



Managing Back-End Storage

The nodes in a cluster view the back-end storage controllers as individual disks, known as managed disks (MDisks).

This chapter includes the following sections:

- [About Managed Disks, page 4-2](#)
- [About MDisk Groups, page 4-2](#)
- [Extents, page 4-2](#)
- [MDisk Modes, page 4-2](#)
- [Isolating Back-end Storage Traffic, page 4-3](#)
- [Verifying Traffic Isolation, page 4-4](#)
- [Configuring LUN Masking, page 4-4](#)
- [Configuring LUN Masking, page 4-4](#)
- [Configuring MDisk Groups, page 4-5](#)
- [Configuring Quorum Disks, page 4-7](#)

About Managed Disks

SVC nodes see the back-end storage controllers presented to SANs as a number of individual disks, known as Managed Disks (MDisks). SVC does not attempt to provide recovery from physical disk failures within the back-end controllers. For this reason, a MDisk is usually, but not necessarily, a RAID array.

All nodes in a cluster must be able to see the same set of back-end storage ports on each back-end controller. Operation in a mode where two nodes see a different set of ports on the same controller is degraded and the system logs errors requesting a repair action. This occurs if inappropriate zoning is configured in the fabric or if inappropriate LUN masking is used. This rule has important implications for back-end storage such as FastT which impose exclusivity rules on the HBA WWNs that a storage partition can be mapped to. For such storage devices, the initiator N-port pWWNs of all nodes in a cluster must be mapped to the same set of LUNs

About MDisk Groups

MDisk Groups are collections of MDisks. A MDisk is contained within exactly one MDisk Group. The MDisk Group acts as a container into which MDisks contribute Extents (Chunks of disk blocks) and from which VDisk (logical disks associated with nodes in an I/O group) consume these extents of storage.

Extents

An extent is the unit of allocation of storage in a MDisk. Each MDisk is broken up logically into a number of extents. A MDisk does not need to be an integer multiple of extent size. SVC supports a partial extent at the end of the MDisk.

However, a VDisk occupies an integer number of extents even if the VDisk size is not an integer multiple of the extent size. The remaining space at the last extent in the VDisk remains unused.

MDisk Modes

The three MDisk modes are Image, Managed, or Unmanaged.

- Image mode

Image Mode provides a direct block-for-block translation from a MDisk to a VDisk virtualization. This mode allows virtualization of MDisks which already contain data. It allows a customer to insert SVC into the data path of an existing storage configuration with minimal downtime. Once SVC is inserted into the data path using image mode, you can use SVC's migration facilities to migrate the data to managed mode and re-arrange the data while an application is accessing the data.
- Managed mode

Disks operating in managed mode allow an arbitrary relationship between the VDisk extents and the MDisk extents. The actual mapping of the extent is based on the VDisk creation policy. The unused extents in a MDisk are available for use in creating new VDIsks data migration.
- Unmanaged mode

Mdisks in this mode do not belong to any Mdisk group

Isolating Back-end Storage Traffic

Setup the traffic isolation for all initiator N-ports in the SampleCluster into the same VSAN or zone as the back-end storage ports.

To isolate initiator N-ports, follow these steps.

Step 1 Create the target N-ports that this cluster can and should access. In this example, the N-ports to be accessed by the SampleCluster are 20:12:00:05:30:00:8d:e0 (interface fc 1/7) and 20:22:00:05:30:00:8d:e0 (interface fc 1/8).

Step 2 Create VSAN 3 for the target traffic for the SampleCluster.

```
switch# conf t
switch(config)# vsan database
switch(config-vsan-db)# vsan 3
```

Step 3 Add the two Fibre Channel N-ports connected to targets into VSAN 3.

```
switch(config-vsan-db)# vsan 3 interface fc1/7
switch(config-vsan-db)# vsan 3 interface fc1/8
switch(config-vsan-db)# exit
switch(config)#
```

Step 4 Add the CSM node's initiator N-port in the SampleCluster in VSAN 3.



Note

Each initiator must be explicitly removed from VSAN 1.

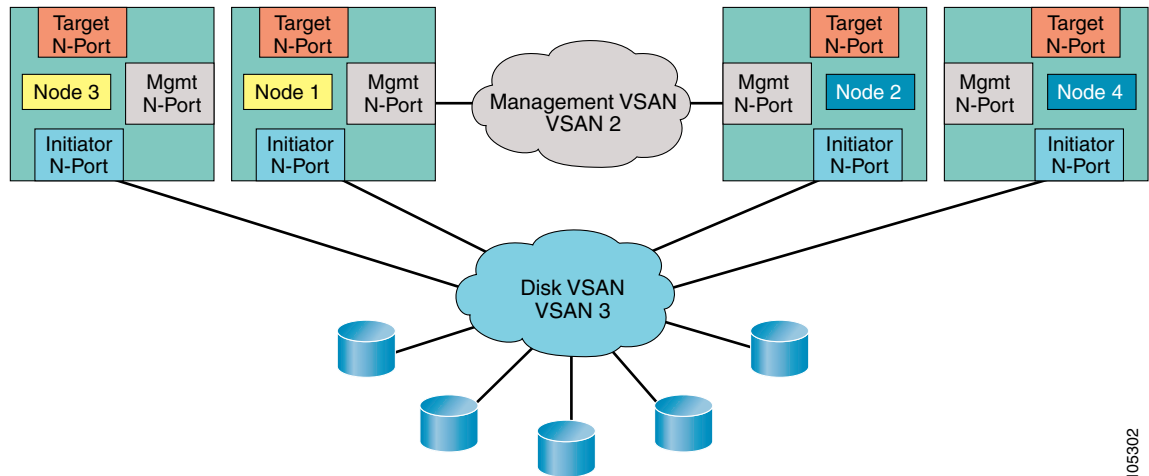
```
switch(config)# interface svc3/1
switch(config-if)# initiator vsan 3
switch(config-if)# no initiator vsan 1
switch(config-if)# exit
switch(config)# interface svc3/2
switch(config-if)# initiator vsan 3
switch(config-if)# no initiator vsan 1
switch(config-if)# exit
switch(config)# interface svc7/1
switch(config-if)# initiator vsan 3
switch(config-if)# no initiator vsan 1
switch(config-if)# exit
switch(config)# interface svc 7/2
switch(config-if)# initiator vsan 3
switch(config-if)# no initiator vsan 1
switch(config-if)# exit
```

Step 5 Configure the default zone configuration as permit for VSAN3.

```
switch(config)# zone default-zone permit vsan 3
```

You have now created a disk VSAN (see [Figure 4-1](#)).

Figure 4-1 Creating Traffic Isolation Using a Target VSAN



105302

Verifying Traffic Isolation

To verify that the target initiator ports are configured and traffic isolation has been implemented issue the `show fcn database` command for VSAN 3. `switch# show fcn database vsan 3`

VSAN 3:

| FCID | TYPE | PWWN | (VENDOR) | FC4-TYPE:FEATURE |
|----------|------|-------------------------|----------|----------------------|
| 0x730000 | N | 50:05:07:63:00:c5:9c:f9 | (IBM) | scsi-fcp:target fc.. |
| 0x730001 | N | 50:05:07:63:00:cf:9c:f9 | (IBM) | scsi-fcp:target fc.. |
| 0x730002 | N | 22:36:00:05:30:00:11:69 | (Cisco) | scsi-fcp:init svc |
| 0x730003 | N | 22:37:00:05:30:00:11:69 | (Cisco) | scsi-fcp:init svc |
| 0x730004 | N | 22:38:00:05:30:00:11:69 | (Cisco) | scsi-fcp:init svc |
| 0x730005 | N | 22:39:00:05:30:00:11:69 | (Cisco) | scsi-fcp:init svc |

Total number of entries = 6

Configuring LUN Masking

The initiator N-port pWWN of all four nodes in the cluster must be used to configure LUN masking in the backend storage controller.

```
switch# show interface svc 3/1
svc2/1 is up
Node WWN is 2e:ab:00:05:30:00:1a:e0
Fabric WWN is 20:01:00:05:30:00:1a:de
Target N-port WWN is 2e:a5:00:05:30:00:1a:e0, vsan is 1, FCID is 0xe80003
Initiator N-port WWN is 21:2e:00:05:30:00:00:21, vsan is 3, FCID is 0xea0004
Mgmt N-port WWN is 2f:af:00:05:30:00:1a:e0, vsan is 2, FCID is 0xe80000
```

```

5 minutes input rate 2392 bits/sec, 299 bytes/sec, 0 frames/sec
5 minutes output rate 2240 bits/sec, 280 bytes/sec, 0 frames/sec
272 frames input, 89764 bytes
0 discards, 0 errors
232 frames output, 84176 bytes
0 discards, 0 errors

```

Identifying MDisks

To identify MDisks, follow these steps:

- Step 1** Display the details of each MDisk using the **show cluster *cluster-name* mdisk** command. This command lists the back-end controller ports that can be accessed by the MDisk. The MDisk still do not belong to any MDisk group.

```

switch# svc-config
switch(svc)# show cluster SampleCluster mdisk
-----
id          nwnn                mdisk-grp          capacity          status
-----
1           50:05:07:63:00:c0:9c:f9          953 MB            online
2           50:05:07:63:00:c0:9c:f9          1.86 GB           online
3           50:05:07:63:00:c0:9c:f9          1.86 GB           online
4           50:05:07:63:00:c0:9c:f9          1.86 GB           online
5           50:05:07:63:00:c0:9c:f9          1.86 GB           online
6           50:05:07:63:00:c0:9c:f9          953 MB            online
7           50:05:07:63:00:c0:9c:f9          1.86 GB           online
8           50:05:07:63:00:c0:9c:f9          953 MB            online
9           50:05:07:63:00:c0:9c:f9          1.86 GB           online
10          50:05:07:63:00:c0:9c:f9          953 MB            online

```

- Step 2** Display details of the required MDisk using the **show cluster *cluster-name* mdisk id** command.

```

switch(svc)# show cluster SampleCluster mdisk id 1
mdisk id 1 is online
Is unmanaged
Controller node WWN is 50:05:07:63:00:c0:9c:f9 --> IBM ESS storage device's nWWN
Controller port WWN is 50:05:07:63:00:cf:9c:f9, LUN 00:00:00:00:00:00:00
Controller port WWN is 50:05:07:63:00:c5:9c:f9, LUN 00:00:00:00:00:00:00
Controller serial number is 07B24417
Capacity is 953 MB

```

Configuring MDisk Groups

To configure MDisk groups, follow these steps:

- Step 1** Obtain a list of candidate MDisks using the **show cluster *cluster-name* mdisk candidate** command. Select the group(s) to add the configured MDisk.

```

switch(svc)# show cluster SampleCluster mdisk candidate
-----
id          nwnn                capacity
-----
1           50:05:07:63:00:c0:9c:f9          953 MB

```

```

2          50:05:07:63:00:c0:9c:f9      1.86 GB
3          50:05:07:63:00:c0:9c:f9      1.86 GB
4          50:05:07:63:00:c0:9c:f9      1.86 GB
5          50:05:07:63:00:c0:9c:f9      1.86 GB
6          50:05:07:63:00:c0:9c:f9      953 MB
7          50:05:07:63:00:c0:9c:f9      1.86 GB
8          50:05:07:63:00:c0:9c:f9      953 MB
9          50:05:07:63:00:c0:9c:f9      1.86 GB
10         50:05:07:63:00:c0:9c:f9      953 MB

```

Step 2 Create a MDisk group called *finance*.

```

switch(svc)# cluster config SampleCluster
switch(svc-cluster)# mdisk-grp add finance extent 16
switch(svc-cluster)# mdisk-grp name finance
switch(svc-cluster-mdisk-grp)#

```

Step 3 Add MDisks 1 to 5 to the *finance* MDisk group.

```

switch(svc-cluster-mdisk-grp)# mdisk id 1
switch(svc-cluster-mdisk-grp)# mdisk id 2
switch(svc-cluster-mdisk-grp)# mdisk id 3
switch(svc-cluster-mdisk-grp)# mdisk id 4
switch(svc-cluster-mdisk-grp)# mdisk id 5

```

Step 4 Exit to the SVC cluster configuration mode.

```

switch(svc-cluster-mdisk-grp)# exit
switch(svc-cluster)#

```

Step 5 Create a MDisk group called *marketing*.

```

switch(svc-cluster)# mdisk-grp name marketing
switch(svc-cluster)# mdisk-grp add marketing extent 32
switch(svc-cluster-mdisk-grp)#

```

Step 6 Add MDisks 6 to 10 to the *marketing* MDisk group.

```

switch(svc-cluster-mdisk-grp)# mdisk id 6
switch(svc-cluster-mdisk-grp)# mdisk id 7
switch(svc-cluster-mdisk-grp)# mdisk id 8
switch(svc-cluster-mdisk-grp)# mdisk id 9
switch(svc-cluster-mdisk-grp)# mdisk id 10

```

Step 7 Exit to the SVC configuration mode.

```

switch(svc-cluster-mdisk-grp)# exit
switch(svc-cluster)# exit
switch(svc)#

```

Step 8 Verify that all configured MDisks are allocated to each MDisk group.

```

switch(svc)# show cluster SampleCluster mdisk

```

```

-----
id          nwwn                                mdisk-grp      capacity      status
-----
1           50:05:07:63:00:c0:9c:f9  finance        953 MB        online
2           50:05:07:63:00:c0:9c:f9  finance        1.86 GB       online
3           50:05:07:63:00:c0:9c:f9  finance        1.86 GB       online
4           50:05:07:63:00:c0:9c:f9  finance        1.86 GB       online
5           50:05:07:63:00:c0:9c:f9  finance        1.86 GB       online
6           50:05:07:63:00:c0:9c:f9  marketing      953 MB        online
7           50:05:07:63:00:c0:9c:f9  marketing      1.86 GB       online
8           50:05:07:63:00:c0:9c:f9  marketing      953 MB        online
9           50:05:07:63:00:c0:9c:f9  marketing      1.86 GB       online
10          50:05:07:63:00:c0:9c:f9  marketing      953 MB        online

```

- Step 9** Verify the details of each MDisk group and confirm that each MDisk group has 5 MDisks. VDIsks have not been assigned at this point.

```
switch(svc)# show cluster SampleCluster mdisk-grp
```

| name | Capacity | free | extent size(MB) | number of mdisks | number of vdisks | status |
|-----------|----------|---------|--------------------|---------------------|---------------------|--------|
| finance | 7.56 GB | 7.56 GB | 16 | 5 | 0 | online |
| marketing | 6.48 GB | 6.48 GB | 16 | 5 | 0 | online |



Note The storage capacity (953 MB) is not an integer multiple of the extent size, the last partial extent is unused.

- Step 10** To identify the number of free extents for MDisk 1, use the `show cluster cluster-name mdisk id` command.

```
switch(svc)# show cluster SampleCluster mdisk id 1
mdisk id 1 is online
  Is member of mdisk-grp finance
  Controller node WWN is 50:05:07:63:00:c0:9c:f9
  Controller port WWN is 50:05:07:63:00:cf:9c:f9, LUN 00:00:00:00:00:00:00:00
  Controller port WWN is 50:05:07:63:00:c5:9c:f9, LUN 00:00:00:00:00:00:00:00
  Controller serial number is 07B24417
  Capacity is 953 MB
  Number of free extents is 42
  Is quorum disk number 2
```



Note MDisk ID 1 (one) is a quorum disk and SVC nodes reserve some storage on the quorum disks for cluster management.

Configuring Quorum Disks

Quorum disks are used when the cluster loses half of the nodes that make up the cluster. The half of the cluster that contains the majority of the quorum disks continues to process the I/O transactions. The other half stops processing the I/O transactions. By establishing quorum disks, you can ensure that both halves of the cluster do not continue to operate.

Quorum disks are set automatically by the software. You can use the `quorum` command to overwrite the system-assigned quorum disk and pick a particular set of managed disks to be a quorum disk. The managed disk that is currently assigned the quorum index number is automatically set to be a non-quorum disk.

You can assign one of 3 possible quorum index numbers in any desired order. The quorum index number ranges from 1 to 3.



Tip We recommend that you set quorum disks on multiple controllers to avoid the possibility of losing all of the quorum disks with a single failure.

To configure quorum disk allocation, follow these steps:

Step 1 Create a cluster called SampleCluster.

```
switch(svc)# cluster config SampleCluster
switch(svc-cluster)#
```

Step 2 Sets the quorum disk ID for the specified MDisk in this cluster.

```
switch(svc-cluster)# quorum disk 2 mdisk 1
```

When you issue this command, the Mdisk that was previously-assigned the quorum index 2 will no longer be a quorum disk.



Caution

If the **quorum** command fails due to the lack of sufficient extents in the new quorum disk, the old quorum disk may no longer be operational.
