



Technical Report

Sybase ASE on NetApp HCI

Bobby Oommen, NetApp

October 2018 | TR-4721

Abstract

SAP Adaptive Server Enterprise (SAP ASE) requires an infrastructure platform that can perform true scaling, especially when customers want to virtualize the infrastructure stack. The NetApp® HCI (Hyper Converged Infrastructure) solution for infrastructure administrators offers an easy way to consume infrastructure resources and scale, starting low and growing incrementally as the application ecosystem grows. This document highlights use cases for SAP ASE and describes how to deploy SAP ASE database components on a NetApp HCI system.

TABLE OF CONTENTS

1	Introduction.....	3
1.1	Configuration.....	3
2	Application Use Cases.....	4
3	SAP ASE Provisioning on NetApp HCI	4
3.1	Storage Layout.....	4
3.2	Storage Configuration	5
3.3	OS Configuration for SAP ASE	8
4	SAP ASE – Backup and Recovery	11
4.1	ASE Snapshot Backup	11
4.2	ASE Snapshot Recovery.....	11
4.3	Disaster Recovery	11
4.4	RTO and RPO Guidelines	12
4.5	Disaster Recovery Process for SAP ASE	12
5	Disaster Recovery Operations for SAP ASE	13
5.1	Failover – Activating the Disaster Recovery Server for Writes.....	13
5.2	Failback – Reversing Replication.....	13
6	Summary	13
	Where to Find Additional Information	13
	Version History	14

1 Introduction

SAP ASE has been widely adopted by many users, especially in the financial services industry. It offers critical value to enterprises by providing high availability and superior management of infrastructure resources. To address increasingly dynamic business requirements, customers are looking for an infrastructure that can be easily provisioned and consumed. NetApp HCI helps you build a hyper converged private-cloud strategy that is capable of transforming and empowering your organization so that you can move faster, drive operational efficiency, and reduce costs.

You can deploy NetApp HCI in minutes, thus eliminating the complex management of traditional three-tier architectures. NetApp HCI is an enterprise-scale hyper converged infrastructure solution that is ideally suited for customers who want to break free from some of the limitations of first-generation HCI, including scaling compute and storage resources independently and avoiding costly and inefficient overprovisioning. With NetApp HCI, compute and storage resources are separate. Therefore, infrastructure administrators can scale compute resources when a specific SAP ASE instance needs more compute power, and they can also independently scale storage resources to increase capacity.

With NetApp HCI, you can run multiple applications with guaranteed performance by using the robust Quality of Service (QoS) feature. QoS allows different application components to safely share resources in a single infrastructure system.

1.1 Configuration

NetApp HCI is available with multiple configuration options—small, medium, and large for compute and storage—and thus serves essentially as a small blade inside a chassis. Figure 1 details the configuration specifications.

From the configuration depicted in Figure 1, you can go from 5TB to 44TB of effective capacity. For compute resources, you can go from 16 cores to 36 cores and from 256Gb to 768Gb of RAM with the flexibility of mixing and matching compute or storage nodes independently. For example, if you have a setup with 2 compute nodes and 4 storage nodes, you can add more compute resources without adding any storage just by adding a compute blade to the chassis. Having this scalability and flexibility enables you to build an efficient and agile cloud in your data center.

Figure 1) NetApp HCI configuration.

Storage Nodes	Small	Medium	Large	Compute Nodes	Small	Medium	Large
	RU	1RU, Half-width			RU	1RU, Half-width	
SSD's	(6) 480GB	(6) 960GB	(6)1.92TB	Core for VM's	16	24	36
NVRAM	8GB	8GB	8GB	GB Memory for VM's	256	512	768
Networking	2x 25/10GbE SFP28 2 x 1GbE RJ45	2x 25/10GbE SFP28 2 x 1GbE RJ45	2x 25/10GbE SFP28 2 x 1GbE RJ45	Networking	4x 25/10GbE SFP8 2 x 1GbE RJ45	4x 25/10GbE SFP28 2 x 1GbE RJ45	4x 25/10GbE SFP28 2 x 1GbE RJ45
Effective Block Capacity@ 5x-10x Efficiency	5.5TB – 11TB	11TB – 22TB	22TB – 44TB				

2 Application Use Cases

NetApp HCI can support a wide range of database application use cases, including database consolidation, dynamic resource allocation for creating development and test environments, integrating cloud infrastructure, and more. The per-volume QoS controls for storage nodes help individual databases get the I/O throughput that they need without being affected by other databases that run in parallel on the same HCI setup. With QoS and data reduction efficiencies, you can achieve higher database density with the shared storage infrastructure by having several database instances. VMware administrators have full control of each storage volume on which the database resides. Therefore, they can perform all maintenance operations, including setting the QoS policy for each database copy through the vCenter NetApp Element Software Plug-In. Administrators can also use REST APIs to achieve full automation and make storage management simpler and easier.

NetApp storage Snapshot™ copies provide a point-in-time view of the contents of an active file system or storage volume. You can use these copies for rapid recovery of corrupted datasets and to create space-efficient copies of datasets for development and testing use cases. The cloning process can be coupled with NetApp SolidFire® Element® software QoS control so that database clones can coexist with the production copies without any performance effects on the upstream applications. NetApp HCI is Data Fabric ready, which increases agility by expanding application resources across a wide range of cloud environments.

3 SAP ASE Provisioning on NetApp HCI

3.1 Storage Layout

The storage layout described in this section was used for this setup and varies depending on the application requirements. To best leverage all the resources storage nodes, consider the following issues:

- NetApp recommends isolating database workloads with separate volumes for each database.
- Separate volumes for each database provide the granularity necessary to set the QoS settings by volume based on throughput requirements.
- NetApp recommends that you not create any Linux or host-side partition for the LUNs presented from the storage nodes.
- NetApp recommends starting with the same minimum QoS for all SAP ASE instances and then making changes to the QoS settings based on workload requirements. The operating system on the storage cluster performs volume balancing based on the minimum QoS, and having the same QoS settings creates an ideal volume distribution across the storage nodes.

Table 1) Storage volumes.

Volume Name	Size in GB	QoS		
		Min	Max	Burst
system	5	1000	10000	15000
sybase-data1	1000	2000	30000	35000
sybase-data2	1000	2000	30000	35000
Sybase-log	1000	2000	30000	35000

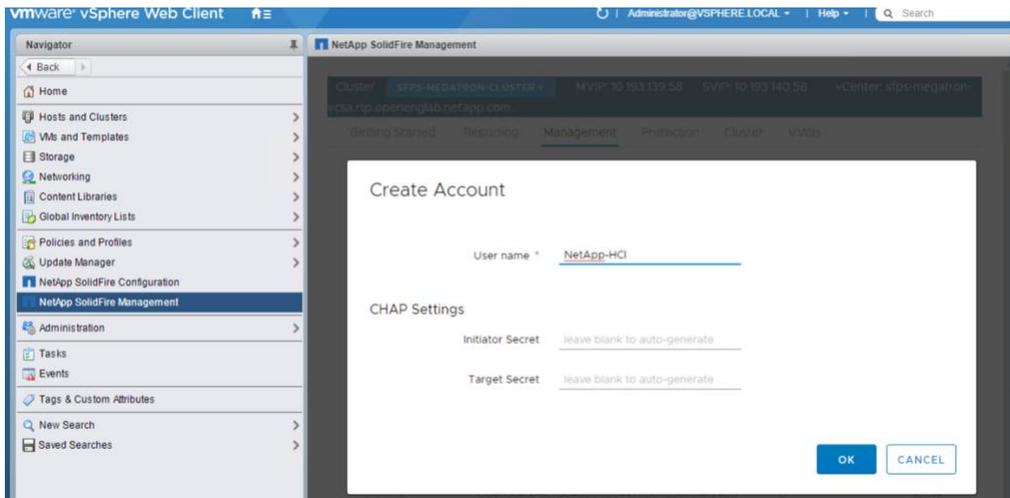
3.2 Storage Configuration

This section shows how to configure storage volumes to support an SAP ASE database. NetApp recommends that you have all the database components on the storage node. NetApp supports presenting the storage in a 4K sector size (native mode) and in a 512-byte sector size (512e).

Create an Account

To create an account, complete the following steps:

1. Log in to the vCenter as an administrator.
2. Select Home → NetApp SolidFire Management.
3. In the NetApp SolidFire Management pane on the right, select Management → Account → Create Account. The Create Account window opens.



4. Enter a user name. In this case, we used NetApp-HCI.
5. In the CHAP Settings section, enter the following:
 - a. The initiator secret for CHAP node session authentication.
 - b. The 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 is returned.

Create a Volume

To create a volume, complete the following steps:

1. Log in to vCenter as an administrator.
2. Select Home → NetApp SolidFire Management.
3. In the NetApp SolidFire Management pane on the right, select Management → Volumes → Create. The Create Volume window opens.

Create Volume

Volume Name * Sybase-data1

Volume Size * 1000 GB

Account * NetApp-HCI

512e

Quality of Service

I/O Size	Min IOPS	Max IOPS	Burst IOPS
4 KB	15000	100000	100000
8 KB	9375 IOPS	62500 IOPS	62500 IOPS
16 KB	5556 IOPS	37037 IOPS	37037 IOPS
256 KB	385 IOPS	2564 IOPS	2564 IOPS

Max Bandwidth 699 MB/s 699 MB/s

[RESET DEFAULT QOS](#)

4. Enter the volume name (from 1 to 64 characters in length). For example, enter the name Sybase-data1.
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 NetApp-HCI.
7. Enter the QoS values. For this setup, the following values were chosen: Min IOPS 15,000, Max IOPS 100,000, and Burst IOPS 100,000.
8. Click Create Volume.

Repeat steps 1 through 8 for all the remaining volumes that are part of the database.

Create Initiators

To create an initiator, complete the following steps:

1. Log in to vCenter as an administrator.
2. Select Home → NetApp SolidFire Management.
3. In the NetApp SolidFire Management pane on the right, select Management → Initiators → Create. The Create Initiator window opens.

4. From the Windows host, get the initiator name and enter it in the IQN/WWPN field. Optionally, you can also provide an alias.
5. Click OK.

Create Volume Access Groups

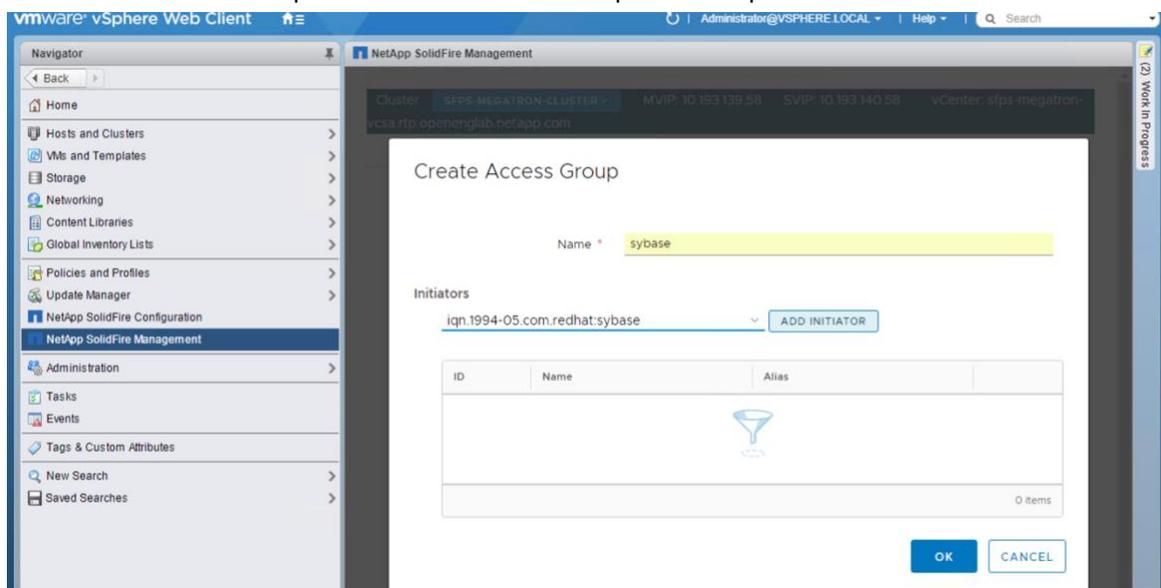
Volume access groups limit connectivity from designated host servers based on a unique identifier, whereas CHAP authentication uses secret keys for unidirectional or bidirectional authentication. In this document, initiator iSCSI Qualified Names (IQNs) are used to access the volumes.

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 vCenter as an administrator.
2. Select Home → NetApp SolidFire Management.
3. In the NetApp SolidFire Management pane on the right, select Management → Access Groups → Create Access Group. The Create Access Group window opens.

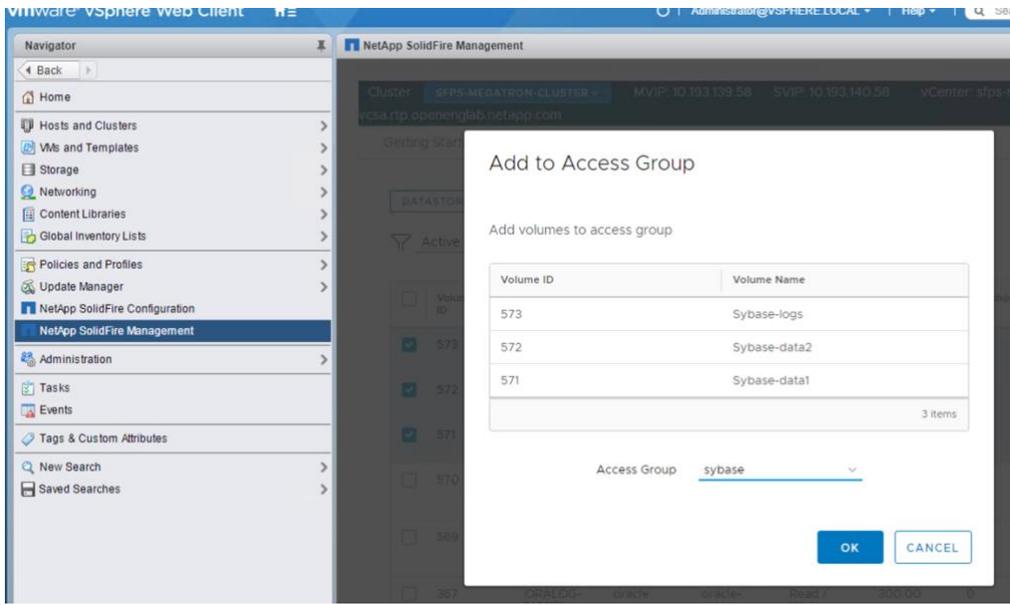


4. Enter the access group name, which is Sybase in this case.
5. From the dropdown list, select the initiator that was created in “Create Initiators,” step 3.
6. Click OK.

Add Volumes to Access Groups

To add the volumes to the access group, complete the following steps:

1. Select Management → Volumes.
2. From the listed volumes, select the volumes that are part of the SAP ASE database.
3. Click the Bulk Actions drop-down list.
4. Select Add to Access Group. The Add to Access Group window opens.



5. From the Access Group drop-down list, select the Access group sybase.
6. Click OK.

3.3 OS Configuration for SAP ASE

The guidelines in this document apply to Red Hat Enterprise Linux 7.x distributions of the SAP ASE Enterprise edition. Alternate distributions can be used if they have full compatibility with the SAP ASE software.

Update Kernel Parameters

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

```
vm.dirty_ratio = 15
vm.dirty_background_ratio = 5
vm.swappiness = 1
net.core.somaxconn = 4096
net.ipv4.tcp_fin_timeout = 30
net.ipv4.tcp_keepalive_intvl = 30
net.ipv4.tcp_keepalive_time = 120
net.ipv4.tcp_max_syn_backlog = 4096
```

Network Tuning

Consider the following guidelines for optimal network performance:

- Enable jumbo frames for all host network interfaces.
- To isolate data traffic, configure the interface that is used for data traffic with a different subnet from the public network.

iSCSI Tuning

Linux iSCSI initiator configuration works with storage volumes in its default configuration. To maximize system throughput, increase the number of sessions per target (`nr_sessions`) from the default of 1 to 8.

1. Make the following changes to the iSCSI daemon in the `/etc/iscsi/iscsid.conf` file:

```
iscsid.startup = /etc/rc.d/init.d/iscsid force-start
```

```

node.startup = automatic
node.leading_login = No
node.session.timeo.replacement_timeout = 120
node.conn[0].timeo.login_timeout = 15
node.conn[0].timeo.logout_timeout = 15
node.conn[0].timeo.noop_out_interval = 5
node.conn[0].timeo.noop_out_timeout = 5
node.session.err_timeo.abort_timeout = 15
node.session.err_timeo.lu_reset_timeout = 30
node.session.err_timeo.tgt_reset_timeout = 30
node.session.initial_login_retry_max = 8
node.session.cmds_max = 128
node.session.queue_depth = 32
node.session.xmit_thread_priority = -20
node.session.iscsi.InitialR2T = No
node.session.iscsi.ImmediateData = Yes
node.session.iscsi.FirstBurstLength = 262144
node.session.iscsi.MaxBurstLength = 16776192
node.conn[0].iscsi.MaxRecvDataSegmentLength = 262144
node.conn[0].iscsi.MaxXmitDataSegmentLength = 0
discovery.sendtargets.iscsi.MaxRecvDataSegmentLength = 32768
node.conn[0].iscsi.HeaderDigest = None
node.session.iscsi.FastAbort = Yes
node.startup = automatic
node.session.nr_sessions = 8

```

2. Make discovery of iSCSI devices persistent over reboots:

```
chkconfig iscsid
```

3. To rescan the new storage volumes, run the following commands:

```

iscsiadm -m discovery -t sendtargets -p <SVIP> --op update -n node.session.nr_sessions
-v 2
iscsiadm -m node -L all

```

I/O Scheduler Tuning

Run the following commands to tune the Linux operating system so that it takes advantage of the performance characteristics of the storage nodes (<devpath> is the device name).

```

echo 0 > /sys/<devpath>/queue/rotational
echo noop > /sys/<devpath>/queue/scheduler
echo 128 > /sys/<devpath>/queue/nr_requests
echo 2 > /sys/<devpath>/queue/rq_affinity
echo 0 > /sys/<devpath>/queue/add_random

```

Multipath Configuration

Configure the Linux multipath driver (multipathd) by making the following changes to the /etc/multipath.conf file.

```

defaults {
    user_friendly_names yes
}

devices {
    device {
        vendor "SolidFir"
        product "SSD SAN"
        path_grouping_policy multibus
        path_checker tur
        hardware_handler "0"
        failback_immediate
        rr_weight uniform
        rr_min_io 10
        rr_min_io_rq 10
        features "0"
    }
}

```

```
        no_path_retry 24
        prio const
    }
}
```

Optionally, you can enable persistent mapping of `/dev/mapper` entries by associating the device's worldwide identifier (WWID) with a specific operating system alias. For this option, make the following additions to the `/etc/multipath.conf` file:

```
multipaths {
    multipath {
        wwid 36f47acc100000000707a646c000003b1
        alias sybase-data1
    }

    multipath {
        wwid 36f47acc100000000707a646c000003b2
        alias sybase-data2
    }

    multipath {
        wwid 36f47acc100000000707a646c000003b3
        alias sybase-logs
    }
}
```

To enable multipathing, run the following command:

```
systemctl enable multipathd.service
```

You can check the status of the multipath daemon with the following command:

```
systemctl list-unit-files|grep multipath
multipathd.service      enabled
```

File System Setup with a Logical Volume Manager

To configure a logical volume manager (LVM) to stripe data across multiple storage volume devices, complete the following steps:

1. Create a volume group by using the multipath devices `mpatha` and `mpathb`.

```
vgcreate datavg /dev/mapper/mpatha /dev/mapper/mpathb
```

2. Create a logical volume on this volume group.

```
lvcreate -l 100%FREE -n datalv datavg
```

3. Create an XFS file system.

```
mkfs.xfs /dev/datavg/datalv
```

4. Create a database directory path for SAP ASE and mount the file system.

```
mkdir -p /data
mount -t ext4 -o nobarrier,discard,noatime /dev/datavg/datalv /data/
```

Set the appropriate permissions so that the ASE user can access the devices on the storage nodes. You can set permissions on multipath devices by creating a `udev` rule file that allows appropriate access to the devices.

4 SAP ASE – Backup and Recovery

SAP supports customers who want to use external data protection processes (for example, storage Snapshot copies) for SAP ASE. Using Snapshot copies for data protection provides faster backup and recovery and helps customers restore applications to meet stringent SLAs. The storage APIs that come with the Element OS are used to create backup and recovery workflows in conjunction with the databases to create application-consistent Snapshot copies. SAP Note [1887068](#) lists the steps needed when using storage Snapshot copies.

4.1 ASE Snapshot Backup

Before taking a storage Snapshot copy, the database is put into a consistent state by quiescing the user database and master database. The following steps show how to perform database backup by using a storage Snapshot copy.

1. Quiesce the ASE database to put the application into a consistent state:

```
quiesce database snap1 hold MYDB for external dump
```

2. Perform a Snapshot operation on all the database volumes on the storage nodes. You can do this either through vCenter or by using Element OS APIs.
3. Put the database back into normal mode:

```
Quiesce database release
```

4. Take a transaction dump of the database to save the log files. This transaction dump can be used later to apply logs during a Snapshot recovery of the database.

```
dump tran MYDB to /stage
```

4.2 ASE Snapshot Recovery

If you must use a Snapshot copy to recover the database, complete the following steps:

1. Shut down the SAP ASE instance by logging in as a Sybase user:

```
shutdown
```

2. Unmount the file system where the database is hosted. For example, to unmount the data and log file systems, the following commands were used:

```
umount /data  
umount /logs
```

3. From vCenter or by using Element APIs, roll back the volume by using the appropriate Snapshot copy.
4. Start the SAP ASE instance by using the `-q` option so that the ASE instance puts the database into the Load Pending state.

4.3 Disaster Recovery

NetApp HCI supports three replication modes, and choosing the one to use depends on application requirements in terms of RPO and RTO.

- Synchronous (sync) replication
- Asynchronous (async) replication
- Snapshot based replication

For sync replication, the application write order is preserved because the writes from the host are acknowledged only after writing to the destination storage cluster that is part of the replication setup.

Application-consistent Snapshot copies can be taken through API calls with scripting during async or Snapshot based replication so that those copies are used to recover the database during a disaster.

The following sequence of steps occurs during a replication process:

- Volume metadata is compared between the source and the target.
- Target metadata is updated to match the source volume.
- Blocks are transferred from the source until the volumes reach an active state.
- After the volumes are active, all future writes are transferred to the target volume.

For Snapshot copy-only replication, the copies are replicated to the target cluster, and data blocks are not written to the target volume until a rollback of the copy is performed on the target.

4.4 RTO and RPO Guidelines

- Use sync replication when the RPO requirement is zero. Sync replication enters a PauseDisconnected state if there is a network issue between the source and target clusters.
- Async replication is recommended when the RPO requirement is less than a minute. Replication enters a PauseDisconnected state if there is a network issue between the source and the target.
- For RPO requirements of more than a minute, Snapshot copy-based replication is recommended. The RPO and RTO for Snapshot copy-based replication depends on the frequency of Snapshot creation at the source as well as the change rate at the source.
- The RTO is not zero because downtime is required for the switchover of the host servers to the disaster recovery server for both sync and async replication. The amount of downtime can be reduced by automating the disaster recovery workflow with Element software APIs.

4.5 Disaster Recovery Process for SAP ASE

When the primary storage cluster goes offline or is not accessible, the destination storage volumes are activated for writes so that applications can be brought online. All the steps can be controlled by using APIs from a system that has access to the primary and destination storage clusters.

Steps for Sync and Async Replication

1. Pause replication from the target cluster.
2. Change the access mode to read/write.
3. Map the volumes to the target host (if this has not been done already).
4. Rescan any iSCSI devices.
5. Start the SAP ASE instances. The database performs a recovery and applications do not lose any data as long as the replication pairs are in the active state.

Steps for Snapshot Copy Only Replication

1. Pause replication from the target cluster.
2. Change the access mode to read/write.
3. Map the volumes to the target host (if this has not been done already).
4. Rescan any iSCSI devices.
5. Roll back to the last known good Snapshot copy for data and log volumes.
6. Start SAP ASE with the `-q` option so that the database is in the Load Pending state.
7. Apply the latest transaction by using the Load command. The transaction dump location (a separate volume from the database volumes) must be part of the replication process in order to complete this step.

8. Start the SAP ASE instances. The database performs a recovery and is current until the last applied transaction log. The database can be kept as current as possible by taking frequent transaction dumps.

5 Disaster Recovery Operations for SAP ASE

For this setup, two HCI clusters were used to configure the disaster recovery (DR) configuration, and the following requirements were met:

- Cluster pairing was completed between the two HCI clusters.
- Volume pairing was completed for all volumes that are part of the SAP ASE database.
- The sync delay between the volume pairs was zero.
- A sample script was executed to create load on the database by doing 100% writes/inserts.

5.1 Failover – Activating the Disaster Recovery Server for Writes

To simulate disaster recovery, the replication link was taken down to create a replication status of `PauseDisconnected`, and the following steps were completed:

1. Stopped SAP ASE on the primary server.
2. Paused replication on volume pairs on the DR server.
3. Changed the volume access to Read/Write from Replication Target on the DR server.

5.2 Failback – Reversing Replication

After the DR server is activated, the database receives updates from production, and the DR server then becomes the new production server. In addition, the original production server becomes the new DR server. The following steps were followed to make the original production server the new DR server.

1. Unmount any file systems on the original production SAP ASE host.
2. Log out of iSCSI sessions on the original production SAP ASE host.
3. Change the volume access to Replication Target from Read/Write on the original production server.

6 Summary

Enterprise customers are looking at simplified infrastructure deployments to meet the changing and dynamic demands of data centers. They are moving away from dedicated platforms for potential workloads and are trying to avoid overprovisioning to increase efficiency and reduce cost. NetApp HCI offers a comprehensive solution that provides granular control over every application and satisfies the most demanding performance requirements. The underlying storage nodes of NetApp HCI use all-flash media coupled with capacity thin provisioning and in-line data efficiency features. These features yield significant efficiency and agility when deploying applications and help businesses consolidate their workloads with confidence.

Where to Find Additional Information

To learn more about the information described in this document, refer to the following resources:

- NetApp HCI Resource page
<https://mysupport.netapp.com/info/web/ECMLP2831412.html>
- Configuring SolidFire on Linux for Element OS
<https://fieldportal.netapp.com/content/468085>

Version History

Version	Date		Document Version History
Version 1.0	September 2018		Initial document creation

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 © 2018 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.