



Technical Report

SQL Server on NetApp HCI

Bobby Oommen, NetApp
January 2019 | TR-4638

Abstract

This document introduces the NetApp® HCI solution to infrastructure administrators and provides important design paradigms to consider when using NetApp HCI solution for SQL Server databases. The document talks about use cases that are ideally suited for NetApp HCI and discusses architecture considerations for applications running in the context of NetApp HCI. By following the guidelines in this document, you can learn how to effectively design, implement, and run SQL Server databases on NetApp HCI.

TABLE OF CONTENTS

1	Introduction	3
1.1	Performance Guarantee	3
1.2	Enterprise Scale	3
1.3	Streamline Operation	3
1.4	Configuration	4
2	Application Use Cases	4
2.1	Consolidation	4
2.2	Development and Testing	5
3	Virtual Machine Configuration	5
3.1	Memory Reservation	5
3.2	Storage Option	5
4	Storage Configuration	5
4.1	Creating an Account	5
4.2	Creating a Volume	6
4.3	Creating Initiators	7
4.4	Creating Volume Access Groups	7
4.5	Adding Volumes to Access Groups	8
5	Operating System Configuration	9
5.1	Enabling Jumbo Frames	9
5.2	Enabling Microsoft iSCSI Service	9
5.3	Enabling Multipath I/O	10
5.4	Configuring iSCSI	12
5.5	Autoconfiguring iSCSI	13
5.6	Initializing Volumes	13
6	SQL Server Clustering	15
7	Sample Sizing	16
8	Summary	16
	Where to Find Additional Information	17
	Version History	17

1 Introduction

NetApp HCI is an enterprise-scale, hybrid converged infrastructure solution that is ideally suited for customers who want to break free from the limitations of the first generation of NetApp HCI.

With NetApp HCI, you can run multiple applications with guaranteed performance to confidently deploy resources across your entire data center. With this architecture, you can deploy your infrastructure by simplifying management and independently scaling both compute and storage resources. NetApp HCI is Data Fabric ready out of the box so you can access all your data across any cloud, whether it is public, private, or hybrid. By moving to NetApp HCI, IT organizations can transform their data center by driving operational efficiency and reducing costs.

The NetApp Data Fabric is a software-defined approach to data management that allows you to connect disparate data management and storage resources. You can also streamline data management between on-premises resources and cloud storage for enhanced data portability, visibility, and protection.

1.1 Performance Guarantee

One of the biggest challenges in any data center is delivering predictable performance, especially when multiple applications share infrastructure. There is always the possibility that one application might interfere with another application, creating a performance disruption that forces administrators to spend valuable time troubleshooting. Mainstream applications, such as virtual desktop infrastructure and database applications, have different I/O patterns that can affect each other when deployed in a shared infrastructure. The robust quality of service (QoS) capabilities that come with storage nodes provide you with fine-grained control of every application, which eliminates noisy neighbors, meets unique performance needs, and satisfies performance SLAs. The storage architecture that is a part of the NetApp HCI solution also eliminates performance variation in the context of data localization, because the data is distributed across all the nodes in the NetApp HCI cluster.

1.2 Enterprise Scale

NetApp HCI scales compute and storage resources independently, avoids costly and inefficient overprovisioning, and simplifies capacity and performance planning. Running on the innovative NetApp Element software, which runs on the storage nodes, and delivered on an architecture designed by NetApp, NetApp HCI is a true enterprise-scale, hybrid converged infrastructure solution. NetApp HCI comes in 2RU 4-node building blocks (chassis) with compute and storage nodes available in small, medium, and large sizes.

1.3 Streamline Operation

The holy grail of IT is to automate all routine tasks, eliminating the risk of user error associated with manual operations, while freeing up resources to focus on higher value assignments that drive business. NetApp HCI allows IT departments to become more agile and responsive by simplifying day-zero deployment and ongoing management from day one forward. The NetApp Deployment Engine eliminates most manual steps needed to deploy infrastructure, while the vCenter plug-in makes management with the VMware environment simple and intuitive. Finally, a robust suite of APIs enables seamless integration into higher-level management, orchestration, backup, and disaster recovery tools.

1.4 Configuration

NetApp HCI is available with multiple configuration options for compute and storage. The nodes are similar to a small blade that sits inside a chassis.

From the configuration information in Table 1, each storage node can deploy from 5.5TB to 80TB of effective capacity. From a compute node perspective, shown in Table 2, 8 to 36 CPU cores and 384GB to 1TB of RAM are available.

Table 1) NetApp HCI configuration storage nodes.

	H410S	H610S
RU	2RU, half-width	1RU, half-width
Drive Capacity	480GB/960GB/1.92TB	960GB/1.92TB/3.84TB
Performance per node	50,000 IOPS - 100,000 IOPS	100,000 IOPS
SSD	6x encrypting or non-encrypting	12x encrypting or non-encrypting
Effective block capacity	5.5TB–44TB	20TB–80TB

Table 2) NetApp HCI configuration compute nodes.

	H410C	H610C
RU	2RU, half-width	1RU
Cores for VMs	8-40	32
CPU	2x Intel Xeon Gold 5122, 4 cores, 3.6GHz 2x Intel Xeon Silver 4110, 8 cores, 2.1GHz 2x Intel Xeon Gold 5120, 14 cores, 2.2GHz 2x Intel Xeon Gold 6138, 20 cores, 2.0GHz	2x Intel Xeon Gold 6130, 16 cores, 2.1GHz, 2x NVIDIA Tesla M10 GPU cards
Memory	384GB – 1TB	512GB
Base networking	4x 25/10GbE SFP28/SFP+ 2x 1GbE RJ-45	2x 25/10GbE (SFP28) 2x 1GbE RJ-45

As an example, a minimum size starting solution would be a configuration with two small compute and four small storage nodes. As requirements change, you can add more compute or storage nodes of any size to the chassis independently of each other. This flexibility of adding only compute or only storage nodes enables unique scalability options for building an efficient and agile cloud in your data center for various use cases.

2 Application Use Cases

NetApp HCI can support a wide range of database application use cases. This section shows how to identify which application use cases are a good fit for the NetApp HCI solution.

2.1 Consolidation

The volume QoS controls for storage nodes provide individual databases with the I/O throughput that they need without being affected by other databases that run in parallel on the same NetApp HCI system. With

QoS and data reduction efficiencies, you can achieve higher database density within the shared storage infrastructure by having several database instances.

Moreover, database or VMware administrators have full control of each storage volume on which the database resides. They can perform all maintenance operations, including setting the QoS for each database copy, through the vCenter SolidFire plug-in. Administrators also can use REST APIs to achieve full automation and make storage management simpler and easier. For architectures that contain multiple user databases with varying resource needs, multiple storage volumes can be used with differing QoS controls or policies. Individual databases can be assigned to specific volumes, using default configuration methods or SQL Server filegroups.

2.2 Development and Testing

NetApp Element Snapshot™ copies provide a point-in-time view of the contents of an active file system or storage volume. You can use Snapshot 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 QoS control on the storage nodes so that database clones can coexist with the production copies without any performance effects on the upstream applications.

The CopyVolume feature on the storage nodes allows you to refresh an existing cloned copy of a database without performing any file system remount operations. In this use case, you can frequently refresh a copy of the database by only taking changes from the production copy.

3 Virtual Machine Configuration

3.1 Memory Reservation

Memory requirements are crucial for the SQL Server application. To avoid memory over commitment, you must reserve the memory allocated to the SQL Server virtual machine to avoid performance issues.

3.2 Storage Option

A SQL Server database can be deployed on a VMware virtual machine with several options for storage configuration such as VMFS, RDM, or direct iSCSI session to the underlying storage. For this paper, a direct iSCSI option was used to provision storage for the SQL Server database.

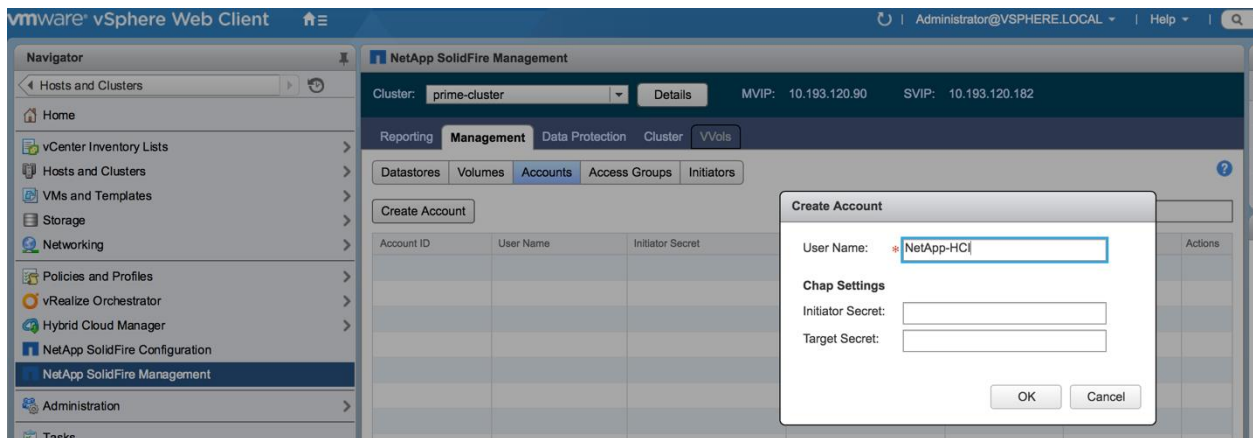
4 Storage Configuration

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

4.1 Creating an Account

To create an account, 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 → Account → Create Account. The Create Account window appears.



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

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

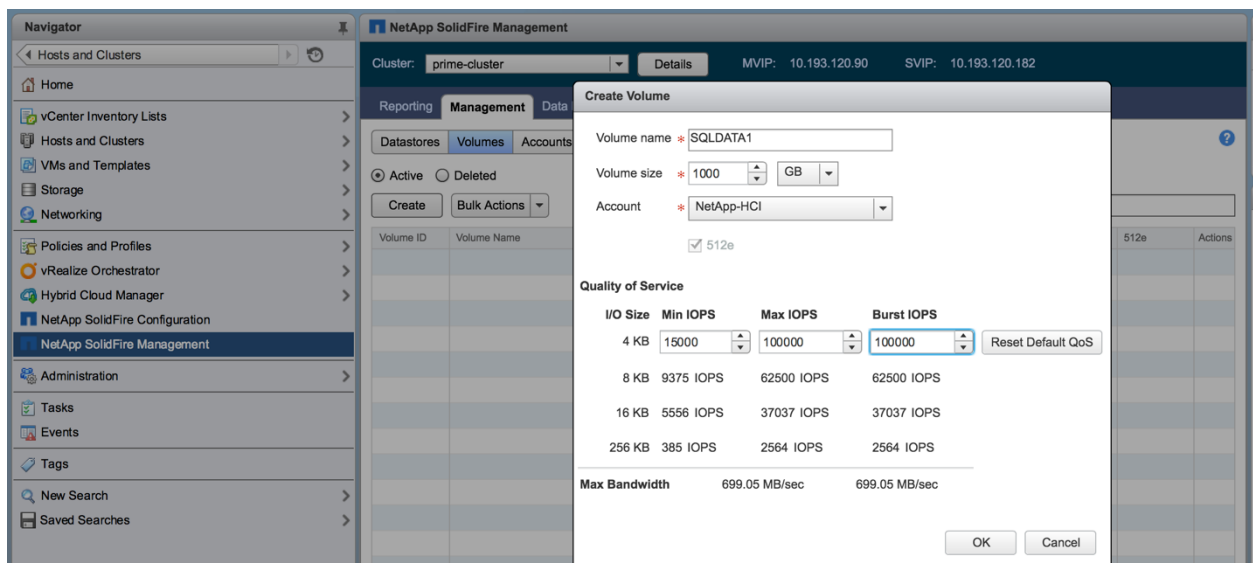
6. Click Create Account.

Note: If an account with the same name exists, you get an error message.

4.2 Creating a Volume

To create a volume, 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 → Volumes → Create. The Create Volume window appears.



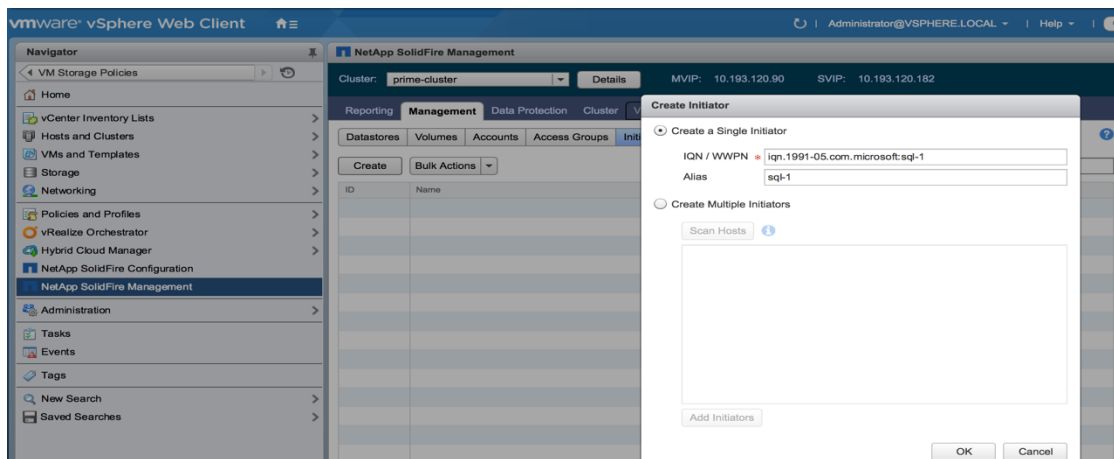
4. Enter the volume name (1 to 64 characters in length). For example, enter the name `SQLDATA1`.

5. Enter the size of the volume.
6. Click the Account drop-down list and select the account that requires access to the volume. In this case, select NetApp-HCI.
7. Enter the QoS values. For this setup, the following values were chosen: Min 15K, Max 100K, and Burst 100K.
8. Click Create Volume.
9. Repeat steps 1 through 8 for all the remaining volumes.

4.3 Creating Initiators

To create a volume, 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 → Initiators → Create. The Create initiator window appears.



4. From the windows host, get the initiator name and type it in the IQN/WWPN field. Optionally, you can provide an alias as well.
5. Click OK.

4.4 Creating Volume Access Groups

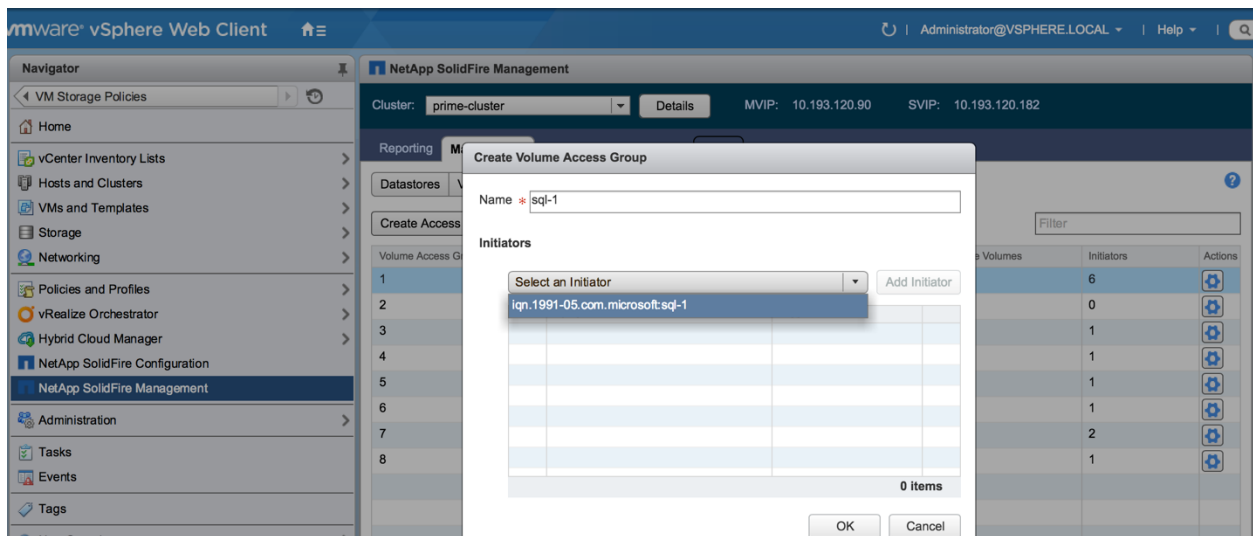
Volume access groups limit connectivity from designated host servers based on a unique identifier, whereas CHAP authentication uses secret keys for unidirectional or bidirectional authentication. In this document, initiator iSCSI qualified names (IQNs) are used to access the volumes.

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 appears.



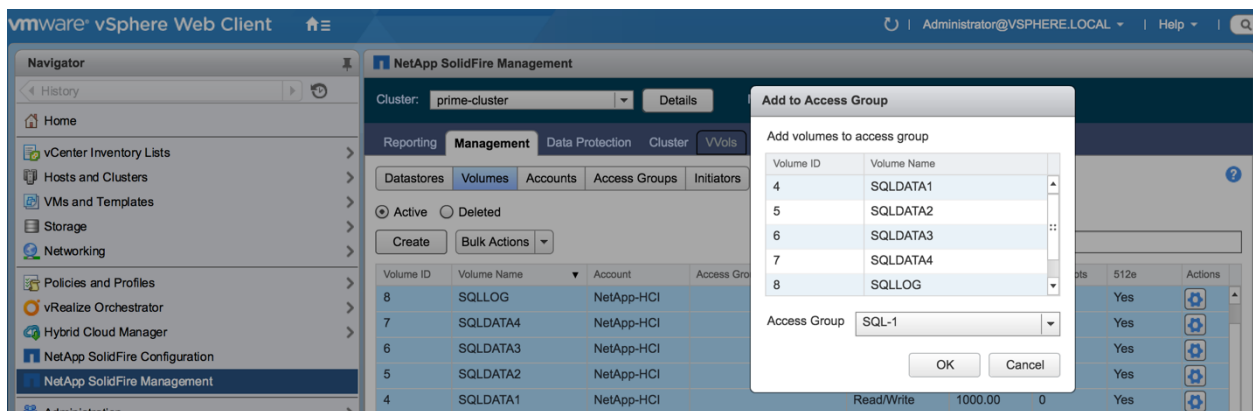
4. Type in the Access group name, `sql-1` in this case.
5. From the drop-down list, select the initiator that was created in step 4.
6. Click OK.

Note: For a SQL Server cluster setup, multiple initiators must be created depending upon the number of nodes in the SQL cluster.

4.5 Adding Volumes to Access Groups

Add volumes to the access group by selecting Management → Volumes.

1. From the listed volumes, select all the volumes that are part of the SQL database.
2. After the volumes are selected, click the Bulk Actions drop-down list.
3. Select Add to Access Group. The Add to Access group window appears.



4. From the Access Group drop-down list, select the access group SQL-1.
5. Click OK.

The SQL Server database volumes are now listed as part of the selected volume access group and are ready to be mapped to the host operating system. At the Windows host, a dynamic disk or striped volume was created to combine multiple data volumes from the storage nodes.

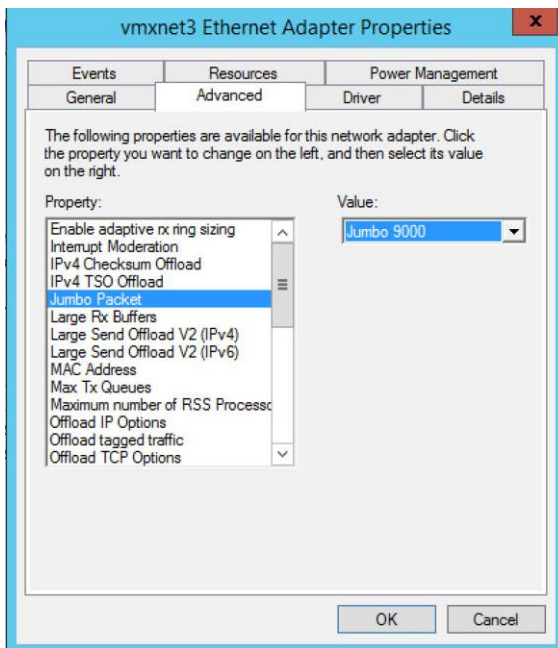
5 Operating System Configuration

The guidelines in this document apply to the Windows 2016 R2 operating system and SQL Server 2016. Alternate Windows versions can be used assuming that they have full compatibility with the SQL Server database software.

5.1 Enabling Jumbo Frames

You must configure the network interface card that is used to access the storage network to support jumbo frames. On the server where SQL Server is running, complete the following steps:

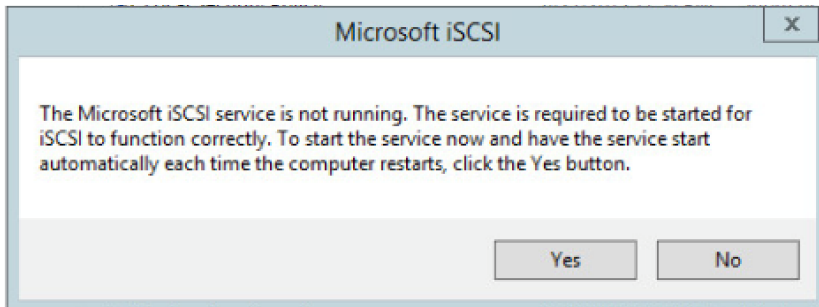
1. Select Control Panel → Hardware → Device Manager. The device manager window opens.
2. Expand Network Adapters.
3. Right-click the Ethernet adapter used to connect to the storage network and select Properties from the context menu. An adapter properties dialog box appears.
4. Select Advanced Tab.
5. In the Property list, select Jumbo Packet and change the value from the drop-down list to Jumbo 9000.



5.2 Enabling Microsoft iSCSI Service

If the Microsoft iSCSI service is not already running, complete the following steps:

1. Select Administration Services → iSCSI Initiator → Right-click → Start.
The Microsoft iSCSI dialog box opens.

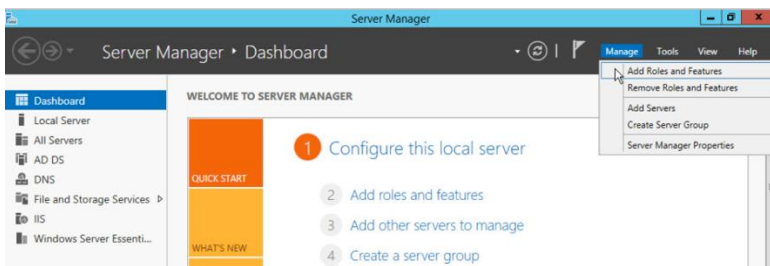


2. Click Yes.

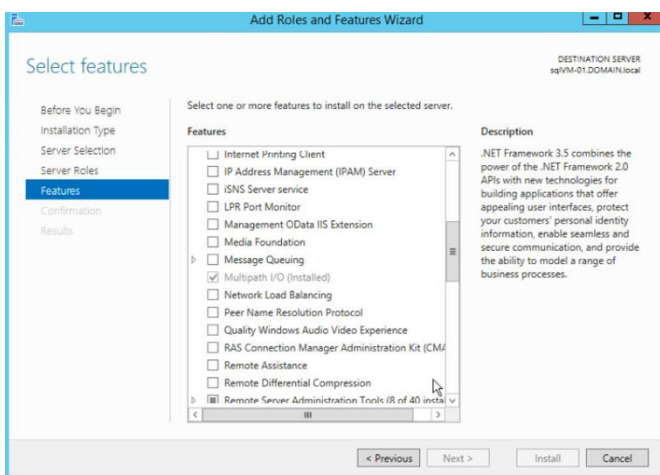
5.3 Enabling Multipath I/O

To enable MPIO and configure the Microsoft device-specific module to recognize the storage network, complete the following steps:

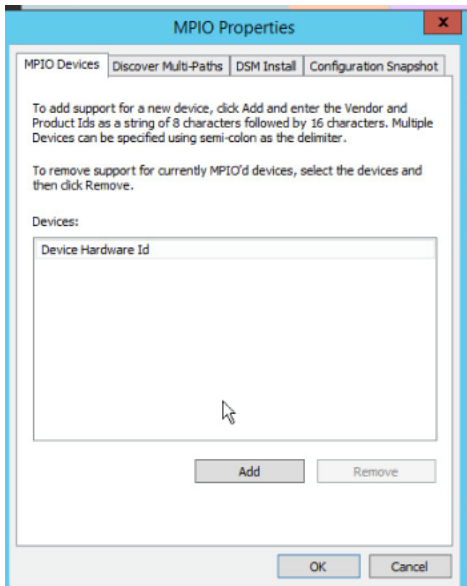
1. Open Server Manager.
2. On the Server Manager dashboard, select Manage → Add Roles and Features.



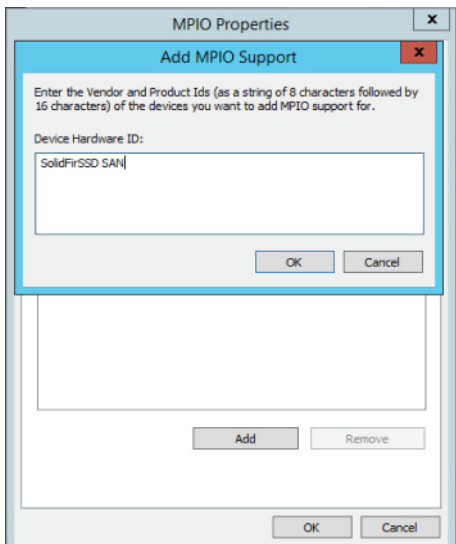
3. In the Add Roles and Features wizard, complete the following steps:
 - a. Before you begin, click Next.
 - b. For Installation Type, select either Role Based or Feature Based and click Next.
 - c. For Server Selection, select your server and click Next.
 - d. For Server Roles, click Next.
 - e. For Features, select MPIO and click Next.



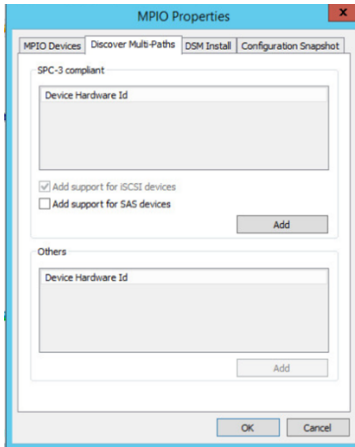
4. Click Install when prompted.
5. After the installation is complete, open MPIO from Administrative Services.



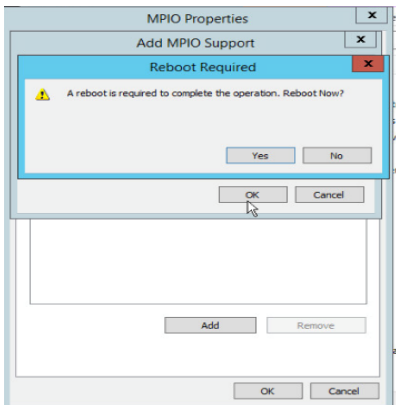
6. Click the MPIO Devices tab and then click Add to add a hardware device ID.
7. Enter SolidFirSSD SAN as the device hardware ID.



8. Click OK.
9. If you are prompted to restart, click No.
10. Select the Discover Multi-Paths tab and then select the checkbox to add support for iSCSI devices.



11. Select the checkbox for Add Support for iSCSI Devices.



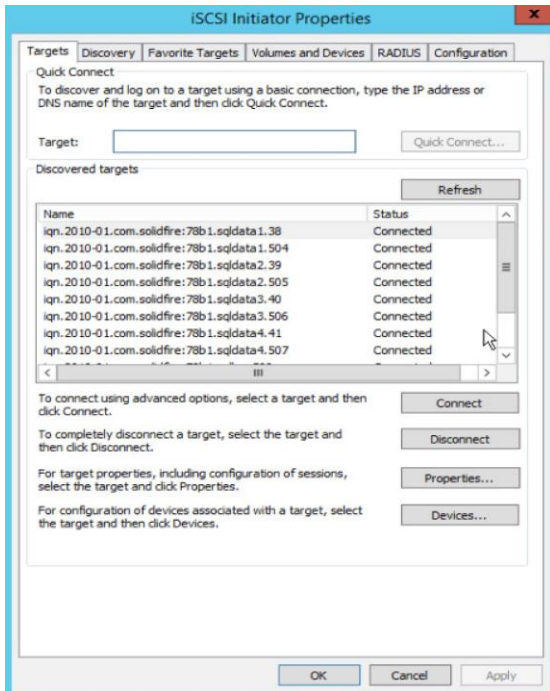
12. When prompted, restart the system.

5.4 Configuring iSCSI

After you enable MPIO, you should configure multiple iSCSI sessions for each storage volume. To do this from the iSCSI initiator utility, complete the following steps:

1. Open the iSCSI initiator utility. The iSCSI Initiator Properties dialog box appears.

Note: The iSCSI initiator performs discovery for each storage volume and returns the IQN strings for the volumes added to the volume access group in the storage NetApp HCI node.



2. For each target IQN string, complete the following steps:
 - a. Select a device and click Properties. The Properties dialog box opens.
 - b. Select the Add This Connection to the List of Favorite Targets checkbox and the Enable Multi-Path checkbox.
 - c. Click OK.

5.5 Autoconfiguring iSCSI

To make sure that your iSCSI volumes are available upon reboot before the SQL Server service starts, you must configure iSCSI for restart access. To do so, complete the following steps:

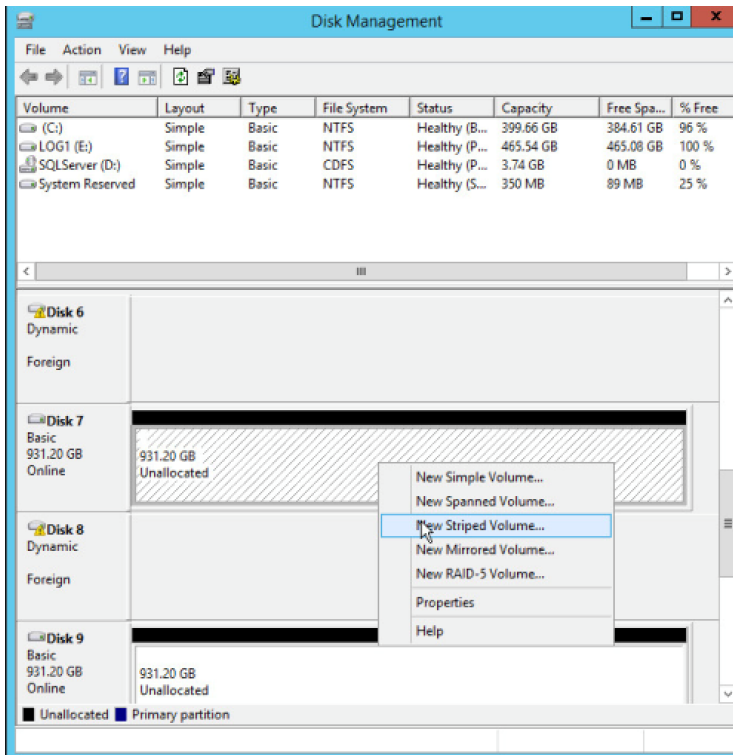
1. Launch the iSCSI Initiator utility. The iSCSI Initiator Properties dialog box appears.
2. Click the Volumes and Devices tab.
3. Click Auto Configure.
4. Click OK.

5.6 Initializing Volumes

The connected iSCSI volumes are presented as standard Windows disks that can be mounted either singly or as a dynamic striped volume in very high-throughput implementations.

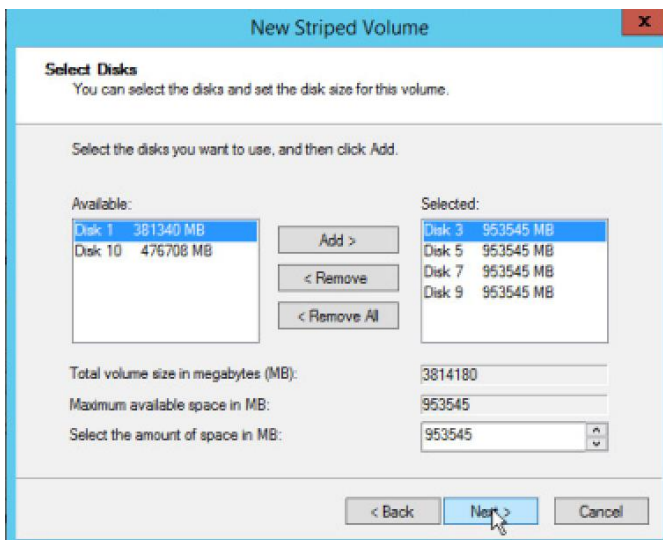
1. Go to Start → Disk Management.
2. Right-click the target disk name and select New Striped Volume.

Note: If only one disk is used, select New Simple Volume.

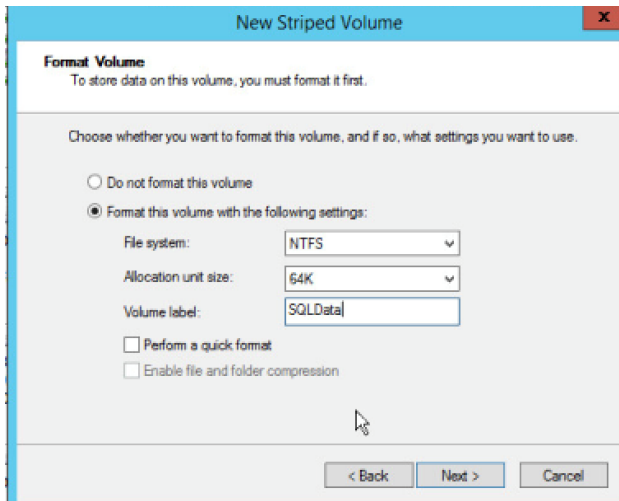


- In the New Striped Volume window, add the required disks and click Next. In this case, four disks are added to create a SQLData striped volume.

Note: For SQL Server clustering, storage pools are recommended to stripe data across the volumes on the storage nodes.



- On the Format Volume screen, set the following parameters and click Next:
 - File system: NTFS
 - Allocation unit size: 64K
 - Volume label: SQLData



Note: The recommended allocation size for the transaction log volume is 4K.

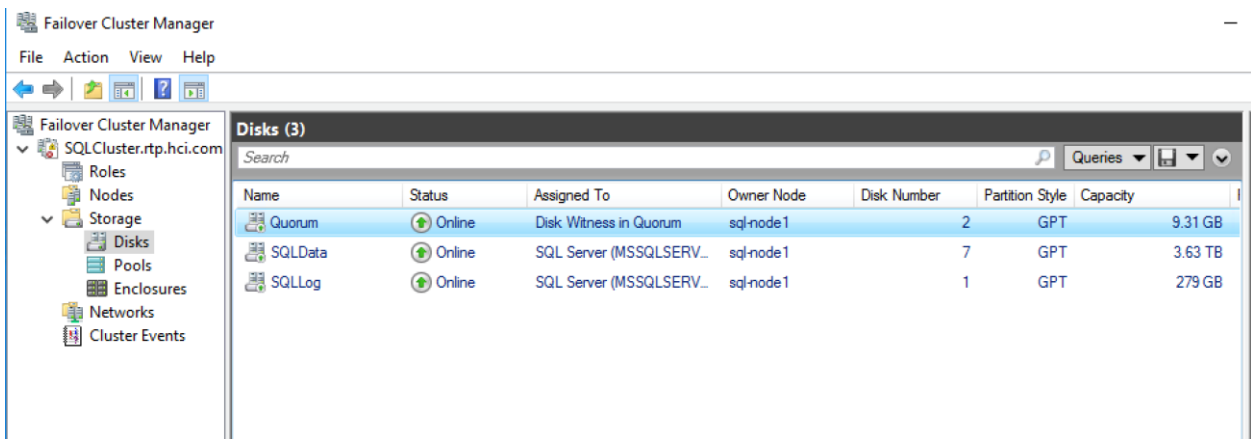
After you configure the iSCSI devices, you can install SQL Server 2016. See the Microsoft SQL Server installation guide for information about installing and configuring SQL Server 2016.

6 SQL Server Clustering

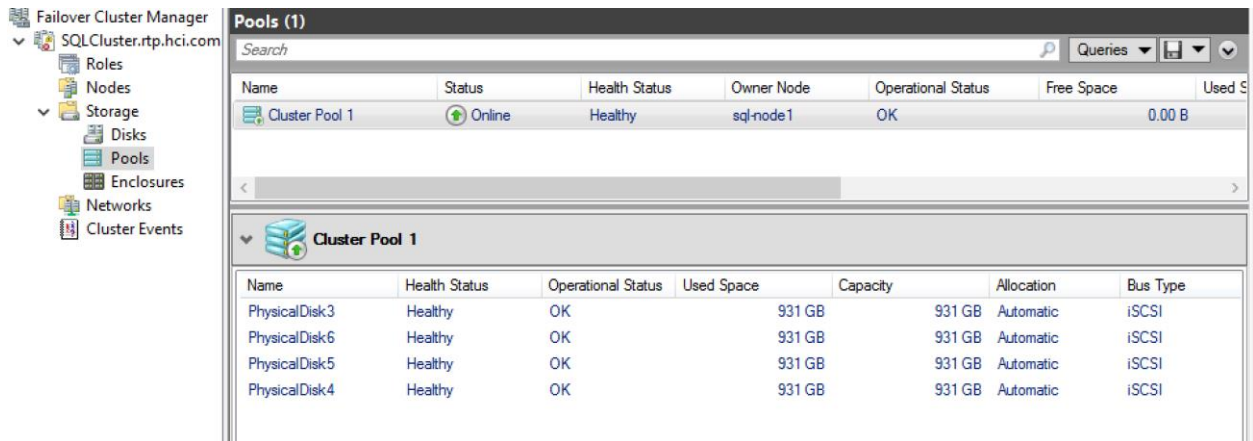
SQL Server clustering is set up by adding one or more SQL Server instances to a Windows failover cluster. A Windows failover cluster uses shared storage. The SQL Server user and system databases must be located on the shared storage. This arrangement allows the cluster to move the SQL instance to any node in the cluster at any time. There is only one copy of the data, but the network name and SQL Server service for the instance can be made active from any cluster node.

Before you set up SQL Server clustering, verify that the following conditions are met:

- Windows failover clustering is set up.
- Each node has Windows Server 2016 installed and is part of the Windows failover cluster. None of the nodes can be a domain controller.
- All the nodes have access to the shared NetApp HCI storage cluster.
- All the iSCSI initiator names from the SQL Server nodes have been added to the NetApp HCI volume access groups.



For this setup, a quorum disk was set up along with a data and log disk (see the previous screen shot). These resources were provisioned from the underlying NetApp HCI storage nodes. In order to increase the throughput for the SQL Server cluster database, multiple volumes were created on the storage nodes and were assigned to a storage pool at the Windows cluster, as depicted in the following screenshot.



For the transaction log, one volume from the storage node was used to create a volume at the Windows host and was then added to the failover cluster.

7 Sample Sizing

The recommended minimum configuration for SQL Server has four compute nodes and four storage nodes.

Table 3) Minimum configuration.

Minimum Configuration	Small	Medium	Large
SQL virtual machines	8	12	8
vCPU	8	8	8
Memory	16GB	32GB	32GB
Database size	Up to 1TB each	Up to 2TB each	Up to 4TB
Maximum IOPS for all databases	200,000	200,000	400,000

8 Summary

Data centers are moving away from dedicated platforms and are trying to avoid overprovisioning to increase efficiency and reduce cost. NetApp HCI provides a solution with robust QoS capabilities, allowing the granular control of every application, eliminating noisy neighbors, and satisfying performance SLAs. The underlying storage nodes of NetApp HCI use all-flash media coupled with capacity thin provisioning and inline data-reduction features. These features yield significant efficiency and agility when deploying applications and help the business consolidate its workloads with confidence. For additional information, contact your NetApp service personnel.

Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- Configuring SolidFire on Windows for Element
<http://www.netapp.com/us/products/storage-systems/all-flash-array/solidfire-web-scale.aspx>
- NetApp Product Documentation
<http://docs.netapp.com>
- NetApp HCI Documentation Resources
<https://www.netapp.com/us/documentation/hci.aspx>

Version History

Version	Date	Document Version History
Version 1.0	October 2017	Initial document creation
Version 1.0	January 2019	Updated configuration options and other minor edits

Copyright Information

Copyright © 2019 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.