✕

Policy Name

Quality of Service

| IO Size | Min IOPS | Max IOPS | Burst IOPS |
|---------|-----------|------------|------------|
| 4 KB | 15000 | 10000 | 10000 |
| 8 KB | 9375 IOPS | 62500 IOPS | 62500 IOPS |
| 16 KB | 5556 IOPS | 37037 IOPS | 37037 IOPS |
| 262 KB | 385 IOPS | 2564 IOPS | 2564 IOPS |

| | | |
|---------------|---------------|---------------|
| Max Bandwidth | 699.05 MB/sec | 699.05 MB/sec |
|---------------|---------------|---------------|

2.3 Create a Volume

To create a volume, complete the following steps:

1. Log in to the Element UI.
2. Select Management > Volumes. The Volumes List window opens.
3. Click Create Volume. The Create a New Volume window opens.

Create a New Volume
✕

Volume Details

Volume Name

Volume Size Block Size 512e 4k

Account [Create Account?](#)

Quality of Service

Policy

Custom Settings

4. Enter the volume name (1 to 64 characters in length). For example, provide the name as ORADATA1.

5. Enter the size of the volume.
6. Click the Account drop-down list and select the account that should have access to the volume. In this case, select oracle.
7. In Quality of Service options, select Policy and select the `oracle` policy created previously.
8. Click Create Volume.
9. Repeat steps 1 through 7 for all the remaining volumes (CRS1, CRS2, CRS3, ORADATA2, ORADATA3, and ORALOG1).

2.4 Create Volume Access Groups

Volume access groups limit connectivity from designated host servers based on a unique identifier, whereas CHAP authentication uses secret keys for unidirectional or bidirectional authentication. In this document, initiator iSCSI qualified names (IQNs) are used to access the volumes.

Volume access groups have the following system limits:

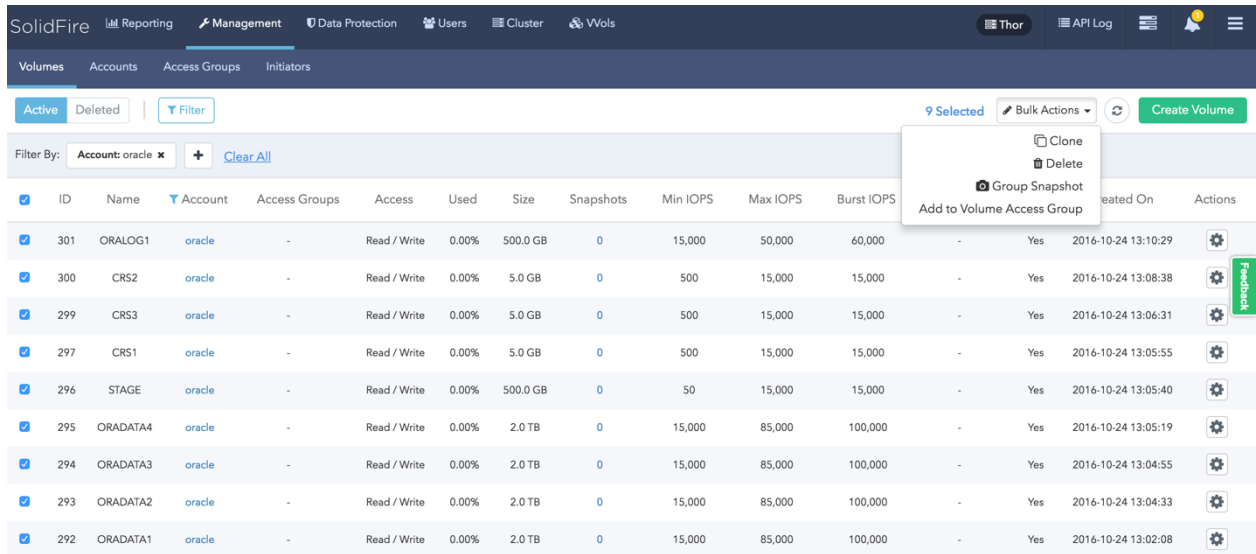
- They can have a maximum of 64 IQNs.
- An IQN can belong to only one access group.
- A single volume can belong to a maximum of four access groups.

To create volume access groups, complete the following steps:

1. Log in to the Element.
2. Select Management > Access Groups. The Access Group window opens.
3. Click Create Access Group. The Create a New Access Group window opens.

4. Enter a name for the volume access group.

5. Select the IQN from the initiator drop-down list or click Create Initiator.
Note: Multiple IQNs are listed for Oracle RAC clusters, depending on the number of nodes.
6. Click Create Access Group.
7. Add the volumes to the access group by selecting Management > Volumes.
8. Select the checkbox to the left of each volume.
9. Near the Create Volume button, click the Bulk Actions drop-down list.
10. Select Add to Volume Access Group.



11. Scroll to the bottom and click Add to Volume Access Group. The Add to Volume Access Group window opens.
12. Select the previously created volume access group from the drop-down list.
13. Click Add to join the selected volumes to the target group.

The Oracle volumes are now listed as part of the selected volume access group. They are now ready for presentation to the host operating system after the host identifier is allowed into the volume access group configuration.

Note: For this configuration, four SolidFire volumes were chosen to distribute data across all four volumes in the ASM disk group. For a web-scale deployment, with hundreds of databases, NetApp recommends having just one volume for data for each individual database and that you control performance with the QoS settings.

3 Operating System Configuration

Before you proceed, follow the recommendations in [TR-4639: Element Software Linux Configuration Guide](#) to configure the operating system. These recommendations apply to Debian, SUSE, and Fedora based Linux distributions. This report assumes that a default installation of a software development workstation exists. You can use alternate distributions if they are fully compatible with the Oracle grid infrastructure and with Oracle Database software.

3.1 Add Packages

After you install the base operating system, you must update it to meet Oracle grid infrastructure and Oracle Database installation requirements. For further information, see the [Oracle documentation](#) and the release requirements, specifically the section “Install Prerequisite Packages.”

3.2 Update Kernel Parameters

Update the kernel parameters for your host operating system to the following values:

```
net.ipv4.ip_local_port_range = 9000 65500
net.core.rmem_default = 4194304
net.core.rmem_max = 16777216
net.core.wmem_default = 262144
net.core.wmem_max = 16777216
net.ipv4.ipfrag_high_thresh = 524288
net.ipv4.ipfrag_low_thresh = 393216
net.ipv4.tcp_rmem = 4096 524288 16777216
net.ipv4.tcp_wmem = 4096 524288 16777216
net.ipv4.tcp_window_scaling = 1
net.core.optmem_max = 524287
net.core.netdev_max_backlog = 2500
net.ipv4.tcp_mem = 16384 16384 16384
fs.aio-max-nr = 1048576
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 0
vm.min_free_kbytes=262144
vm.swappiness=10
```

3.3 Optimize Network Performance

Consider the following guidelines for optimal network performance:

- Enable jumbo frames for all data network interfaces.
- The RAC interconnect should be configured with a different subnet than the RAC virtual IP address (VIP address) so that cluster interconnect traffic is fully isolated.
- Use separate virtual LANs (VLANs) for the management and data traffic.
- Add two virtual network interface cards (vNICs) for data traffic.

3.4 Set User Limits

Oracle recommends that you set the user limits to an optimal number because user limits affect application performance. For this solution, the following user limits were applied for both the `grid` and the `oracle` user. The same settings were applied on all RAC nodes. See the [Oracle documentation](#) for any updates to the following settings:

```
grid soft nproc 2047
grid hard nproc 16384
grid soft nofile 1024
grid hard nofile 65536
oracle soft nproc 2047
oracle hard nproc 16384
oracle soft nofile 1024
oracle hard nofile 65536
```

3.5 Set Permissions

Before the Oracle installation, you should set appropriate permissions so that Oracle can access the underlying devices on the NetApp HCI storage. You can set the device permissions on multipath devices by creating an `udev` rule file that allows the owner of the Oracle Database software to access the devices.

4 Oracle Configuration

4.1 Install Prerequisite Packages

NetApp recommends that you run the `runcluvfy.sh` script that is part of the Oracle installation package under the `grid` directory. This script thoroughly verifies all the required packages for each RAC node. You can run the script from the first RAC node, and it validates all the other nodes that are part of the cluster. The script can be executed as follows and is available in the `installer` directory of the grid installation software:

```
./runcluvfy.sh stage -pre crsinst -n rac1,rac2 -fixup -verbose -asm -asmdev  
/dev/mapper/crs1,/dev/mapper/crs2,/dev/mapper/crs3
```

4.2 Configure Oracle ASM

The storage nodes automatically distribute and protect Oracle Database data across all nodes in the NetApp HCI cluster. For optimal distribution of data, use ASM to stripe data over all the storage volumes. This implementation helps confirm that all nodes can process Oracle writes and that Oracle Database I/O requests are uniformly distributed over all the nodes. As your capacity and performance requirements increase, you can dynamically create and add storage volumes to Oracle ASM disk groups.

If the overall throughput requirement of the database goes beyond 600MBps, NetApp recommends that you create more than one storage volume. You should also configure the ASM software to use a large allocation unit (64MB). The larger allocation unit reduces the amount of metadata and metadata updates that ASM requires to support a given amount of storage.

The storage nodes provide complete fault isolation through a comprehensive shared-nothing architecture. Therefore, ASM mirroring within a single cluster is not appropriate. For extreme fault tolerance, ASM mirroring can be used between clusters.

Note: ASMLib is a requirement on Linux when you use a 4KB native volume. Any volumes that you create with the default 512e are accessible by Oracle ASM.

4.3 Configure ASM Library

Oracle ASMLib simplifies storage deployment and management for Linux OS. Before you assign any LUN or volume to ASMLib, you must set the parameter `ORACLEASM_USE_LOGICAL_BLOCK_SIZE` to `true`.

5 Microsoft Windows Configuration

Follow the recommendations in [TR-4643: Element Software Windows Configuration Guide](#). The Windows OS automatically aligns disk partitions to storage sector boundaries. No specific tuning of the partition table layout or device queue depth settings is required with Windows on NetApp HCI. To configure Microsoft Windows, complete the following tasks:

1. Enable jumbo frames on the network interface cards (NICs).
2. To identify storage devices, enable Microsoft Multipath I/O (MPIO) and configure the Microsoft device-specific module (DSM).
3. Configure at least four iSCSI sessions per storage device.

6 Backup and Recovery Using Storage Snapshot Copies

You can take point-in-time (PiT) snapshots in Element for the volumes that are part of the Oracle Database. The application-consistent snapshots can be used to do database recovery in the event of

data corruption or media failure. SolidFire supports having multiple volumes for the database, and all volumes that are part of the database have the same consistency point during the group snapshot.

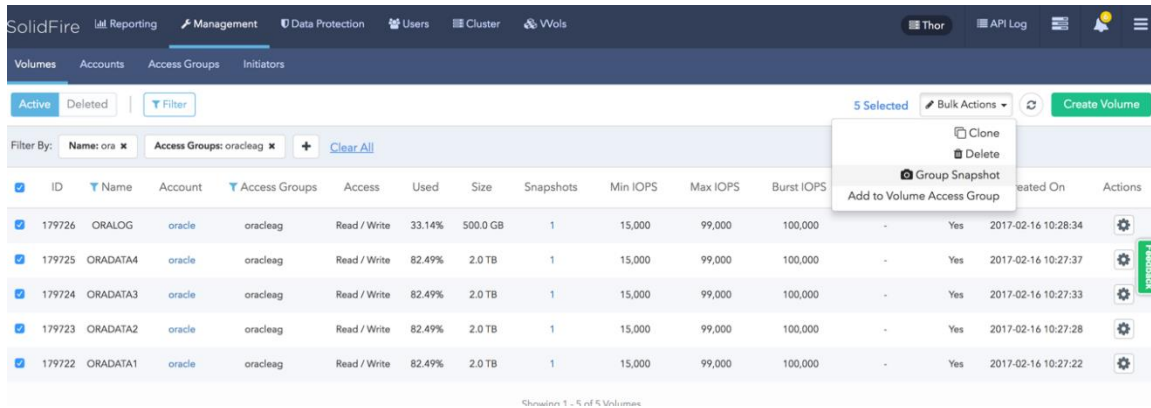
6.1 Create Snapshots (Backup)

To create backup snapshots, complete the following steps:

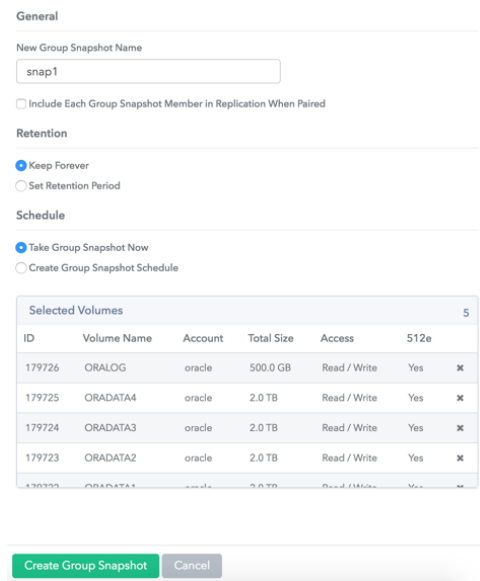
1. Log in to the Oracle instance. To put the database into backup mode, run the following commands:

```
alter system archive log current;
alter database begin backup;
```

2. Log in to the Element UI.
3. Select Management > Volumes. The Volumes List window opens.



4. Select the volumes that are part of the database.
5. Select Bulk Actions > Group Snapshot.



6. In the Create Group Snapshot window, enter a name for the snapshot (Snap1 in this case).
7. Set the desired retention time.
8. Select Take Group Snapshot Now or select Create Group Snapshot Schedule.
9. Scroll to the bottom and click Create Group Snapshot.

10. Take the database out of backup mode.

```
alter database end backup;
```

6.2 Restore Snapshots (Recovery)

Database recovery can be performed by reverting the volumes on the SolidFire array using a snapshot. To recover the database, complete the following steps:

1. Log in to the Oracle Database server, stop the Oracle instance, and unmount the file systems that are part of the database. If the database is provisioned using ASM, stop the ASM instance.
2. Log in to the Element UI.
3. Select Data Protection > Group Snapshots. The Group Snapshot list opens.

| <input checked="" type="checkbox"/> | ID | Name | Account | Access Groups | Access | Used | Size | Snapshots | Min IOPS | Max IOPS | Burst IOPS | Attributes | 512e | Created On | Actions |
|-------------------------------------|--------|----------|---------|---------------|--------------|--------|--------|-----------|----------|----------|------------|------------|------|---------------------|---------|
| <input checked="" type="checkbox"/> | 179725 | ORADATA4 | oracle | oracleag | Read / Write | 82.49% | 2.0 TB | 0 | 15,000 | 99,000 | 100,000 | - | Yes | 2017-02-16 10:27:37 | |
| <input checked="" type="checkbox"/> | 179724 | ORADATA3 | oracle | oracleag | Read / Write | 82.49% | 2.0 TB | 0 | 15,000 | 99,000 | 100,000 | - | Yes | 2017-02-16 10:27:33 | |
| <input checked="" type="checkbox"/> | 179723 | ORADATA2 | oracle | oracleag | Read / Write | 82.49% | 2.0 TB | 0 | 15,000 | 99,000 | 100,000 | - | Yes | 2017-02-16 10:27:28 | |
| <input checked="" type="checkbox"/> | 179722 | ORADATA1 | oracle | oracleag | Read / Write | 82.49% | 2.0 TB | 0 | 15,000 | 99,000 | 100,000 | - | Yes | 2017-02-16 10:27:22 | |

4. Select the volumes that are part of the database.
5. Select Bulk Actions > Group Snapshot.
6. Click Actions.
7. Select Rollback Volumes to Group Snapshot. The Rollback to Group Snapshot window opens.

Rollback To Group Snapshot

Group Snapshot Details
ID: 20 Name: snap1
UUID: c4bf3fc-80fb-47e9-9d5c-f5dcc75e278d
Created On: 2017-03-14 14:01:05

Save volumes' current state as a group snapshot

New Group Snapshot Name

| ID | Name | Size | Volume ID | Volume Name | Account | Access | Retain Until | Replication Enabled | Replicated |
|--------|-------|----------|-----------|-------------|---------|--------------|--------------|---------------------|------------|
| 182717 | snap1 | 2.0 TB | 179723 | ORADATA2 | oracle | Read / Write | - | No | - |
| 182718 | snap1 | 2.0 TB | 179724 | ORADATA3 | oracle | Read / Write | - | No | - |
| 182716 | snap1 | 2.0 TB | 179722 | ORADATA1 | oracle | Read / Write | - | No | - |
| 182719 | snap1 | 2.0 TB | 179725 | ORADATA4 | oracle | Read / Write | - | No | - |
| 182720 | snap1 | 500.0 GB | 179726 | ORALOG | oracle | Read / Write | - | No | - |

8. Click Rollback Group Snapshot.
9. After the volume is reverted successfully, log in to the Oracle instance.
10. Rescan the iSCSI devices.
11. Mount the file systems.
12. If the database is provisioned using ASM, start the ASM instance.
13. Log in to the Oracle instance and run these commands:

```
startup mount;
recover database;
alter database open;
```

7 RMAN Merge

RMAN merged incremental backups can reduce the time for Oracle Database recovery and are suitable primarily for disk-based backup and recovery. This method is widely used by database administrators. As part of the RMAN setup, a level 0 (full image) copy is created. Afterward, the subsequent RMAN runs and creates a level 1 incremental backup of the data file by using the system change number (SCN) as a reference. RMAN then applies all level 1 backups taken since the last backup to the image. The RMAN process uses a tag to link the incremental backups to the image copy. The configuration for this document is done with level 0 and incremental backups using two different locations on the RMAN target file system.

The target file system of the RMAN dump is created on the SolidFire array and is mapped to the database server with permission set so that an Oracle user can perform reads and writes. Four SolidFire volumes were used for this configuration, and the QoS settings were configured to meet the appropriate throughput requirements. On the Linux host, LVM was used to stripe data across multiple volumes.

7.1 Enable Block Change Tracking

Enabling the block change tracking feature improves the performance of incremental backups by recording changed blocks in the block change tracking file. During a backup operation, RMAN uses this file to identify the changed blocks that need to be backed up instead of scanning all data blocks to identify which blocks have changed. To enable block change tracking, connect to the RMAN target database instance and run the following commands:

```
ALTER SYSTEM SET DB_CREATE_FILE_DEST = '+ORADATA'; #The destination of the file changes depending
on the your setup
ALTER DATABASE ENABLE BLOCK CHANGE TRACKING;
```

7.2 Create RMAN Incremental

To perform the RMAN merge process incorporating SolidFire snapshots, complete the following steps:

1. Set the RMAN configuration parameters for the level 0 image.

```
CONFIGURE DEFAULT DEVICE TYPE TO DISK;
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE BACKUP OPTIMIZATION ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO
'/rmanfs/controlfile/control.%F.ctl';
CONFIGURE DEVICE TYPE DISK PARALLELISM 8 BACKUP TYPE TO BACKUPSET;
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/rmanfs/dbimage/full/%U.dbf';
```

2. Run the RMAN backup with the following commands. The first backup is a level 0 backup and takes some time, depending on the amount of data to be copied.

```
RUN {
  BACKUP INCREMENTAL LEVEL 1 FOR RECOVER OF COPY WITH TAG 'ORADB_INCR' DATABASE;
  RECOVER COPY OF DATABASE WITH TAG 'ORADB_INCR';
}
```

3. After a successful level 0 backup, create a snapshot on the SolidFire storage for all the volumes that are part of the RMAN file system.
4. Configure the RMAN parameter to update the location for a level 1 backup.

```
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/rmanfs/dbimage/incremental/%U.dbf';
```

5. Create a level 1 backup image by running the same commands from step 2. To simplify this step, you can save the commands to a script file to use on the RMAN command line. You can take the subsequent level 1 backups, depending on the backup schedule.
6. Back up the archive logs. The following command backs up all the archive logs currently known in the control file that have not been backed up to disk yet. After the backup of the archive logs is finished, the control file is updated.

```
BACKUP ARCHIVELOG ALL NOT BACKED UP 1 TIMES;
```

7. Schedule the snapshot on the SolidFire storage to run right after step 6 for all the volumes that are part of the RMAN file system.

7.3 Use Cases for RMAN Merge

- Users can switch to the RMAN merge file, recover the database to a desired point in time, and significantly reduce the recovery time.
- Users can also use a disk-based image copy.
- The snapshots that are taken on the RMAN image can be used to create clone copies of the database easily. The clone copy of the database is created using the final merge image, which has all the changes from the last level 1 backup operation.
- The archive logs can also be applied to the clone copy to apply the latest changes from production.

A sample output of the RMAN target file system is provided in the appendix.

8 Database Cloning

Creating usable, space-efficient, and time-efficient database copies quickly and with virtually no effect on the production system is important. Element volume cloning is a proven technology that helps database and system administrators deliver a near-instantaneous, space-efficient, point-in-time copy of the production database. Traditional methods of copy processing pose various challenges, including system downtime and degraded performance during the cloning process. Also, a large amount of storage space is required to store each clone. The Element volume cloning process is completed quickly, with virtually no performance effect on the production system.

8.1 Clone a Volume

To clone a volume, complete the following steps:

1. Log in to the Element UI.
2. Select Data Protection > Group Snapshots. Right-click the Actions button for the snapshots. In this case, we select Snap1, the snapshot taken in section 8.1.
3. The Clone Volumes from Group Snapshot window opens.

| ID | UUID | Name | Create Time | Status | # Volumes | Retain Until | Remote Replication | Actions |
|----|--------------------|-------------------|---------------------|--------|-----------|--------------|--------------------|-------------|
| 20 | c4bf3fc-80fb-47... | snap1 | 2017-03-14 14:01:05 | done | 5 | - | disabled | [Gear Icon] |
| 18 | 952f293e-470f-4... | Snap1-forClone | 2017-02-16 18:52:50 | done | 5 | - | | [Gear Icon] |
| 16 | d331fa1e-aadc-... | 2017-02-13T16:... | 2017-02-13 11:29:50 | done | 6 | - | | [Gear Icon] |
| 14 | 22570848-1a4b-... | 2017-02-10T00:... | 2017-02-09 19:45:20 | done | 6 | - | | [Gear Icon] |

4. Click the gear icon and then select Clone Volumes from Group Snapshot.

- Enter a prefix (CL in this example).
- Provide the account and access info, depending on what level of access you require for the clone volumes.

Clone Volumes From Group Snapshot ✕

New Volume Clone Details

New Volume Name Prefix

Account

Access

| Group Snapshot Members | | | | | | | | | |
|------------------------|-------|----------|-----------|-------------|---------|--------------|--------------|---------------------|------------|
| ID | Name | Size | Volume ID | Volume Name | Account | Access | Retain Until | Replication Enabled | Replicated |
| 182717 | snap1 | 2.0 TB | 179723 | ORADATA2 | oracle | Read / Write | - | No | - |
| 182718 | snap1 | 2.0 TB | 179724 | ORADATA3 | oracle | Read / Write | - | No | - |
| 182716 | snap1 | 2.0 TB | 179722 | ORADATA1 | oracle | Read / Write | - | No | - |
| 182719 | snap1 | 2.0 TB | 179725 | ORADATA4 | oracle | Read / Write | - | No | - |
| 182720 | snap1 | 500.0 GB | 179726 | ORALOG | oracle | Read / Write | - | No | - |

- Click Start Cloning.
- The cloned volumes can be viewed in the volumes list.

Filter By: Name: CLO ✕ + [Clear All](#)

| <input type="checkbox"/> | ID | Name | Account | Access Groups | Access | Used | Size | Snapshots | Min IOPS | Max IOPS | Burst IOPS |
|--------------------------|--------|------------|---------|---------------|--------------|--------|----------|-----------|----------|----------|------------|
| <input type="checkbox"/> | 179733 | CLORALOG | oracle | oraVM-CL | Read / Write | 32.51% | 500.0 GB | 0 | 15,000 | 99,000 | 100,000 |
| <input type="checkbox"/> | 179732 | CLORADATA4 | oracle | oraVM-CL | Read / Write | 82.49% | 2.0 TB | 0 | 15,000 | 99,000 | 100,000 |
| <input type="checkbox"/> | 179731 | CLORADATA3 | oracle | oraVM-CL | Read / Write | 82.49% | 2.0 TB | 0 | 15,000 | 99,000 | 100,000 |
| <input type="checkbox"/> | 179730 | CLORADATA2 | oracle | oraVM-CL | Read / Write | 82.49% | 2.0 TB | 0 | 15,000 | 99,000 | 100,000 |
| <input type="checkbox"/> | 179729 | CLORADATA1 | oracle | oraVM-CL | Read / Write | 82.49% | 2.0 TB | 0 | 15,000 | 99,000 | 100,000 |

- After the clone process is complete, add the newly created volumes to the volume access group of the database server (target) where you are planning to host the cloned copy.
- Log in to the target host.
- To rescan the iSCSI devices to present the newly cloned volumes, run the following commands:

```
iscsiadm -m discovery -t sendtargets -p <SolidFire SVIP> --op update -n node.session.nr_sessions -v 2 -> where "n" is the number of paths
```



```
iscsiadm -m node -L all
multipath -r
```

12. Set the permission for the Oracle user or the grid user.
13. Mount the file systems or start the ASM instance.
14. Log in to the Oracle instance.

```
startup mount;
recover database;
alter database open;
```

Appendix A: RMAN Merge Output

```
-rw-r-----. 1 oracle oinstall 6597069774848 Mar 13 09:33 data_D-ORADB_I-2689361952_TS-SLOB_FNO-
5_09ruruln.dbf
-rw-r-----. 1 oracle oinstall      17080320 Mar 13 09:33 0orv12rp_1_1.dbf
-rw-r-----. 1 oracle oinstall      40960 Mar 13 09:33 0prv12rp_1_1.dbf
-rw-r-----. 1 oracle oinstall      8839168 Mar 13 09:33 0nrv12rp_1_1.dbf
-rw-r-----. 1 oracle oinstall     56811520 Mar 13 09:33 0mrv12rp_1_1.dbf
-rw-r-----. 1 oracle oinstall     1671168 Mar 13 10:14 0trv157p_1_1.dbf
-rw-r-----. 1 oracle oinstall      40960 Mar 13 10:14 0urv157p_1_1.dbf
-rw-r-----. 1 oracle oinstall     286720 Mar 13 10:14 0srv157p_1_1.dbf
-rw-r-----. 1 oracle oinstall     3702784 Mar 13 10:14 0rrv157p_1_1.dbf
-rw-r-----. 1 oracle oinstall     5251072 Mar 13 11:43 data_D-ORADB_I-2689361952_TS-USERS_FNO-
6_0druruln.dbf
-rw-r-----. 1 oracle oinstall     157294592 Mar 13 11:43 data_D-ORADB_I-2689361952_TS-
UNDOTBS1_FNO-4_0cruruln.dbf
-rw-r-----. 1 oracle oinstall     838868992 Mar 13 11:43 data_D-ORADB_I-2689361952_TS-SYSTEM_FNO-
1_0bruruln.dbf
-rw-r-----. 1 oracle oinstall     1027612672 Mar 13 11:43 data_D-ORADB_I-2689361952_TS-SYSAUX_FNO-
3_0aruruln.dbf
-rw-r-----. 1 oracle oinstall      1900544 Mar 13 11:43 12rv1aep_1_1.dbf
-rw-r-----. 1 oracle oinstall      40960 Mar 13 11:43 13rv1aep_1_1.dbf
-rw-r-----. 1 oracle oinstall      294912 Mar 13 11:43 11rv1aep_1_1.dbf
-rw-r-----. 1 oracle oinstall     4153344 Mar 13 11:43 10rv1aep_1_1.dbf
-rw-r-----. 1 oracle oinstall      40960 Mar 13 13:08 0lrv12rp_1_1.dbf
-rw-r-----. 1 oracle oinstall      40960 Mar 13 14:02 0qrv157p_1_1.dbf
-rw-r-----. 1 oracle oinstall      40960 Mar 13 15:07 0vrv1aep_1_1.dbf
```

Where to Find Additional Information

To learn more about the information that is described in this document, review the following documents and websites:

- DS-3881: NetApp HCI Datasheet
<https://www.netapp.com/us/media/ds-3881.pdf>
- TR-4644: NetApp SolidFire Quality of Service (QoS)
<https://www.netapp.com/us/media/tr-4644.pdf>
- TR-4651: NetApp SolidFire SnapMirror Architecture and Configuration
<https://www.netapp.com/us/media/tr-4651.pdf>
- TR-4641: NetApp HCI Data Protection
<https://www.netapp.com/us/media/tr-4641.pdf>
- TR-4639: Element Software Linux Configuration Guide
www.netapp.com/us/media/tr-4639.pdf
- NVA-1122: VMware Private Cloud on NetApp HCI
www.netapp.com/us/media/nva-1122-design.pdf
www.netapp.com/us/media/nva-1122-deploy.pdf
- TR-4748: Build Your NetApp Data Fabric
www.netapp.com/us/media/tr-4748.pdf

- TR-4643: Element Software Windows Configuration Guide
<https://www.netapp.com/us/media/tr-4643.pdf>
- NetApp product documentation
<https://docs.netapp.com>

Version History

| Version | Date | | Document Version History |
|-------------|------------|--|--|
| Version 1.0 | April 2019 | | Initial document release. |
| Version 2.0 | July 2020 | | Merged Oracle on SolidFire (TR-4606) material; reviewed content. |

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2019–2020 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4768-0820