

iSER host attachment from Microsoft Windows servers to IBM SVC/ IBM Storwize family controllers

Overview

Challenge

Block storage access over Ethernet using traditional iSCSI protocol from standalone and clustered Windows hosts is not scalable for high-speed Ethernet networks.

Solution

IBM SVC and Storwize controllers have come up with support for high speed Ethernet ports accessible using an iSCSI Extensions over RDMA (iSER) protocol for high performance and low latency access to storage that is accessible from stand-alone Microsoft Windows hosts and clustered Windows hosts in a MSCS cluster.

iSER configuration guide for initiator and target


The purpose of this paper is to provide detailed, step-by-step instructions on how to discover and establish connections from stand-alone and clustered Microsoft® Windows® hosts to storage from IBM® System Storage® SAN Volume Controller (SVC) and IBM Storwize family systems using the iSCSI Extensions over RDMA (iSER) protocol over high speed Ethernet links.

Using the steps in this paper, Cluster Shared Volumes (CSVs) in a Microsoft Cluster Service (MSCS) can be accessed over iSER connections from SVC or IBM Storwize® storage system.

This paper is intended for users who are familiar with SVC, Internet Small Computer System Interface (iSCSI), iSER interconnects, and MSCS clustering but need a practical example to guide them through the steps of *how to* export the iSER-attached disk to a Windows cluster environment and create a CSV volume an existing iSCSI storage area network (SAN) environment to an SVC virtualized environment.

Getting started

IBM SVC and IBM Storwize family systems were used as iSCSI targets until the 8.1.0.0 release using 1 Gb and 10 Gb Ethernet ports. Release 8.1.1.0 added support for 25 Gb RoCE and Internet Wide-area RDMA Protocol (iWARP) ports which could be used as iSCSI targets. This means that only iSCSI hosts could connect to IBM SVC or IBM Storwize family controllers using the 25 Gb ports.



In the 8.2.0 release, IBM introduced support for iSER target on SVC and Storwize platforms using the 25 Gb RoCE and iWARP ports. This feature enables hosts to connect to SVC and Storwize using the iSER protocol enabling high performance and lower latency connectivity from hosts to SVC/Storwize.

In the 8.2.1 release, IBM introduced support for iSER host attach from stand-alone and clustered Windows hosts.

Introduction

The iSCSI protocol based on the TCP/IP stack has been the primary protocol used for block data transfers over Ethernet fabrics. Data transfers using traditional iSCSI stacks use processor cycles on host for data transfers and involve multiple copies of data in the I/O path and high latency. iSER is based on a data mover architecture, which can work on various remote direct memory access (RDMA) technologies such as RoCE, iWARP and InfiniBand. iSER enables direct remote movement of data across application buffers on different hosts and can help to achieve low latency data transfers with a high throughput.

Advantages of iSER

iSER has multiple advantages:

- It frees up the host processor off of any data copy work, increasing the overall system performance.
- It eliminates the need to create multiple copies of data across OS and network boundaries.
- It is vendor and technology independent. Data transfers happen through an API layer called `ib_verbs`, which is available both in user space and kernel space and it is independent of the underlying Ethernet technology. Any Ethernet protocol that supports the `ib_verbs` API layer and conforms to the RDMA semantics can be used.
- iSER offers latency in the range of 70 microseconds for 4KB block transfers to 300 microseconds for 256 KB transfer sizes
- 25 Gbps ports offer 2.5 times the bandwidth of 10 Gbps ports with reduced and simplified cabling.
- 25 Gbps ports can be connected to a 10 Gbps fabric using SFP+ connectors. The same ports can serve data over iSCSI or iSER protocols from different hosts. Hence, it is possible for legacy 10 Gb iSCSI hosts to establish an iSCSI connection to the SVC target through the 25 Gbps ports connected to 10 Gb switches.

Architecture

Software

- IBM SVC software product version 8.2.1.0 or later

Hardware

- SAN Volume Controller model SV1 8.2.1 or later
- Storwize V7000 model 600 or later / IBM FlashSystem® 9100 or later
- Windows 2016 Server with Chelsion T6 adapter

Network

- Cisco/Mellanox 100Gig Switch or other supported 25Gb/ 50Gb/ 100Gb switch
-

Environment setup

This section describes the lab environment used to perform the tests.

The setup used has two Microsoft Windows 2016 servers in a MSCS 2016 Windows cluster that has a 50 GB logical unit number (LUN) allocated from an IBM Storwize storage controller (refer to Figure 1).

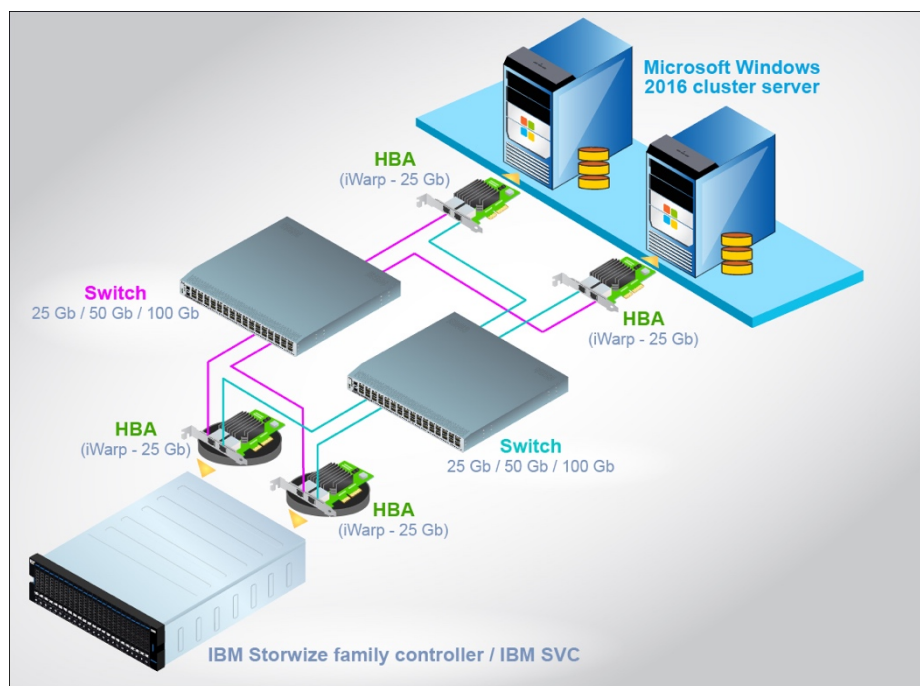


Figure 1: Lab setup with IBM Storwize as a storage controller

The iSCSI initiator configuration is same on both SVC and Storwize systems. In this example, the test team used SVC as the iSCSI initiator. Same steps are applicable for Storwize also.

iSER initiator and iSER target configuration for MSCS support

We will first look at the steps need to configure the IBM SVC and/or IBM Storwize family system as an iSER Target.

Configuring iSER Storwize target controller

Administrative model for configuring iSER is exactly the same as iSCSI. Perform the following steps to establish iSER host-attach connectivity:

1. Find out the IQN of the servers in the host MSCS cluster.

For host 1, on the iSCSI Initiator Properties page, click the **Configuration** tab. In the Initiator Name section, find the specified IQN.

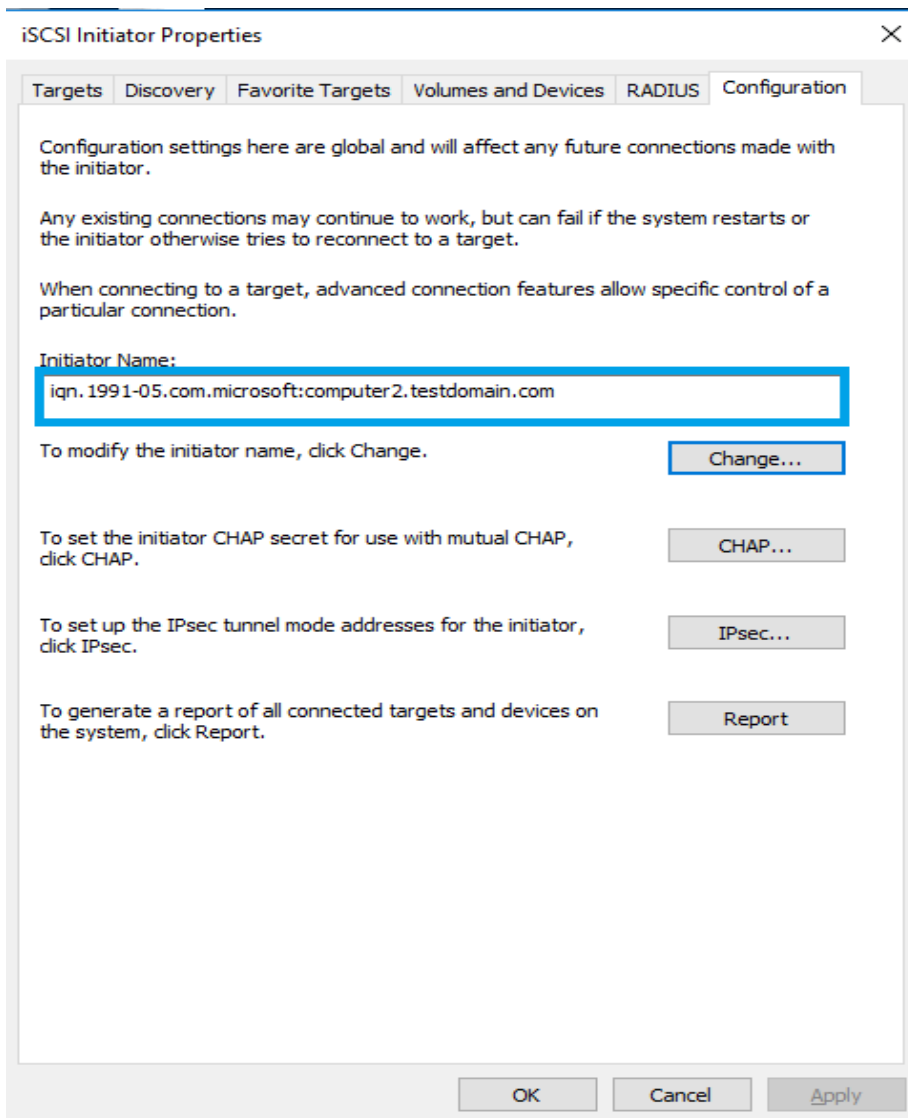


Figure 2: Retrieving the iSCSI / iSER IQN for host 1 from the iSCSI Initiator tab

On Storwize controller, create a host object and add all the SVC initiator's iSCSI qualified name (IQN) in the host object.

```
# svctask mkhost -iscsiname<iqn> -name <host_obj_name>  
  
Example:  
svctask mkhost -iscsiname  
iqn.199105.com.microsoft:computer2.testdomain.com -name  
winCluster
```

- 2. Find out the IQN of the second server in the host MSCS cluster.

On the iSCSI Initiator Properties page of host 2, click the **Configuration** tab and find the specified IQN.

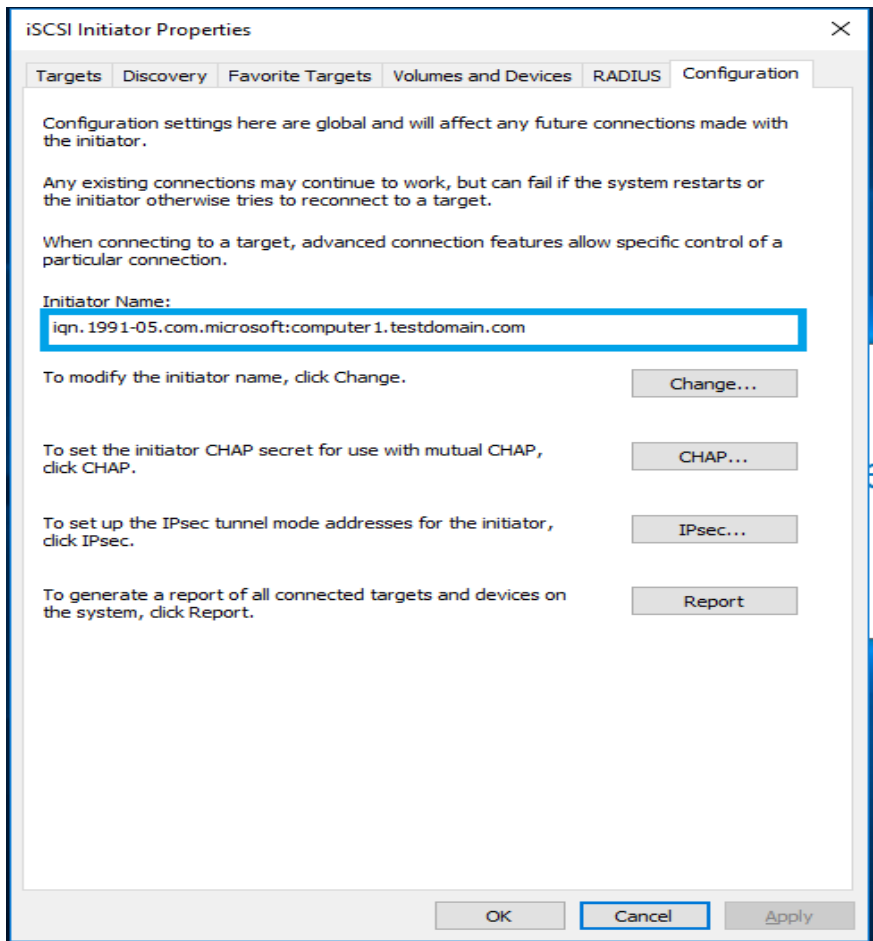


Figure 3: Retrieving the iSCSI / iSER IQN for host 2 from the iSCSI initiator tab

3. Use the `host_obj_name` which is given in `mkhost` command in step 1 and use the same in the following CLI to add other IQNs of the iSCSI initiators cluster.

```
svctask addhostport -iscsiname <iqn> <host_obj_name>
```

Example:

```
svctask addhostport -iscsiname  
iqn.199105.com.microsoft:computer2.testdomain.com  
winCluster
```

4. Verify the host object from the Storwize target GUI.

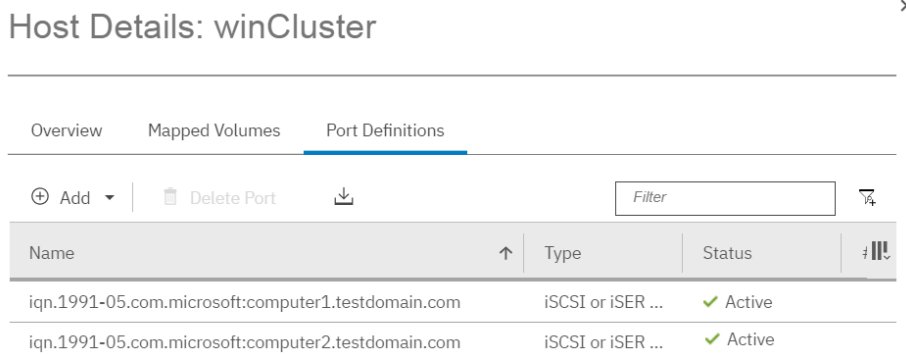


Figure 4: Viewing the iSCSI/ iSER host object and host port details in the GUI

5. To enable iSCSI authentication on the back-end Storwize controller, specify a Challenge Handshake Authentication Protocol (CHAP) secret to the host object created in step 1.

```
chhost -chapsecret<chap secret><host object id / name>
```

Example:

```
svctask chhost -chapsecret hostchap winCluster
```

6. Map the volumes to the MSCS cluster that the user wants to access from the MSCS hosts.

```
svctask addhostport -iscsiname <iqn> <host_obj_name>
```

Example:

```
svctask addhostport -iscsiname  
iqn.199105.com.microsoft:computer2.testdomain.com  
winCluster
```

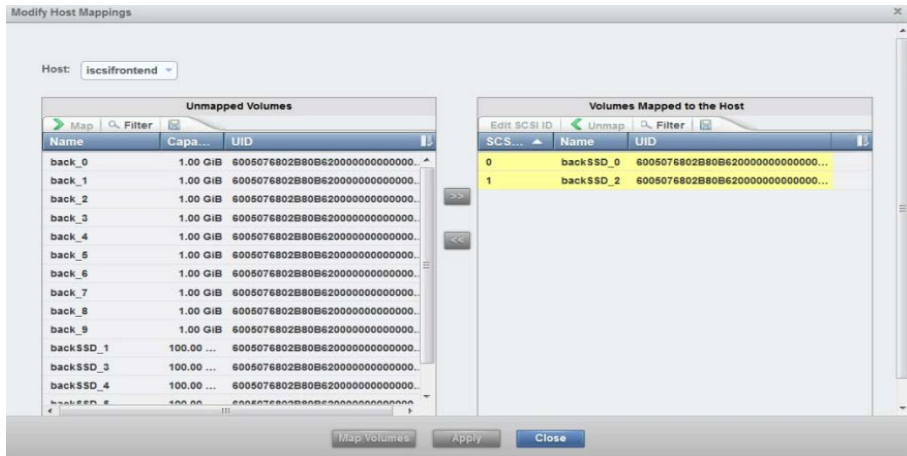


Figure 2: Map VDisk to the newly created host object

7. Configure the target IPs on the Storwize controller using the `-host` flag to allow login from the iSER initiator, if not done already.

```
svctask cfgportip -node <node id> -ip<ipv4 addr> -  
gw<ipv4 gwaddr> -mask <mask> -host yes <port id>  
  
svctask cfgportip -node <node id> -ip_6 <ipv6 addr> -  
prefix_6 <prefix> -gw_6 <ipv6 gwaddr> -host_6 -yes  
<port id>
```

Example:

```
svctask cfgportip -node 1 -ip 172.168.33.100 -gw  
172.168.33.1 -mask 255.255.255.0 -host yes 6  
  
svctask cfgportip -node 1 -ip 172.168.34.100 -gw  
172.168.34.1 -mask 255.255.255.0 -host yes 7  
  
svctask cfgportip -node 2 -ip 172.168.33.200 -gw  
172.168.33.1 -mask 255.255.255.0 -host yes 6  
  
svctask cfgportip -node 2 -ip 172.168.34.200 -gw  
172.168.34.1 -mask 255.255.255.0 -host yes 7
```

If the IPs are already configured, change the host flag to yes using the following command in the CLI:

```
svctask cfgportip -node <node id> -host yes <port id>  
  
svctask cfgportip -node <node id> -host_6 -yes <port id>
```

Example:

```
svctask cfgportip -node 1 -host yes 6  
svctask cfgportip -node 1 -host yes 7  
svctask cfgportip -node 2 -host yes 6  
svctask cfgportip -node 2 -host yes 7
```

Figure 6 shows a sample view of the Storwize system with IPs configured.

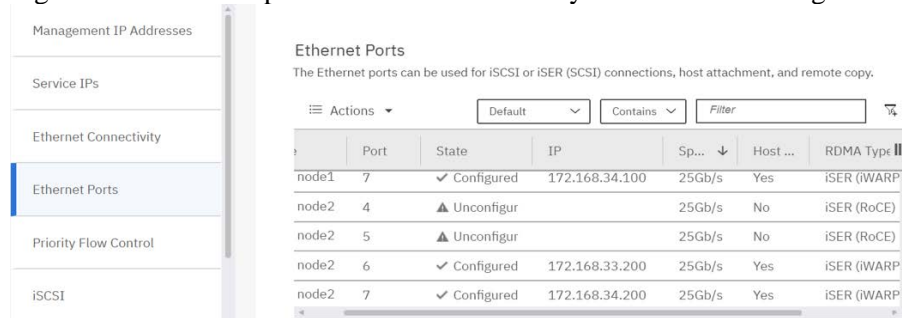


Figure 3: iSCSI IP configured at the IBM Storwize Controller for iSER connectivity

iSER initiator configuration on MSCS cluster servers

Prerequisites:

- Install and configure two Microsoft Windows Server 2016 systems with MSCS host configuration.
- Make sure that all servers are preconfigured with Microsoft Windows Active Directory (AD) services.
- Currently iSER host-attach support to SVC and Storwize is available over 25 Gb iWARP ports on host and target. The standard iSER target port for iSER login over iWARP ports is 860.
- Before configuring iSER at the host side, ensure that all configured target IPs are reachable to and from the iSER Windows host.

Configure the iSCSI/iSER initiator on the Windows host by performing the following steps:

1. Install the updated iSER driver from host adapter vendor website. You can download the driver package for the currently supported iWARP adapters from <https://service.chelsio.com>.
2. After successfully installing the driver, expand the **Computer Management** console and click **Device Manger**. On the right pane, verify the drivers in the Network adapters section, Storage controllers section, and the System devices section.

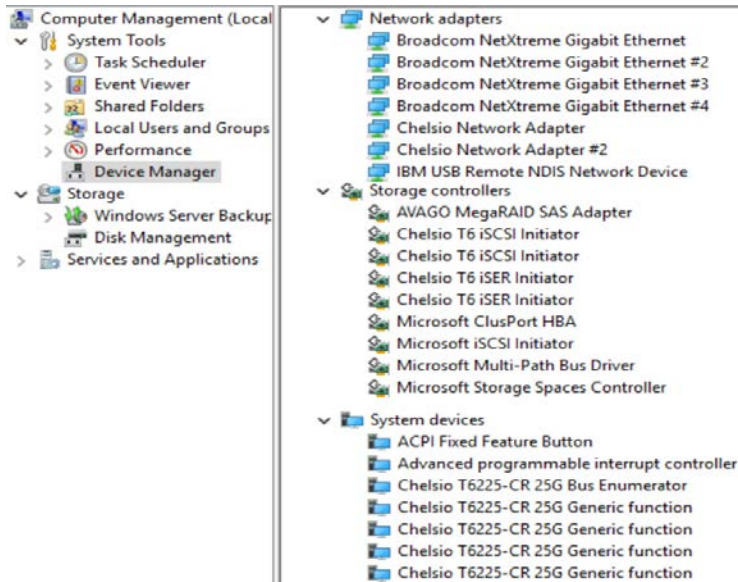


Figure 7: Verify successful installation for the drivers

3. Open the native iSCSI initiator and click the **Discovery** Tab. If no iSCSI / iSER sessions have been established to any target, no entries will be seen in the Target portals section.

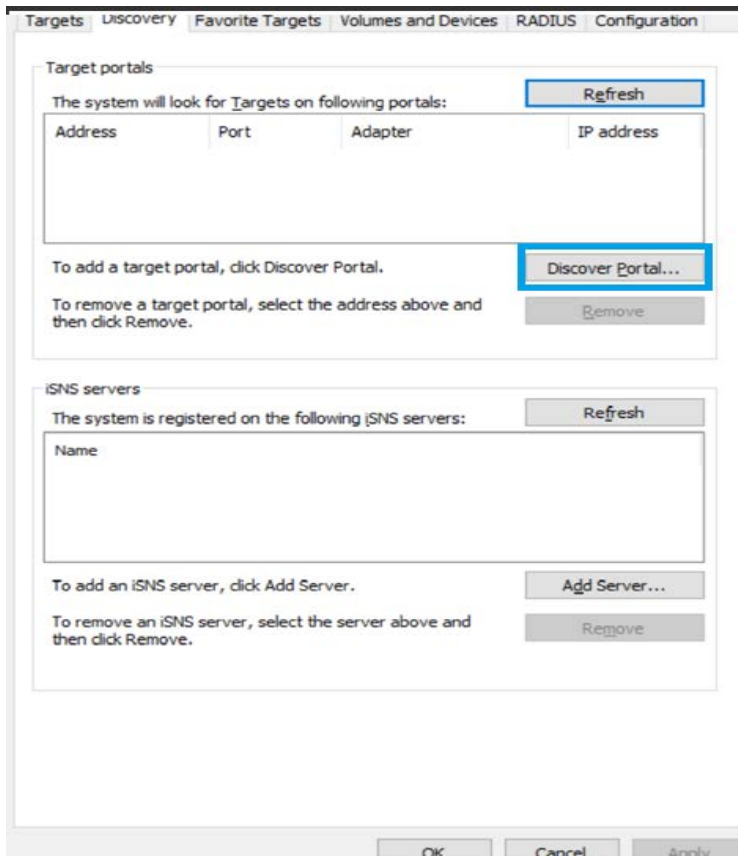


Figure 8: Discovery tab on the iSCSI initiator tool

4. Click **Discover Target**, as shown in Figure 8. In the Discover Target Portal window, click **Advanced**.

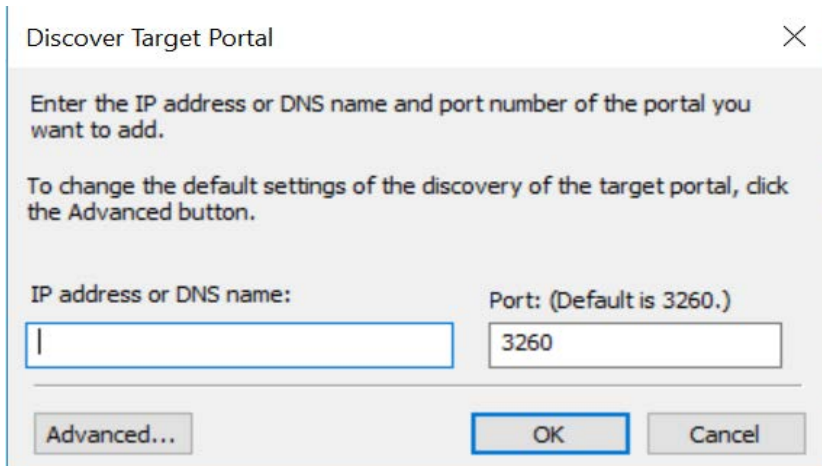


Figure 9: Discovery Target Portal window

5. In the Advanced Setting dialog box, select the corresponding local adapter and initiator IP. Optionally, you can also configure CHAP. Figure 10 shows a screen capture without CHAP configured. Select the local iSER adapter from the drop-down list and its respective IP from Initiator IP drop-down list (according to configured iSER port) and click **OK**.

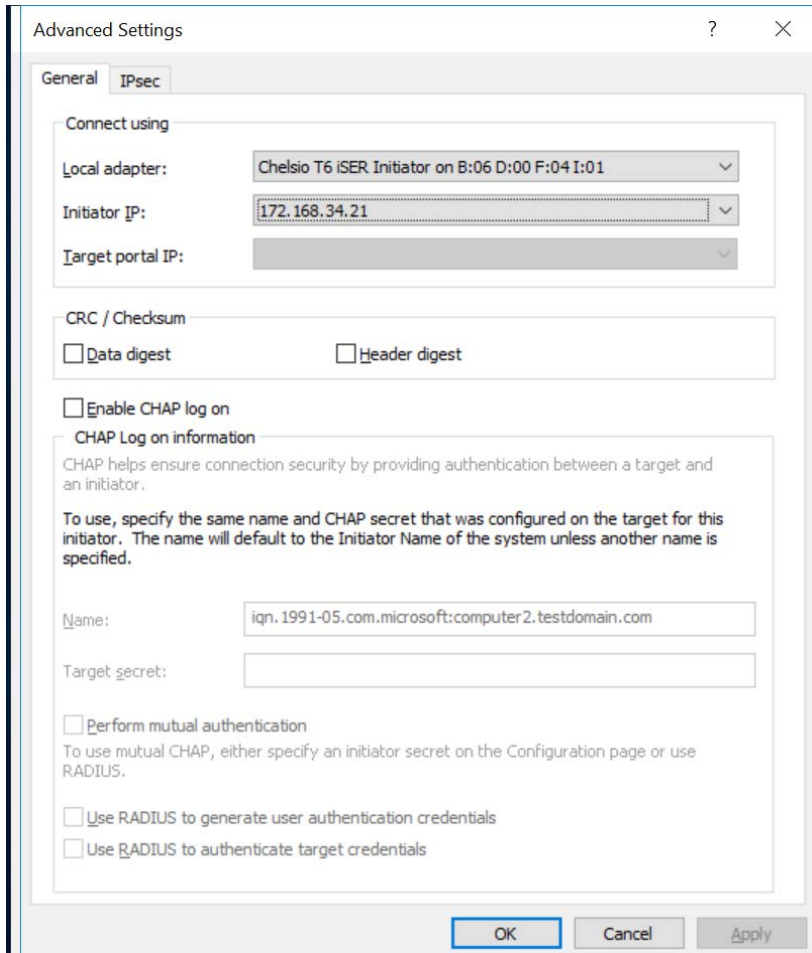


Figure 10: Configure corresponding local adapter and initiator IP

6. You will be redirected to the previous Discover Target Portal dialog box. Assign **860** as port number for iWARP target discovery along with a target IP address. The example in Figure 11 shows the target IP address of the node 2, port 7 configured on Storwize earlier.

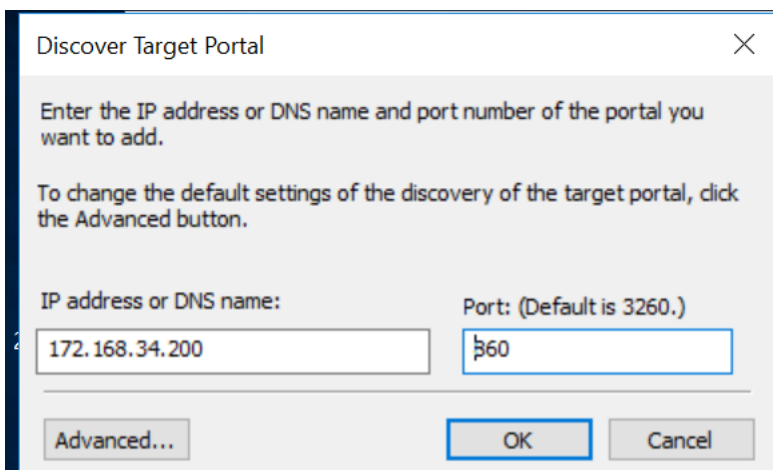


Figure 11: Configured target IP and iSER port

After successful IP discovery, it is listed in the Discovery tab in the Target portals section.

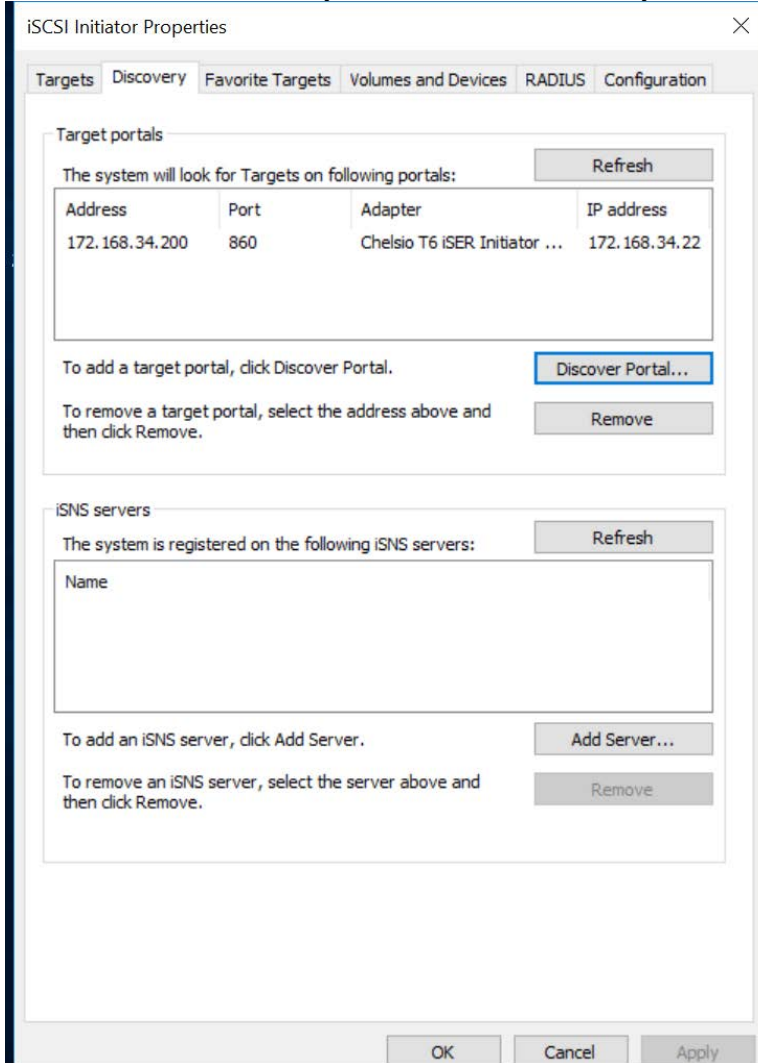


Figure 12: iSER discovered IP listed in the Discovery tab

- Repeat [step 3](#) to [step 6](#) to discover target endpoints using different initiator ports by selecting a different initiator adapter port and IP address in [step 5](#) and a different target portal in [step 6](#).

Note: It is a best practice to establish one to one connectivity from each initiator port to a SVC / Storwize target port.

In the Targets tab, the corresponding discovered target IQN will be shown. Make sure that the status for this is **Inactive**.

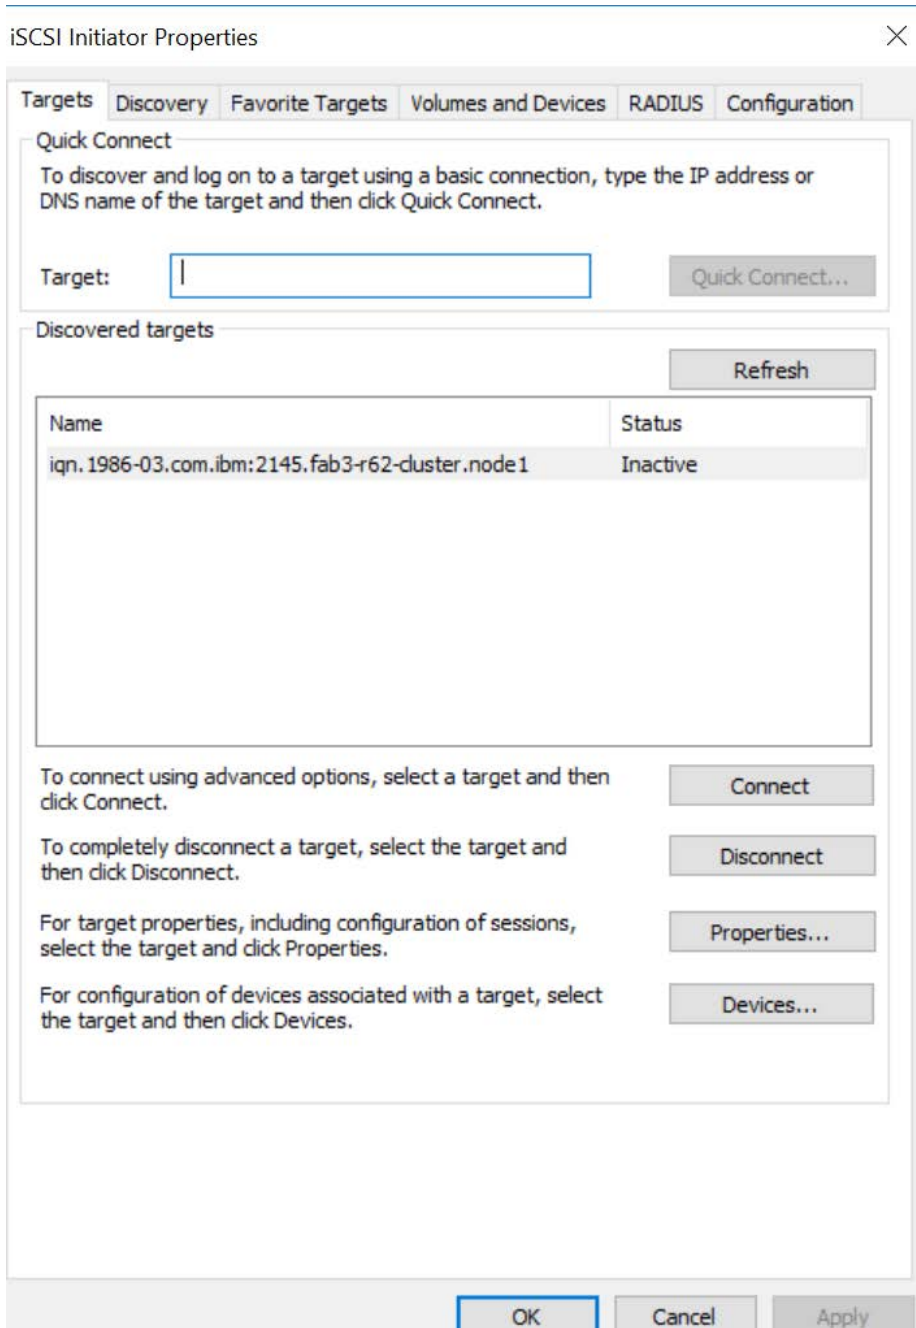


Figure 12: Targets tab

8. Select the name of the discovered target IQN and click **Connect**. In the Connect To Target dialog box, select the **Enable multi-path** check box and click **OK**.

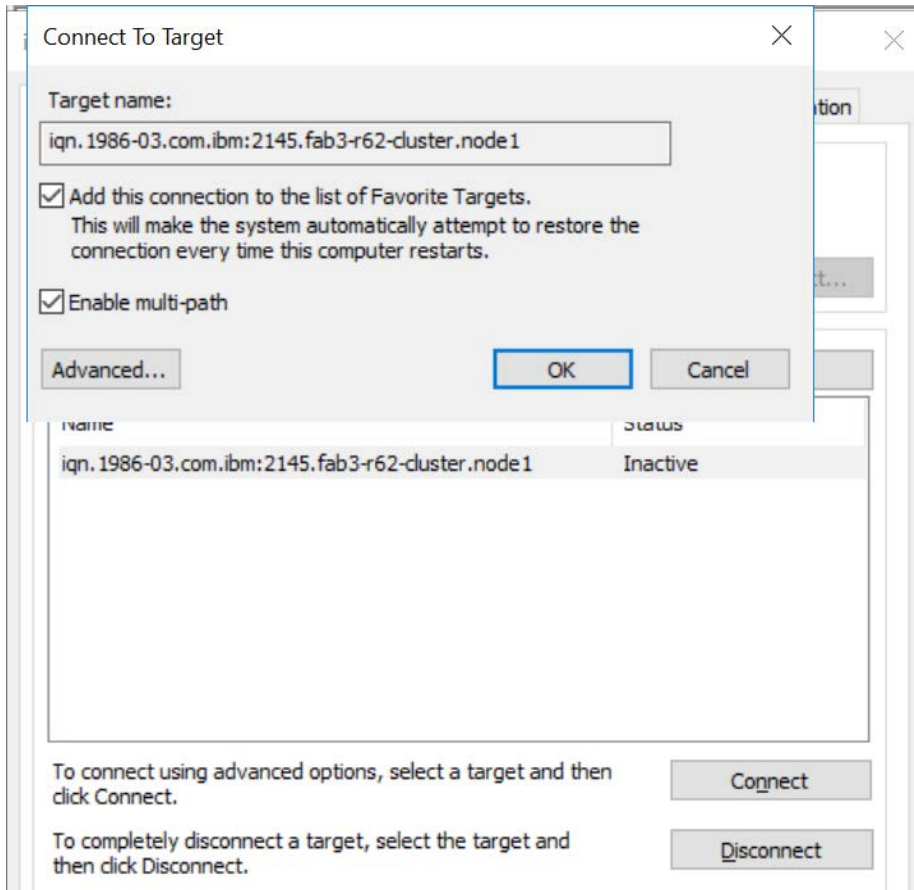


Figure 13: Connect to target

9. Session will successfully be established with the iSER target In the Targets tab, verify that the status of the discovered target IQN is **Connected**.

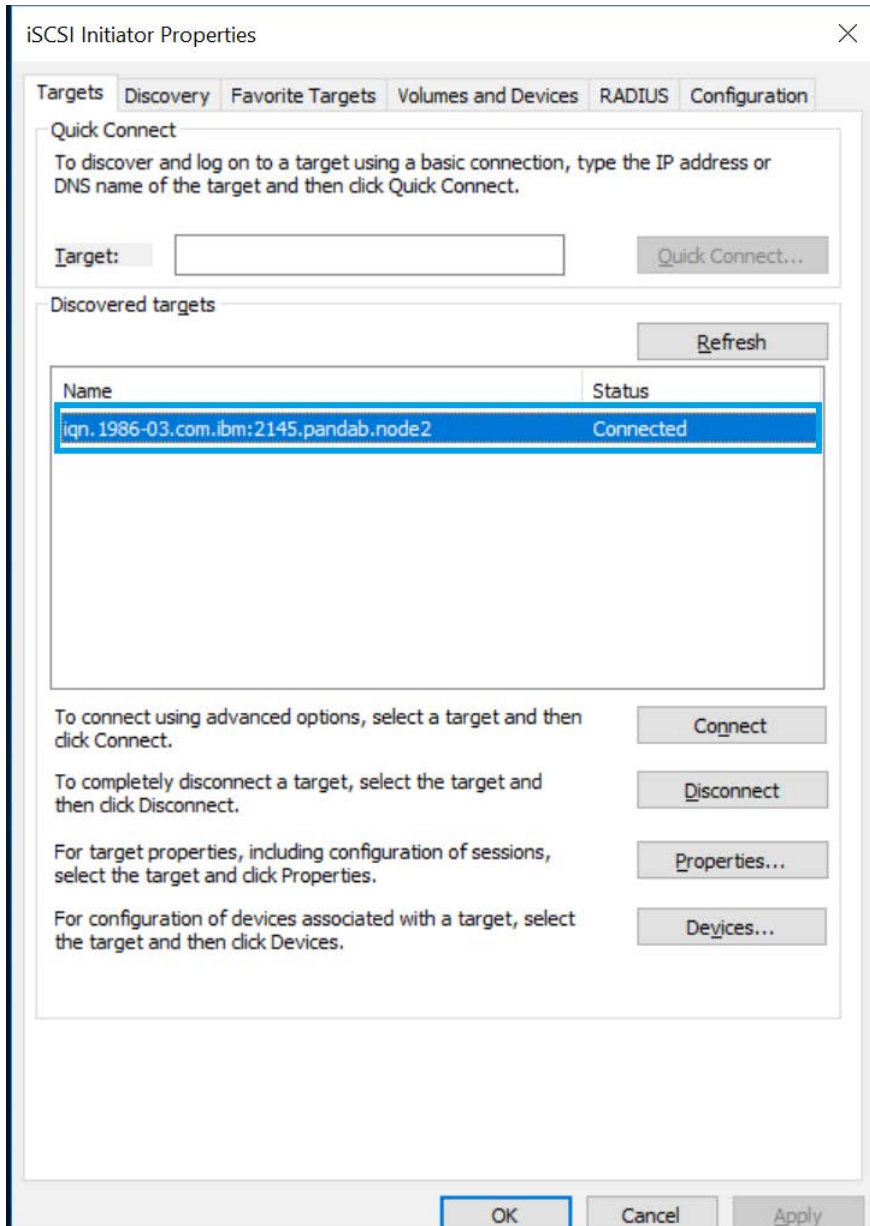


Figure 14: Target connected

10. Repeat same operation from [step 1](#) to [step 9](#) for the second host on the MSCS cluster and make sure that both the hosts are able to successfully establish the session with the SVC target node.
11. Verify the session information details from the target SVC management GUI.
In the SVC cluster GUI click the **Host** tab and select the **WinCluster** host object.

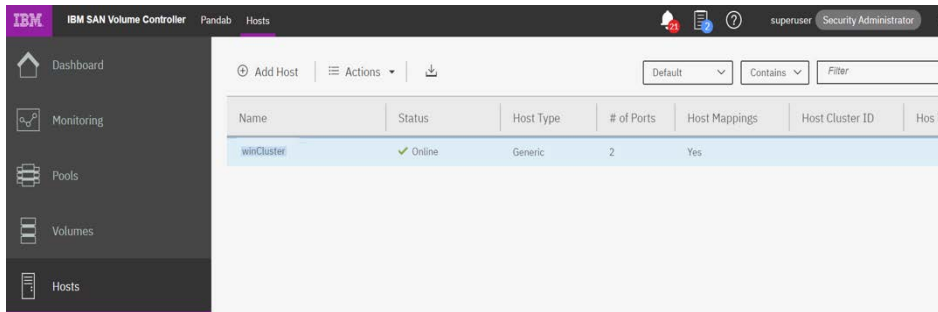


Figure 15: Host object created with status Online and configured host port IQN 2

12. Click the **Hosts** tab. Double-click on a listed host object to view the details (refer to Figure 16).

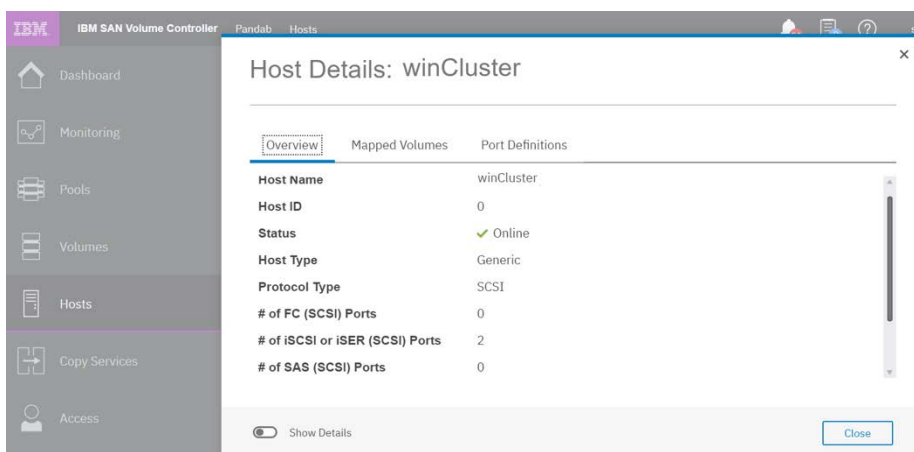


Figure 16: Host object details showing online information

Note: The iSCSI initiator IQN that you configure on the host side should never be the same as the iSER target IQN on the system node.

Configuring Cluster Shared Volume

For MSCS clustering, you need at least two hosts to configure and share the CSV exported from a target. In the previous section, two Windows hosts were used and the iSER initiator settings were completed on both the hosts. Following steps show a sample configuration of Cluster Shared volumes.

1. Configure two VDisks and export to the Windows MSCS cluster by mapping the two VDisks of 50 GB each to the newly created winCluster host object.
2. Go to Disk Management and create a file system on it. The same file system should reflect on the second host, which is part of the MSCS cluster.

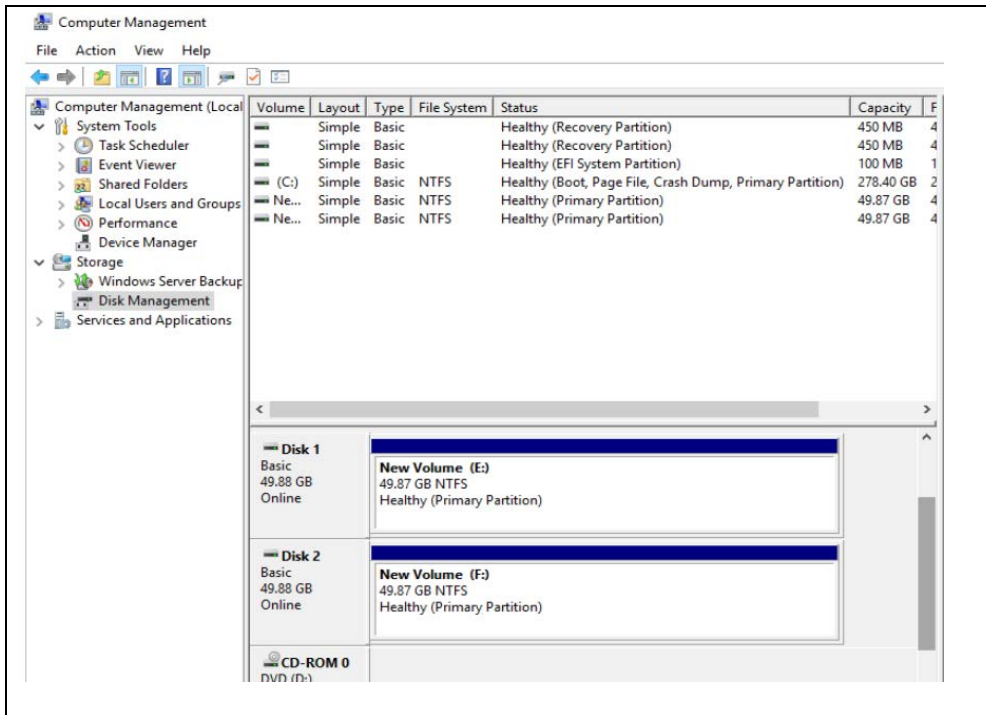


Figure 17: Disk Management

3. Make sure that you have configured the failover cluster. In the Failover Cluster Manager, verify that the Node status is **Up**.

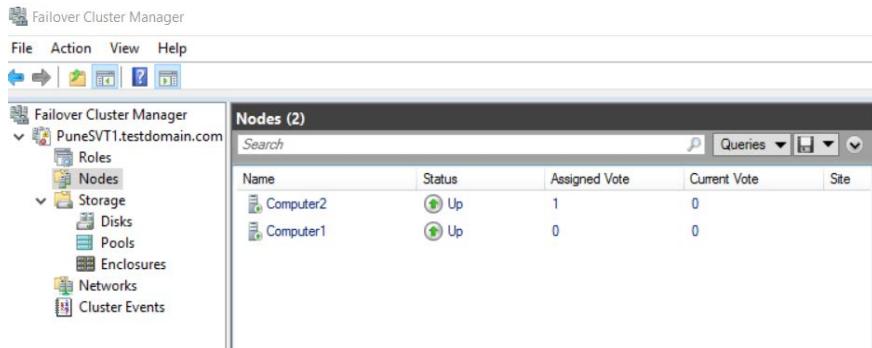


Figure 18: Failover Cluster Manager – Check Node status

4. In Failover Cluster Manager, expand **Storage**, and click **Disks**. Then, add those two disks which was been exported from SVC node to the cluster as shown in Figure 19.

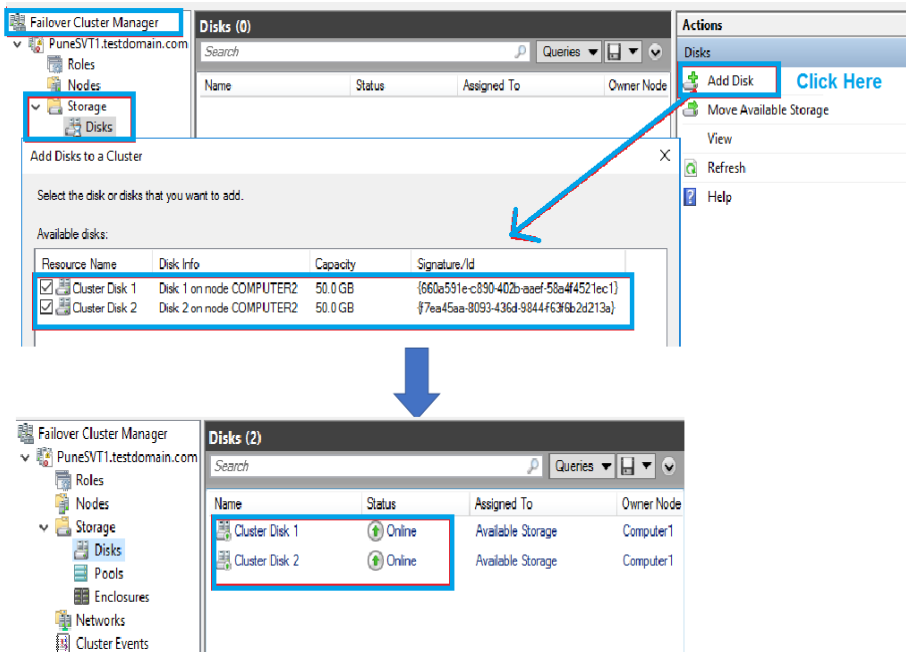


Figure 19: Failover Cluster Manager – Add disks

5. After adding the exported VDisks to the cluster, they are shown as available storage. Now, select the added disk and configure it to be a Cluster Shared Volume.

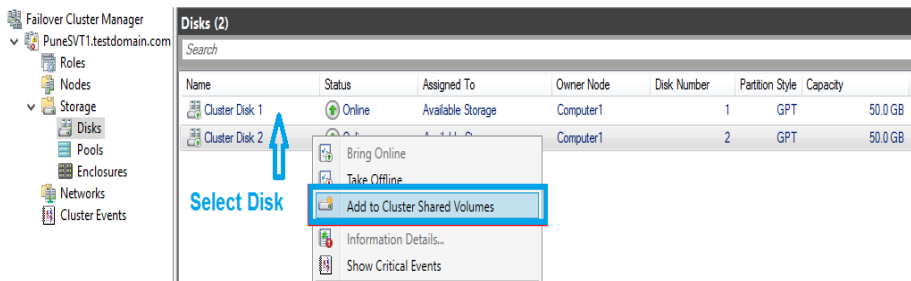


Figure 20: Failover Cluster Manager – Add to Cluster Shared Volumes

Figure 21 shows the exported disk after getting added to a CSV.

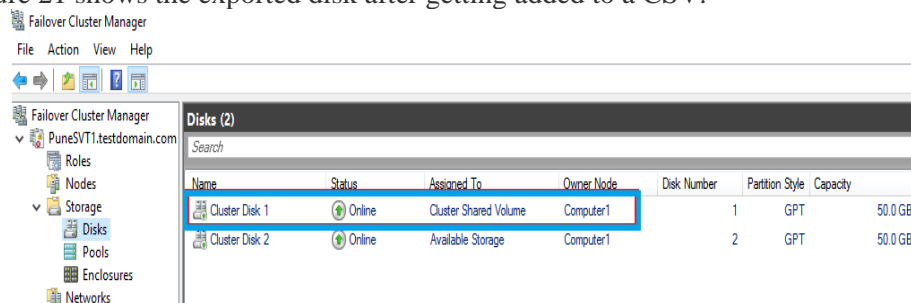


Figure 21: Failover Cluster Manager – Cluster Shared Volume configured

Best practices for iSER host attach connectivity

The following best practices are recommended for iSER host-attach connectivity:

- Same ports can be used for iSCSI and iSER host-attach concurrently.
- Single host can establish either iSCSI or an iSER host-attach session to the SVC / Storwize target.
 - iSCSI sessions should be logged out before attempting iSER host-attach. Similarly, iSER sessions should be logged out before attempting iSCSI host-attach.
 - iSER host attach sessions will be rejected if iSCSI host attach sessions from same host are present. Similarly, iSCSI host-attach sessions will be rejected if iSER host-attach sessions from same host are present.
- A maximum of up to four iSER host-attach sessions from an iSER initiator to each target node can be established.
- A one to one connectivity between an initiator host port to the target port is recommended.
- More than one session from one or more initiator host ports to the same target port will be rejected.
- A single initiator host port can establish sessions to multiple target ports (not recommended).
- A single iSER host can use either IPv4 or IPv6 to establish iSER-host attach sessions
 - IPv6 session requests will be rejected if hosts have existing IPv4 host-attach sessions. Similarly, IPv4 session requests will be rejected if hosts have existing IPv6 host-attach sessions.
 - This restriction does not apply for iSCSI sessions. iSCSI hosts can still have concurrent IPv4 and IPv6 host-attach sessions to the same target.
- RoCE / iWARP ports are iSER un-interoperable
 - A Windows iWARP initiator can establish a iSER host-attach session to iWARP target ports only. Because the target supports both RoCE and iWARP ports, it is necessary to initiate target discovery and log in from the Windows initiator host ports to SVC/Storwize target ports only.
 - Attempting to establish iSER host-attach connections from Windows iWARP ports to SVC/ Storwize RoCE ports will fail.

Features and limits

This section covers support for authentication, session limit for iSER sessions, and SCSI reservation support over iSER sessions.

Authentication

Only unidirectional CHAP is supported for iSER host-attach sessions.

Session limits

Different number of total iSER host-attach sessions is supported on different SVC/Storwize platforms. Refer to the Limits and Restrictions page for product and release to understand the exact limits.

Refer to [V8.2.1.x Configuration Limits and Restrictions for IBM System Storage SAN Volume Controller](#). A different link is available for each major release of IBM SVC and IBM Storwize family controller.

Reservation support

iSER supports SCSI reservations and provides the following features:

- SCSI2 and SCSI3 reservation are supported
- SCSI3 Persistent reservation over multiple iSER host-attach sessions is supported

Resources

The following websites provide useful references to supplement the information contained in this paper:

- IBM Storwize V7000 Support page
ibm.com/storage/support/storwize/v7000
- IBM Storwize family home page
www.ibm.com/it-infrastructure/storage/storwize
- IBM Storwize V7000 Supported hardware list
<https://www-01.ibm.com/support/docview.wss?uid=ssg1S1003703>
- IBM Storwize V7000 Knowledge Center
ibm.com/support/knowledgecenter/ST3FR7
- IBM SVC Knowledge Center
<https://www.ibm.com/support/knowledgecenter/STPVGU>
- IBM Redbooks®
ibm.com/redbooks/redbooks/pdfs/sg247432.pdf

About the authors

Rakesh Kumar Chinhara is a software engineer in the SVC team at IBM Systems Labs India. He has around 10 years of experience in IT. He holds a bachelor's degree in Electrical Engineering from the Biju Patnaik University of Technology. You can reach Rakesh at rchinhar@in.ibm.com.

Anuj Chandra is a software architect in the SVC team at IBM Systems Labs India. He has more than 19 years of experience in IT. He holds a bachelor's degree in Computer Science from the University of Pune and has studied Business Management from IIM Calcutta. You can reach Anuj at anuj.chandra@in.ibm.com



©Copyright IBM Corporation 2019
IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America
All Rights Reserved

IBM, the IBM logo and ibm.com are trademarks or registered trademarks of the Internal Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked items are marked on their first occurrence in the information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

Other product, company or service names may be trademarks or service marks of others.

References in the publication to IBM products or services do not imply that IBM intends to make them available in all countries in the IBM operates.



Please recycle
