IBM

Redpaper

Roger Eriksson
Markus Oscheka
Brian Sherman
Stephen Solewin

# IBM XIV Gen3 with IBM System Storage SAN Volume Controller and Storwize V7000

This IBM® Redpaper™ publication describes preferred practices for attaching the IBM XIV® Storage System, Gen3, to either an IBM System Storage® SAN Volume Controller (SVC) or IBM Storwize® Version 7000 virtual storage. It also explains what to consider in XIV Storage System data migration when you use it in combination with an SVC or a Storwize V7000 storage.

This information is based on the assumption that you have an SVC or Storwize V7000 and that you are replacing back-end disk controllers with a new XIV system or adding an XIV as a new managed disk controller.

## Benefits of combining storage systems

By combining the IBM XIV Storage System with either the IBM SVC or the IBM Storwize V7000, you gain the benefit of the high-performance grid architecture of the XIV and retain the business benefits of the SVC or Storwize V7000. Because the SVC and Storwize V7000 have virtualization layers that can overlay multiple homogeneous and non-homogenous storage systems, virtualizing an XIV can provide the following benefits:

- Non-disruptive data movement between multiple storage systems
- IBM FlashCopy® consistency groups across multiple storage systems
- IBM Metro Mirror and IBM Global Mirror relationships between multiple storage systems
- High availability and multisite mirroring with SVC stretched cluster and VDisk mirroring
- Support for operating systems that do not offer native multipath capability or that XIV does not support (such as HP Tru64 UNIX)
- Enhanced performance by using solid-state drives (SSDs) in the SVC or Storwize V7000 or other external storage when used in combination with IBM Easy Tier®
- Use of VMware Array API Integration across multiple storage systems to allow VMware vMotion to exploit the VAAI hardware accelerated storage feature
- Use of IBM Real-time Compression™

The sections that follow address each of the requirements of an implementation plan in the order in which they arise. However, this chapter does not cover physical implementation requirements (such as power requirements) because they are already addressed in *IBM XIV Storage System Architecture and Implementation*, SG24-7659.

## Summary of steps for attaching XIV to an SVC or Storwize V7000 and migrating volumes to XIV

Review the following topics when you are placing a new XIV behind an SVC or Storwize V7000:


- “XIV and SVC or Storwize V7000 interoperability” on page 2
- “Zoning setup” on page 5
- “Volume size for XIV with SVC or Storwize V7000” on page 10
- “Using an XIV system for SVC or Storwize V7000 quorum disks” on page 15
- “Configuring an XIV for attachment to SVC or Storwize V7000” on page 17
- “Data movement strategy overview” on page 24


## XIV and SVC or Storwize V7000 interoperability

Because SVC-attached or Storwize V7000-attached hosts do not communicate directly with the XIV, only two interoperability considerations are covered in this section:

- Firmware versions
- Copy functions

## Firmware versions

The SVC or Storwize V7000 and the XIV have minimum firmware requirements. Although the versions cited in this paper were current at the time of writing, they might have changed since then. To verify the current versions, see the IBM Systems Storage Interoperation Center (SSIC) and the SVC Interoperation websites:

http://www.ibm.com/systems/support/storage/config/ssic/index.jsp
http://www.ibm.com/systems/storage/software/virtualization/svc/interop.html

### SVC firmware

The first SVC firmware version that supported XIV was 4.3.0.1. However, the SVC cluster needs to be on at least SVC firmware Version 4.3.1.4, or preferably the most recent level available from IBM.

### Storwize V7000 firmware

The Storwize V7000 was first released with Version 6.1.*x.x* firmware. Because the Storwize V7000 firmware uses the same base as the SVC, that XIV support was inherited from the SVC and is essentially the same. You can display the SVC firmware version by viewing **Monitoring** → **System** in the SVC GUI as shown in Figure 1.

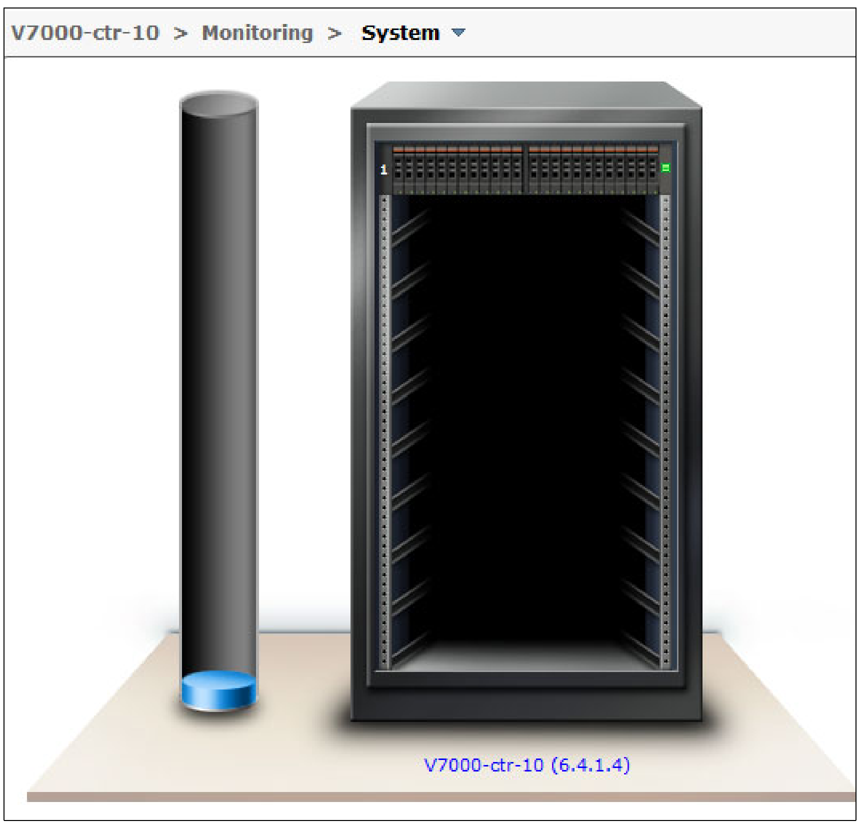


*Figure 1 Displaying Storwize V7000 firmware version*

Or you can use the **lssystem** command in the CLI as shown in Example 1. The Storwize V7000 is on code level 6.4.1.4.

*Example 1 Displaying the Storwize V7000 firmware version*

```
IBM_2145:SVC_A_B:superuser>lssystem
....
code_level 6.4.1.4 (build 75.3.1303080000)
....
```

### XIV firmware

The XIV needs to be on at least XIV firmware Version 10.0.0.a. This is an earlier level, so it is highly unlikely that your XIV is on this level. The XIV firmware version is shown on the **All Systems** → **Connectivity** panel of the XIV GUI, as shown in Figure 2. If you focused on one XIV in the GUI, you can also use **Help** → **About XIV GUI**. At this writing, the XIV is using Version 11.4.0 (circled at the upper right, in red).

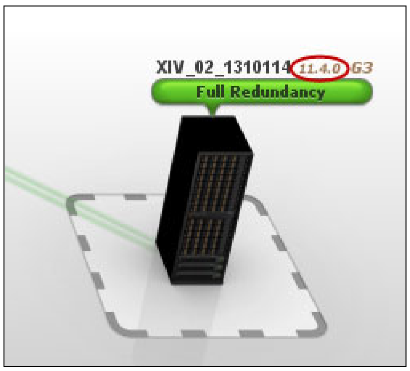


*Figure 2 Checking the XIV version*

You can check the XIV firmware version by using an XCLI command, as shown in Example 2, where the example machine uses XIV firmware Version 11.4.0.

*Example 2 Displaying the XIV firmware version*

```
XIV_02_1310114>>version_get
Version
11.4.0
```

## Copy functions

The XIV Storage System has many advanced copy and remote mirror capabilities, but for XIV volumes that are being used as SVC or Storwize V7000 MDisks (including Image mode VDisks and MDisks), none of these functions can be used. If copy and mirror functions are necessary, perform them by using the equivalent functional capabilities in the SVC or Storwize V7000 (such as SVC or Storwize V7000 FlashCopy and SVC or Storwize V7000 Metro and Global Mirror). This is because XIV copy functions do not detect write cache data resident in the SVC or Storwize V7000 cache that is not destaged. Although it is possible to disable SVC or Storwize V7000 write-cache (when creating VDisks), this method is not supported for VDisks on the XIV.

### IBM Tivoli Storage Productivity Center with XIV and SVC or Storwize V7000

IBM Tivoli® Storage Productivity Center Version 4.1.1.74 was the first version to support the XIV by using an embedded CIM object manager (CIMOM) within the XIV. The CIMOM was added in XIV code level 10.1.0.a.

Tivoli Storage Productivity Center Version 4.2 enhances XIV support by using native commands to communicate with the XIV rather than the embedded CIMOM. This enables Tivoli software for provisioning, for the Data Path Explorer view, and for performance management reporting.

Version 4.2.1 added support to detect more XIV management IP addresses and failover to these addresses (even if only one address was defined to Tivoli Storage Productivity Center).

Version 4.2.1.163 adds enhanced performance metrics when combined with XIV Firmware Version 10.2.4 and later.

Be sure to upgrade your Tivoli Storage Productivity Center to at least Version 4.2.1.163 when you combine it with XIV and SVC or Storwize V7000.

## Zoning setup

One of the first tasks of implementing an XIV system is to add it to the storage area network (SAN) fabric so that the SVC or Storwize V7000 cluster can communicate with the XIV over Fibre Channel technology. The XIV can have up to 24 Fibre Channel host ports. Each XIV reports a single worldwide node name (WWNN) that is the same for every XIV Fibre Channel host port. Each port also has a unique and persistent (WWPN). Therefore, you can potentially zone 24 unique worldwide port names (WWPNs) from an XIV to an SVC or Storwize V7000 cluster. However, the current SVC or Storwize V7000 firmware requires that one SVC or Storwize V7000 cluster cannot detect more than 16 WWPNs per WWNN, so there is no value in zoning more than 16 ports to the SVC or Storwize V7000. Because the XIV can have up to six interface modules with four ports per module, it is better to use just two ports on each module (up to 12 ports, total).

For more information, see the V6.4 Configuration Limits and Restrictions for IBM System Storage SAN Volume Controller web page:

http://www.ibm.com/support/docview.wss?uid=ssg1S1004115

When a partially populated XIV hardware is upgraded to add usable capacity, more data modules are added. At particular points in the upgrade path, the XIV gets more usable Fibre Channel ports. In each case, use half of the available ports on the XIV to communicate with an SVC or Storwize V7000 cluster (to allow for growth as you add modules). Depending on the total usable capacity of the XIV, not all interface modules have active Fibre Channel ports. Table 1 on page 6 shows which modules have active ports as capacity grows. You can also see how many XIV ports are zoned to the SVC or Storwize V7000 as capacity grows.

*Table 1 XIV host ports as capacity grows with different drive capacity*

| XIV modules | Total usable capacity in TB decimal (1 TB) | Total usable capacity in TB decimal (2 TB) | Total usable capacity in TB decimal (3 TB) | Total usable capacity in TB decimal (4 TB) | Total XIV host ports | XIV host ports to zone to an SVC or Storwize V7000 cluster | Active interface modules | Inactive interface modules |
|---|---|---|---|---|---|---|---|---|
| 6 | 28 | 55 | 84 | 112 | 8 | 4 | 4:5 | 6 |
| 9 | 44 | 88 | 132 | 177 | 16 | 8 | 4:5:7:8 | 6:9 |
| 10 | 51 | 102 | 154 | 207 | 16 | 8 | 4:5:7:8 | 6:9 |
| 11 | 56 | 111 | 168 | 225 | 20 | 10 | 4:5:7:8:9 | 6 |
| 12 | 63 | 125 | 190 | 254 | 20 | 10 | 4:5:7:8:9 | 6 |
| 13 | 67 | 134 | 203 | 272 | 24 | 12 | 4:5:6:7:8:9 | |
| 14 | 75 | 149 | 225 | 301 | 24 | 12 | 4:5:6:7:8:9 | |
| 15 | 80 | 161 | 243 | 325 | 24 | 12 | 4:5:6:7:8:9 | |

Table 2 shows another way to view the activation state of the XIV interface modules. As more capacity is added to an XIV, more XIV host ports become available. Where a module is shown as inactive, this refers only to the host ports, not the data disks.

*Table 2 XIV host ports as capacity grows*

| Module | 6 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|---|---|---|---|---|---|---|---|---|
| Module 9 host ports | **Not present** | **Inactive** | **Inactive** | Active | Active | Active | Active | Active |
| Module 8 host ports | **Not present** | Active | Active | Active | Active | Active | Active | Active |
| Module 7 host ports | **Not present** | Active | Active | Active | Active | Active | Active | Active |
| Module 6 host ports | **Inactive** | **Inactive** | **Inactive** | **Inactive** | **Inactive** | Active | Active | Active |
| Module 5 host ports | Active | Active | Active | Active | Active | Active | Active | Active |
| Module 4 host ports | Active | Active | Active | Active | Active | Active | Active | Active |

## Capacity on demand

If the XIV has the Capacity on Demand (CoD) feature, all active Fibre Channel interface ports are usable at the time of installation, regardless of how much usable capacity you purchased. For instance, if a 9-module machine is delivered with six modules active, you can use the interface ports in modules 4, 5, 7, and 8 even though, effectively, three of the nine modules are not yet activated through CoD.

## Determining XIV WWPNs

The XIV WWPNs are in the *50:01:73:8x:xx:xx:RR:MP* format, which indicates the following specifications:

| | |
|---|---|
| **5** | The WWPN format (1, 2, or 5, where XIV is always format 5) |
| **0:01:73:8** | The IEEE object identifier (OID) for IBM (formerly registered to XIV) |
| **x:xx:xx** | The XIV rack serial number in hexadecimal |
| **RR** | Rack ID (starts at 01) |
| **M** | Module ID (ranges from 4 through 9) |
| **P** | Port ID (0 to 3, although port numbers are 1 through 4) |

The module/port (MP) value that makes up the last two digits of the WWPN is shown in each small box in Figure 3 on page 8. The diagram represents the patch panel that is at the rear of the XIV rack.

To display the XIV WWPNs, use the Back view on the XIV GUI or the XIV command-line interface (XCLI) `fc_port_list` command.

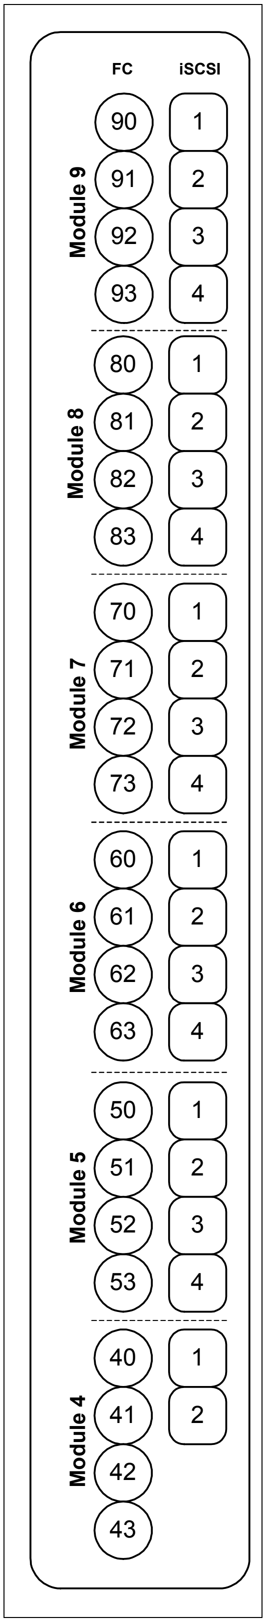


*Figure 3 XIV WWPN determination*

The output that is shown in Example 3 lists the four ports in Module 4.

*Example 3 Listing of XIV Fibre Channel host ports*

```
XIV_02_1310114>>fc_port_list module=1:Module:4
Component ID   Status   Currently Functioning   WWPN ....
1:FC_Port:4:1  OK       yes                     5001738027820140 ....
1:FC_Port:4:2  OK       yes                     5001738027820141 ....
1:FC_Port:4:3  OK       yes                     5001738027820142 ....
1:FC_Port:4:4  OK       yes                     5001738027820143 ....
```

## Hardware dependencies

There are two Fibre Channel host bus adapters (HBAs) in each XIV interface module. They are in the following locations:

- Ports 1 and 2 are on the left HBA (viewed from the rear).
- Ports 3 and 4 are on the right HBA (viewed from the rear).

Consider the following configuration information:

- Ports 1, 2, and 3 are in *SCSI target* mode by default.
- Port 4 is set to *SCSI initiator* mode by default (for XIV replication and data migration).

Use Ports 1 and 3 for SVC or Storwize V7000 traffic because both ports are on different HBAs. If you have two fabrics, place Port 1 in the first fabric and Port 3 in the second fabric.

## Sharing an XIV

It is possible to share XIV host ports between an SVC or Storwize V7000 cluster and non-SVC or non-Storwize V7000 hosts, or between two different SVC or Storwize V7000 clusters. Simply zone the XIV host ports 1 and 3 on each XIV module to either SVC or Storwize V7000 or any other hosts as required.

## Zoning rules

The XIV-to-SVC or Storwize V7000 zone needs to contain all of the XIV ports and all of the SVC or Storwize V7000 ports in that fabric. This is known as *one big zone*. This preference is relatively unique to SVC or Storwize V7000. If you zone individual hosts directly to the XIV (rather than through SVC or Storwize V7000), always use single-initiator zones, where each switch zone contains only one host (initiator) HBA WWPN and up to six XIV host port WWPNs.

For SVC or Storwize V7000, follow these rules:

- With current SVC or Storwize V7000 firmware levels, do not zone more than 16 WWPNs from a single WWNN to an SVC or Storwize V7000 cluster. Because the XIV has only one WWNN, zone no more than 16 XIV host ports to a specific SVC or Storwize V7000 cluster. If you use the suggestions in Table 1 on page 6, this restriction is not a concern.
- All nodes in an SVC or Storwize V7000 cluster must be able to see the same set of XIV host ports. Operation in a mode where two nodes see a different set of host ports on the same XIV results in the controller showing as degraded on the SVC or Storwize V7000, so the system error log requests a repair.

# Volume size for XIV with SVC or Storwize V7000

There are several considerations when you are attaching an XIV system to an SVC or Storwize V7000. Volume size is an important one. The optimum volume size depends on the maximum SCSI queue depth of the SVC or Storwize V7000 MDisks.

## SCSI queue depth considerations

Before firmware Version 6.3, the SVC or Storwize V7000 uses one XIV host port as a preferred port for each MDisk (assigning them in a round-robin fashion). Therefore, the preferred practice is to configure sufficient volumes on the XIV to ensure that the following situations are met:

- Each XIV host port receives closely matching I/O levels.
- The SVC or Storwize V7000 uses the deep queue depth of each XIV host port.

Ideally, the number of MDisks presented by the XIV to the SVC or Storwize V7000 is a multiple of the number of XIV host ports, from one to four.

Because Version 6.3 SVC or Storwize V7000 uses round-robin for each MDisk, it is no longer necessary to balance the load manually. But it is still necessary to have several MDisks because of the following queue depth limitation of SVC and Storwize V7000.

The XIV can handle a queue depth of 1400 per Fibre Channel host port and a queue depth of 256 per mapped volume per host port:target port:volume tuple. However, the SVC or Storwize V7000 sets the following internal limits:

- The maximum queue depth per MDisk is 60.
- The maximum queue depth per target host port on an XIV is 1000.

Based on this knowledge, you can determine an ideal number of XIV volumes to map to the SVC or Storwize V7000 for use as MDisks by using the following algorithm:

```
Q = ((P x C) / N) / M
```

The algorithm has the following components:

**Q** Calculated queue depth for each MDisk

**P** Number of XIV host ports (unique WWPNs) that are visible to the SVC or Storwize V7000 cluster (use 4, 8, 10, or 12, depending on the number of modules in the XIV)

**N** Number of nodes in the SVC or Storwize V7000 cluster (2, 4, 6, or 8)

**M** Number of volumes presented by the XIV to the SVC or Storwize V7000 cluster (detected as MDisks)

**C** 1000 (the maximum SCSI queue depth that an SVC or Storwize V7000 uses for each XIV host port)

If a 2-node SVC or Storwize V7000 cluster is being used with four ports on an IBM XIV System and 17 MDisks, this yields the following queue depth:

```
Q = ((4 ports*1000)/2 nodes)/17 MDisks = 117.6
```

Because 117.6 is greater than 60, the SVC or Storwize V7000 uses a queue depth of 60 per MDisk.

If a 4-node SVC or Storwize V7000 cluster is being used with 12 host ports on the IBM XIV System and 50 MDisks, this yields the following queue depth:

```
Q = ((12 ports*1000)/4 nodes)/50 MDisks = 60
```

Because 60 is the maximum queue depth, the SVC or Storwize V7000 uses a queue depth of 60 per MDisk. A 4-node SVC or Storwize V7000 is a good reference configuration for all other node configurations.

Starting with firmware Version 6.4, SVC or Storwize V7000 clusters support MDisks greater than 2 TB from the XIV system. When you use earlier versions of the SVC code, smaller volume sizes for 2 TB, 3 TB, and 4 TB drives are necessary.

This leads to the suggested volume sizes and quantities for an SVC or a Storwize V7000 system Version 6.4 or higher on the XIV with different drive capacities, as shown in Table 3. These suggestions are valid for SVC Version 6.4 or later.

*Table 3 XIV volume size and quantity recommendations*

| Modules | XIV host ports | Volume size (GB) 1 TB drives | Volume size (GB) 2 TB drives | Volume size (GB) 3 TB drives | Volume size (GB) 4 TB drives | Volume quantity | Ratio of volumes to XIV host ports |
|---|---|---|---|---|---|---|---|
| 6 | 4 | 1600 | 3201 | 4801 | 6401 | 17 | 4.3 |
| 9 | 8 | 1600 | 3201 | 4801 | 6401 | 27 | 3.4 |
| 10 | 8 | 1600 | 3201 | 4801 | 6401 | 31 | 3.9 |
| 11 | 10 | 1600 | 3201 | 4801 | 6401 | 34 | 3.4 |
| 12 | 10 | 1600 | 3201 | 4801 | 6401 | 39 | 3.9 |
| 13 | 12 | 1600 | 3201 | 4801 | 6401 | 41 | 3.4 |
| 14 | 12 | 1600 | 3201 | 4801 | 6401 | 46 | 3.8 |
| 15 | 12 | 1600 | 3201 | 4801 | 6401 | 50 | 4.2 |

**Note:** Because firmware Version 6.3 for SVC or Storwize V7000 uses round-robin scheme for each MDisk, it is not necessary to balance the load manually. Therefore, the volume quantity does not need to be a multiple of the XIV ports.

Using these volume sizes leaves free space. You can use the space for testing or for non-SVC or non-Storwize V7000 direct-attach hosts. If you map the remaining space to the SVC or Storwize V7000 as an odd-sized volume, VDisk striping is not balanced. Therefore, that I/O might not be evenly striped across all XIV host ports.

**Tip:** If you provision only part of the usable XIV space to be allocated to the SVC or Storwize V7000, the calculations no longer work. Instead, size your MDisks to ensure that at least two (up to four) MDisks are created for each host port on the XIV.

## XIV volume sizes

All volume sizes that are shown on the XIV GUI use decimal counting ($10^9$), so 1 GB = 1,000,000,000 bytes. However, a GB that is using binary counting (by using $2^{30}$ bytes) counts 1 GiB as 1,073,741,824 bytes. (It is called a *GiB* to differentiate it from a GB where size is calculated by using decimal counting.)

- By default, the SVC and Storwize V7000 use MiB, GiB, and TiB (binary counting method) for MDisk and VDisk (volume) size displays. However, the SVC and Storwize V7000 still use the terms MB, GB, and TB in the SVC or Storwize V7000 GUI and CLI output for device size displays. The SVC or Storwize V7000 CLI displays capacity in the unit is the most readable by humans.
- By default, the XIV uses GB (the decimal counting method) in the XIV GUI and CLI output for volume size displays, although volume sizes can also be shown in blocks (which are 512 bytes).

It is important to understand that a volume created on an XIV is created in 17 GB increments that are not exactly 17 GB. The size of an XIV 17-GB volume can be described in four ways:

**GB** Decimal sizing where 1 GB is 1,000,000,000 bytes

**GiB** Binary counting where 1 GiB = $2^{30}$ bytes or 1,073,741,824 bytes

**Bytes** Number of bytes

**Blocks** Blocks that are 512 bytes

Table 4 shows how these values are used in the XIV.

*Table 4 XIV space allocation in units*

| Measure | XIV |
|---|---|
| GB | 17 GB (rounded down) |
| GiB | 16 GiB (rounded down) |
| Bytes | 17,208,180,736 bytes |
| Blocks | 33,609,728 blocks |

Therefore, XIV is using binary sizing when creating volumes but displaying it in decimals and then rounding it down.

The suggested size for XIV volumes presented to the SVC or Storwize V7000 for 2 TB drives, where only 1 TB is used, is 1600 GB on the XIV. Although there is nothing special about this volume size, it divides nicely to create, on average, four to eight XIV volumes per XIV host port (for queue depth). Table 5 lists suggested volume sizes.

*Table 5 Suggested volume sizes on the XIV for 2 TB drives presented to SVC or Storwize V7000*

| Measure | XIV |
|---|---|
| GB | 1600 GB |
| GiB | 1490.452 GiB |
| Bytes | 1,600,360,808,448 bytes |
| Blocks | 3,125,704,704 blocks |

The SVC and Storwize V7000 report each MDisk presented by XIV by using binary GiB. Figure 4 shows what the XIV reports.

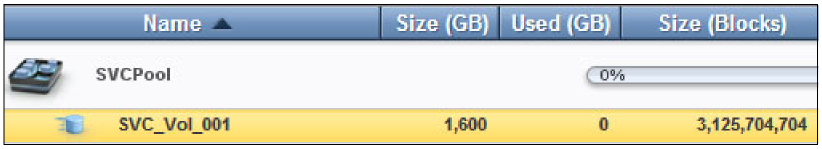


*Figure 4 An XIV volume that is sized for use with SVC or Storwize V7000*

This volume is 3,125,704,704 blocks in size. If you multiply 3,125,704,704 by 512 (because there are 512 bytes in a SCSI block), you get 1,600,360,808,448 bytes. That is exactly what the SVC or Storwize V7000 reports for the same volume (MDisk), as shown in Example 4.

*Example 4 XIV MDisk*

```
IBM_2076:V7000-ctr-10:superuser>lsmdisk -bytes
id name   status mode      mdisk_grp_id mdisk_grp_name capacity ....
0  mdisk0 online unmanaged                             1600360808448 ....
```

## Creating XIV volumes that are the same size as SVC or Storwize V7000 VDisks

To create an XIV volume that is the same size as an existing SVC or Storwize V7000 VDisk, you can use the process that is documented in “Create Image mode destination volumes on the XIV” on page 37. This is only for a transition to or from Image mode.

## SVC or Storwize V7000 2 TB volume limit with firmware earlier than 6.4

For the XIV, you can create volumes of any size up to the entire capacity of the XIV. However, in Version 6.3 or earlier of SVC or Storwize V7000 firmware, the largest XIV-presented MDisk that an SVC or Storwize V7000 can detect is 2 TiB (which is 2048 GiB).

## Creating managed disk groups

All volumes that are presented by the XIV to the SVC or Storwize V7000 are represented on the SVC or Storwize V7000 as MDisks, which are then grouped into one or more managed disk groups (MDisk groups or pools). Your decision is how many MDisk groups to use.

If you are virtualizing multiple XIVs (or other storage devices) behind an SVC or Storwize V7000, create at least one managed disk group for each additional storage device. Except for SSD-based MDisks that are used for Easy Tier, do not have MDisks from different storage devices in a common managed disk group.

In general, create only one managed disk group for each XIV, because that is the simplest and most effective way to configure your storage. However, if you have many managed disk groups, you need to understand the way that the SVC and Storwize V7000 partition cache data when they accept write I/O. Because the SVC or Storwize V7000 can virtualize storage from many storage devices, you might encounter an issue if there are slow-draining storage controllers. This occurs if write data is entering the SVC cache faster than the SVC can destage write data to the back-end disk. To avoid a situation in which a full write cache affects all storage devices that are being virtualized, the SVC partitions the cache for writes on a managed disk group level. Table 6 on page 14 shows the percentage of cache that can be used for write I/O by one managed disk group. It varies, based on the maximum number of managed disk groups that exist on the SVC or Storwize V7000.

Table 6 Upper limit of write cache data

| Number of managed disk groups | Upper limit of write cache data |
|---|---|
| 1 | 100% |
| 2 | 66% |
| 3 | 40% |
| 4 | 30% |
| 5 or more | 25% |

For example, this happens if three managed disk groups exist on an SVC or Storwize V7000, where two of them represent slow-draining, older storage devices and the third is used by an XIV. The result is that the XIV can be restricted to 20% of the SVC cache for write I/O. This might become an issue during periods of high write I/O. The solution in that case might be to have multiple managed disk groups for a single XIV. For more information, see the IBM Redpaper titled *IBM SAN Volume Controller 4.2.1 Cache Partitioning*, REDP-4426.

## SVC or Storwize V7000 MDisk group extent sizes

SVC or Storwize V7000 MDisk groups have a fixed extent size. This extent size affects the maximum size of an SVC or Storwize V7000 cluster. When you migrate SVC or Storwize V7000 data from other disk technology to the XIV, change the extent size to 1 GB (the default extent size since SVC or Storwize V7000 firmware Version 7.1). This allows for larger SVC or Storwize V7000 clusters and ensures that the data from each extent uses the striping mechanism in the XIV optimally. The XIV divides each volume into 1 MB partitions, so the MDisk group extent size in MB must exceed the maximum number of disks that are likely to exist in a single XIV footprint. For many IBM clients, this means that an extent size of 256 MB is acceptable (because 256 MB covers 256 disks, but a single XIV rack has only 180 disks). However, consider using an extent size of 1024 MB, because that size covers the possibility of using multiple XIV systems in one extent pool. Do not expect to see any difference in overall performance by using a smaller or larger extent size.

For the available SVC or Storwize V7000 extent sizes and the effect on maximum SVC or Storwize V7000 cluster size, see Table 7.

Table 7 SVC or Storwize V7000 extent size and cluster size Striped mode VDisks

| MDisk group extent size | Maximum SVC cluster size |
|---:|---:|
| 16 MB | 64 TB |
| 32 MB | 128 TB |
| 64 MB | 256 TB |
| 128 MB | 512 TB |
| 256 MB | 1 PB |
| 512 MB | 2 PB |
| 1024 MB | 4 PB |
| 2048 MB | 8 PB |
| 4096 MB | 16 PB |
| 8192 MB | 32 PB |

Create all VDisks in an XIV-based managed disk group as *striped* and striped across all MDisks in the group. This ensures that you stripe the SVC or Storwize V7000 host I/O evenly across all of the XIV host ports. Do not create sequential VDisks, because they result in uneven host port use. Use Image mode VDisks only for migration purposes.

# Using an XIV system for SVC or Storwize V7000 quorum disks

The SVC or Storwize V7000 cluster uses three MDisks as quorum disk candidates; one is active. Starting with SVC or Storwize V7000 Version 6.3, the quorum disks are selected automatically from different storage systems, if possible. The Storwize V7000 can also use internal SAS drives as quorum disks. It uses a small area on each of these MDisks or drives to store important SVC or Storwize V7000 cluster management information.

## Using an XIV for SVC or Storwize V7000 quorum disks before V6.3

If you are replacing non-XIV disk storage with XIV, ensure that you relocate the quorum disks before you remove the MDisks. Review the IBM Technote tip titled *Guidance for Identifying and Changing Managed Disks Assigned as Quorum Disk Candidates*:

http://www.ibm.com/support/docview.wss?uid=ssg1S1003311

To determine whether removing a managed disk controller requires quorum disk relocation, use the **svcinfo lsquorum** command, as shown in Example 5.

*Example 5 Using the svcinfo lsquorum command on an SVC code level 5.1 and later*

```
IBM_2145:mycluster:admin>lsquorum
quorum_index  status   id  name    controller_id  controller_name  active
0             online    0  mdisk0  0              DS6800_1         yes
1             online    1  mdisk1  1              DS6800_1         no
2             online    2  mdisk2  2              DS4700           no
```

To move the quorum disk function, specify three MDisks to become quorum disks. Depending on your MDisk group extent size, each selected MDisk must have between 272 and 1024 MB of free space. Run the **setquorum** commands before you start migration. If all available MDisk space is allocated to VDisks, you cannot use that MDisk as a quorum disk. Table 8 shows the amount of space that is required on each MDisk.

*Table 8 Quorum disk space requirements for each of the three quorum MDisks*

| Extent size | Number of extents needed by quorum | Amount of space per MDisk needed by quorum |
|---|---|---|
| 16 MB | 17 | 272 MB |
| 32 MB | 9 | 288 MB |
| 64 MB | 5 | 320 MB |
| 128 MB | 3 | 384 MB |
| 256 MB | 2 | 512 MB |
| 1024 MB or more | 1 | One extent |

## Understanding SVC and Storwize V7000 controller path values

If you display the detailed description of a controller as seen by SVC, for each controller host port, you see a path value. The **path_count** is the number of MDisks that are using that port multiplied by the number of SVC or Storwize V7000 nodes, which equals 2 in this example. In Example 6, the Storwize V7000 cluster has two nodes and can access three XIV volumes (MDisks), so 3 volumes times 2 nodes equals 6 paths per WWPN.

You can confirm that the Storwize V7000 is using all six XIV interface modules. In Example 6, because the WWPN ending in 70 is from XIV Module 7, the module with a WWPN that ends in 60 is from XIV Module 6, and so on. XIV interface modules 4 – 9 are zoned to the SVC. To decode the WWPNs, use the process described in, “Determining XIV WWPNs” on page 7.

*Example 6 Path count as seen by an SVC*

```
IBM_2076:V7000-ctr-10:superuser>lscontroller 0
id 0
controller_name XIV_02_1310114
WWNN 5001738027820000
mdisk_link_count 3
max_mdisk_link_count 4
degraded no
vendor_id IBM
product_id_low 2810XIV-
product_id_high LUN-0
product_revision 0000
ctrl_s/n 27820000
allow_quorum yes
WWPN 5001738027820150
path_count 6
max_path_count 6
WWPN 5001738027820140
path_count 6
max_path_count 6
WWPN 5001738027820160
path_count 6
max_path_count 6
WWPN 5001738027820170
path_count 6
max_path_count 6
WWPN 5001738027820180
path_count 6
max_path_count 6
WWPN 5001738027820190
path_count 6
max_path_count 6
```

## Configuring an XIV for attachment to SVC or Storwize V7000

This section presents the steps to configure the XIV and Storwize V7000 combination.

### XIV setup

Use the following steps to set up the XIV Storage System. The steps are shown for using either the XIV GUI or CLI, and you can use either method.

1. Click **Hosts and Clusters** → **Hosts and Clusters**, and then define the SVC or Storwize V7000 cluster either by clicking **Add Cluster**, as shown in Figure 5, or by using the XCLI, as shown Example 7. An SVC cluster consists of several nodes, with each SVC node defined as a separate host. Leave the default for Cluster Type. All volumes that you are mapping to the SVC or Storwize V7000 are mapped to the cluster to avoid problems with mismatched logical unit number (LUN) IDs.

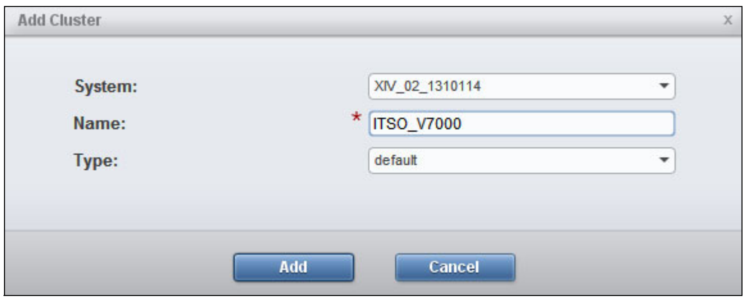


*Figure 5 Define SVC or Storwize V7000 cluster to map to the XIV*

*Example 7 Define the SVC cluster to map to the XIV*

```
cluster_create cluster="ITSO_V7000"
special_type_set cluster="ITSO_V7000" type="default"
```

2. After you click **Add Host**, define the SVC or Storwize V7000 nodes mapped to the XIV (as members of the SVC or Storwize V7000 cluster), as shown in Figure 6 and in Example 8 on page 18. By defining each node as a separate host, you can get more information about individual SVC nodes from the XIV performance statistics display. You might need to do this up to eight times, depending on how many nodes you have.

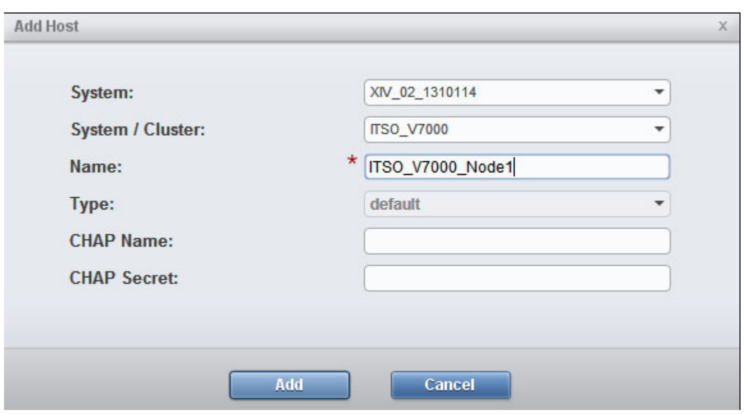


*Figure 6 Adding SVC or Storwize V7000 nodes to the cluster*

*Example 8   Define the Storwize V7000 nodes mapped to the XIV*

```
host_define host="ITSO_V7000_Node1" cluster="ITSO_V7000"
host_define host="ITSO_V7000_Node2" cluster="ITSO_V7000"
```

3. Add the SVC or Storwize V7000 host ports to the host definition of each SVC or Storwize V7000 node by right-clicking it and selecting **Add Port**, as shown in Figure 7 and in Example 9. Define up to four ports per node. If you do not know what the WWPNs for the node are, you can use the **`svcinfo lsnode`** command against each node (use an SSH session to the SVC or Storwize V7000).

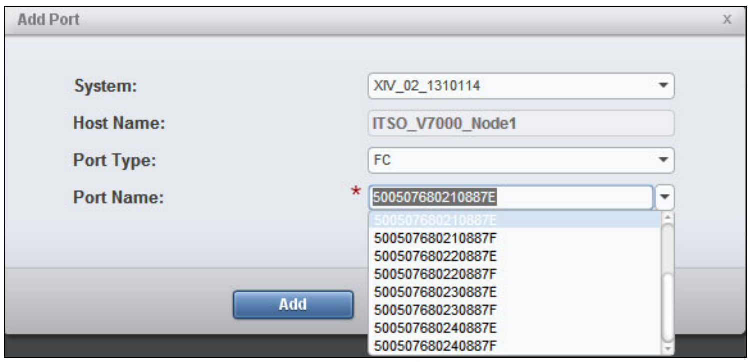


*Figure 7   Adding SVC or Storwize V7000 host ports to the hosts defined on the XIV*

*Example 9   Define the WWPNs of the first SVC node*

```
host_add_port host="ITSO_V7000_Node1" fcaddress="500507680210887E"
host_add_port host="ITSO_V7000_Node1" fcaddress="500507680220887E"
host_add_port host="ITSO_V7000_Node1" fcaddress="500507680230887E"
host_add_port host="ITSO_V7000_Node1" fcaddress="500507680240887E"
```

4. Add the SVC or Storwize V7000 host ports to the host definition of the second node.
5. Repeat Steps 3 and 4 for each SVC I/O group or Storwize V7000 cluster. If you have only two nodes (or one single-control enclosure Storwize V7000), you have only one I/O group.
6. In the XIV menu, click **Pools** → **Volumes by Pools**, and then create a storage pool by clicking **Create Pool**. Figure 8 on page 19 shows creating a pool with 10015 GB of space and no snapshot space. The total size of the pool is determined by the volume size that you choose to use.

> **Note:** The pool size example that is shown in Figure 8 on page 19 is not an ideal size, so do not use it as a guide.

You do not need the snapshot space because you cannot use XIV snapshots with SVC MDisks (instead, use SVC or Storwize V7000 snapshots at the VDisk level).

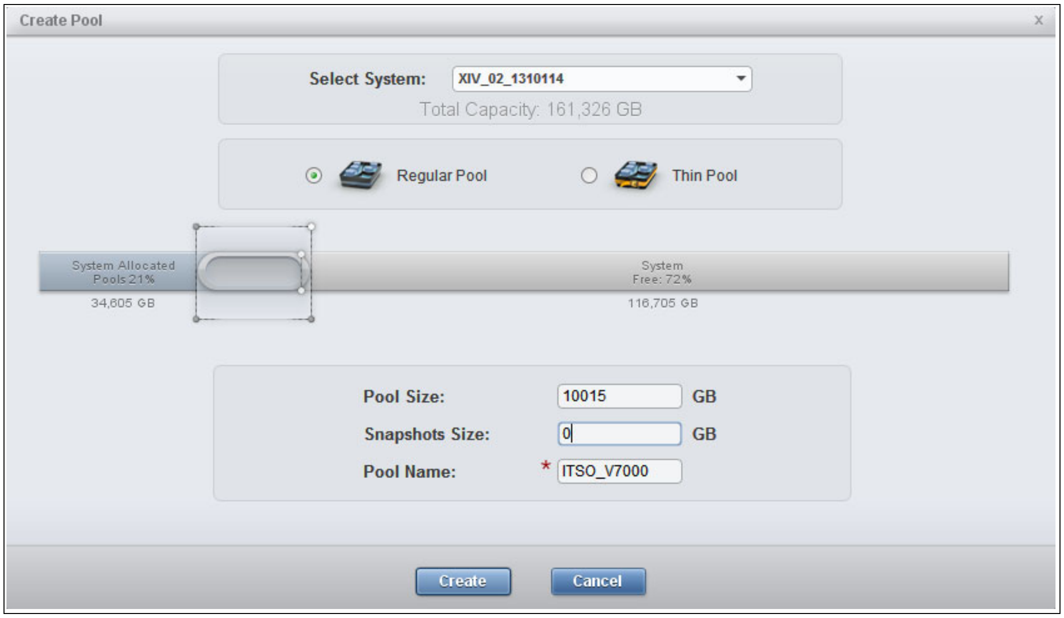


*Figure 8 Create a pool on the XIV*

*Example 10 Create a pool on the XIV*

```
pool_create pool="ITSO_V7000" size=10015 snapshot_size=0
```

> **Important:** It is not supported to use an XIV Thin Pool with the SVC or Storwize V7000. You must use a Regular Pool. Example 10 and Figure 8 show creation of a Regular Pool (where the soft size is the same as the hard size).

7. Create the volumes in the pool on the XIV by clicking **Create Volumes**, as shown in Figure 9 and in Example 11 on page 20. This example creates volumes that are 3200 GB in size because that is the suggested volume size on an XIV system with 2-TB drives.

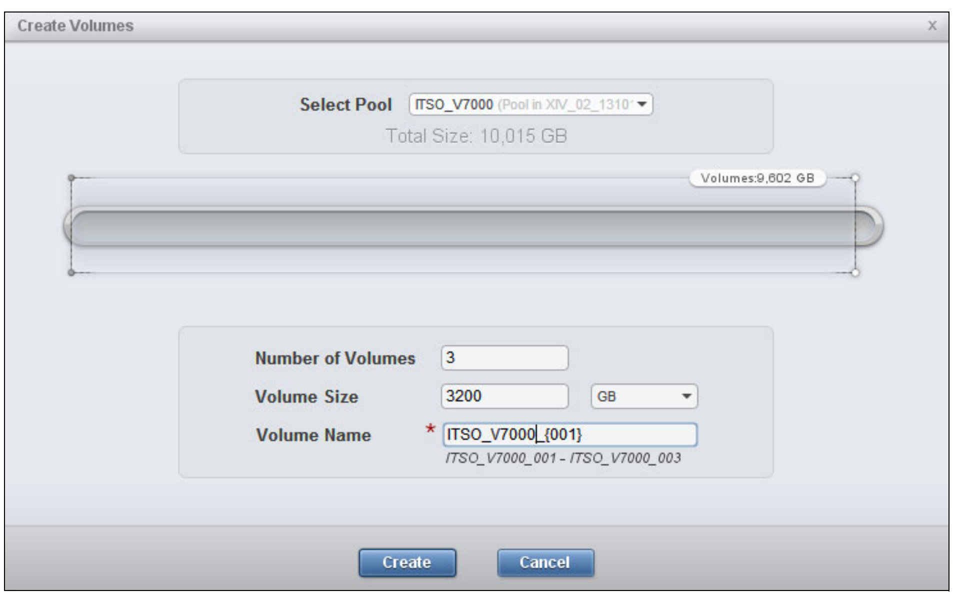


*Figure 9 Add XIV volumes*

*Example 11 Commands to create XIV volumes for use by the SVC*

```
vol_create size=3200 pool="ITSO_V7000" vol="ITSO_V7000_001"
vol_create size=3200 pool="ITSO_V7000" vol="ITSO_V7000_002"
vol_create size=3200 pool="ITSO_V7000" vol="ITSO_V7000_003"
```

8. Map the volumes. Select and right-click the volumes, and then select **Map selected volumes**, and choose either the SVC or Storwize V7000 cluster, as shown in Figure 10 and Example 12.

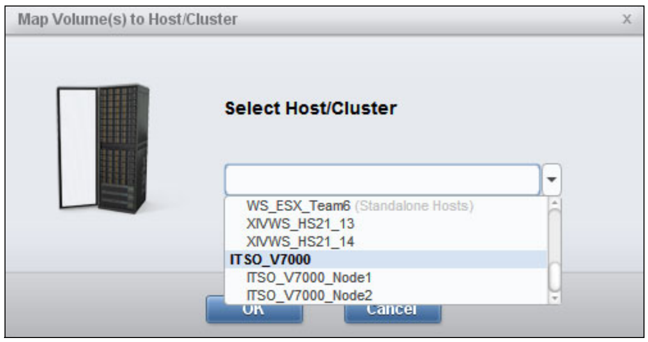


*Figure 10 Map XIV volumes to the SVC or Storwize V7000*

*Example 12 Commands to map XIV volumes to the SVC cluster*

```
map_vol cluster="ITSO_V7000" override=yes vol="ITSO_V7000_001" lun="1"
map_vol cluster="ITSO_V7000" override=yes vol="ITSO_V7000_002" lun="2"
map_vol cluster="ITSO_V7000" override=yes vol="ITSO_V7000_003" lun="3"
```

**Important:** All volumes must be mapped with the same LUN ID to all nodes of an SVC or Storwize V7000 cluster. Therefore, map the volumes to the cluster, not to individual nodes of the cluster.

**Tip:** The XIV GUI normally reserves LUN ID 0 for in-band management. The SVC or Storwize V7000 cannot take advantage of this, but neither is affected.

9. If necessary, change the system name for XIV so that it matches the controller name that is used on the SVC or Storwize V7000. If you are using the XIV GUI, click **Systems** → **System Settings** → **System** to change the XIV system name. Example 13 on page 21 uses the XCLI **`config_get`** command to determine the machine type and serial number. Then, you can use the XCLI **`config_set`** command to set the **`system_name`**. If your SVC is using code Version 5.1, limit the name to 14 characters. However, SVC and Storwize V7000 firmware Version 6.1 and later allow names of up to 63 characters.

*Example 13 Command to set the XIV system name*

```
XIV_02_1310114>>config_get
Name                                 Value
dns_primary
dns_secondary
system_name                          XIV_02_1310114
snmp_location                        Unknown
snmp_contact                         Unknown
snmp_community                       XIV
snmp_trap_community                  XIV
system_id                            10114
machine_type                         2810
machine_model                        214
machine_serial_number                1310114
email_sender_address
email_reply_to_address
email_subject_format                 {severity}: {description}
iscsi_name                           iqn.2005-10.com.xivstorage:010114
ntp_server
support_center_port_type             Management
isns_server
ipv6_state                           enabled
ipsec_state                          disabled
ipsec_track_tunnels                  no
impending_power_loss_detection_method   UPS
```

The XIV configuration tasks are now complete.

## SVC setup steps

Assuming that the SVC or Storwize V7000 is zoned to the XIV, switch to the SVC or Storwize V7000 and run the following commands:

1. Detect the XIV volumes either by using the **`svctask detectmdisk`** CLI command, or use the SVC or Storwize V7000 GUI shown in Figure 11 to click **Pools** → **MDisks by Pools** → **Detect MDisks**.

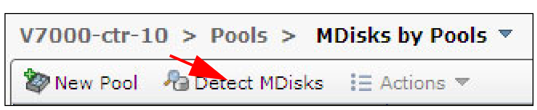


*Figure 11 Detecting MDisks by using the SVC or Storwize V7000 GUI*

2. List the newly detected MDisks, shown in Figure 12 on page 22 and Example 14 on page 22, where there are three free MDisks. They are 2.91 TiB in size (3201 GB).

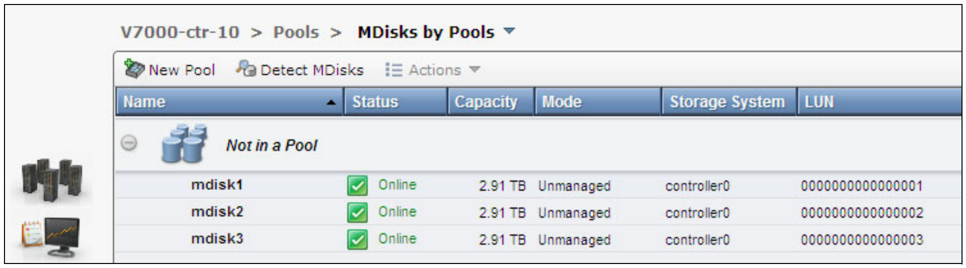


*Figure 12 XIV MDisks detected by the SVC or Storwize V7000*

*Example 14 Command for New XIV MDisks detected by SVC*

```
IBM_2076:V7000-ctr-10:superuser>lsmdisk -filtervalue mode=unmanaged
id name   status mode      mdisk_grp_id mdisk_grp_name capacity ....
1  mdisk1 online unmanaged                             2.9TB ....
2  mdisk2 online unmanaged                             2.9TB ....
3  mdisk3 online unmanaged                             2.9TB ....
```

3. Create a pool (or MDisk group) on the SVC or Storwize V7000 by clicking **New Pool**. Figure 13 shows the Create Pool window.

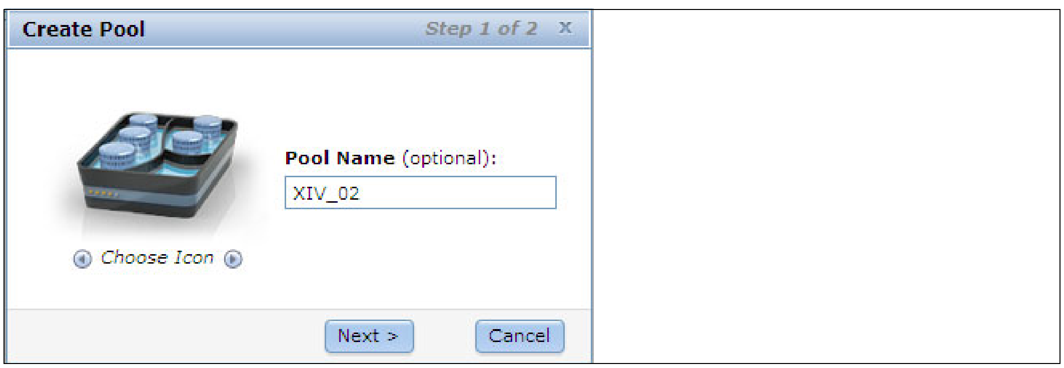


*Figure 13 Create an MDisk group or pool by using the XIV MDisks*

4. When you are prompted, select which MDisks to place into the pool, as shown in Figure 14 on page 23.

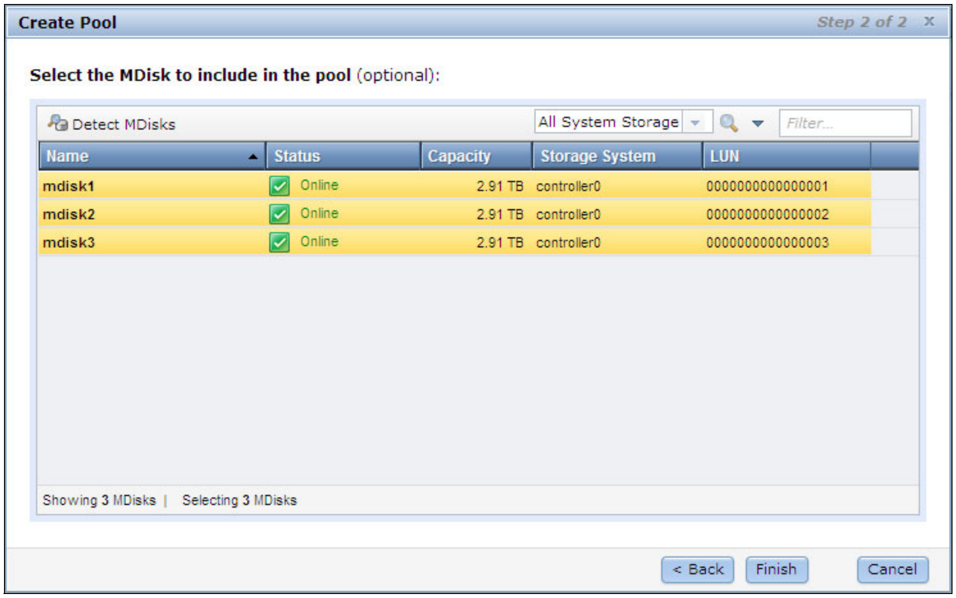


*Figure 14 Select MDisks to include in the SVC or Storwize V7000 pool*

In the Storwize V7000 or SVC GUI, there is no way to define the extent size for the pool. It defaults to 256 MB for SVC or Storwize V7000 firmware versions before 7.1. With 7.1, it defaults to 1024 MB. The SVC CLI offers the choice to specify a different extent size, but it also defaults to 256 MB before firmware Version 7.1. This is not an issue. However, if you choose to use the CLI, you can specify an extent size of 1024 MB, as shown in Example 15.

*Example 15 Command to create the MDisk group*

```
IBM_2076:V7000-ctr-10:superuser>svctask mkmdiskgrp -name XIV_02 -mdisk 1:2:3 -ext 1024
MDisk Group, id [0], successfully created
```

**Important:** Adding a new managed disk group to the SVC might result in an SVC or Storwize V7000 report that you exceeded the virtualization license limit, with an event code 3025 or a message, such as `CMMVC6373W: The virtualized storage capacity that the cluster is using exceeds the virtualized storage capacity that is licensed`. Although this does not affect operation of the SVC, you continue to receive this error message until you correct the situation either by removing the MDisk group or by increasing the virtualization license. If you are not replacing the non-XIV disk with the XIV, ensure that you purchased an additional license. Then, increase the virtualization limit by using the **`svctask chlicense -virtualization`** ***`xx`*** command (where *xx* specifies the new limit in TB). Storwize V7000 licenses external storage by enclosure rather than TB. A 15-module XIV requires 15 enclosure licenses.

5. Relocate quorum disks, if necessary, as documented in “Using an XIV system for SVC or Storwize V7000 quorum disks” on page 15.

6. Rename the default controller name by clicking the name in **Pools** → **External Storage**, as shown in Figure 15.

   A managed disk controller is given a default name by the SVC or Storwize V7000, such as controller0 or controller1 (depending on how many controllers are detected). Because the XIV can have a system name that is defined for it, match the two names closely. The controller name used by SVC Version 5.1 code cannot have spaces and cannot be more than 15 characters long. SVC firmware Version 6.1 and later and Storwize V7000 firmware have a 63-character name limit.


   *Figure 15 Rename the controller*

   In Example 16, controller number 2 is renamed to match the system name that is used by the XIV (which was set in Example 13 on page 21).

   *Example 16 Command to rename the XIV controller definition at the SVC*

```
IBM_2076:V7000-ctr-10:superuser>lscontroller
id controller_name ctrl_s/n             vendor_id            product_id_low
product_id_high
0  controller0     27820000             IBM                  2810XIV-
LUN-0
IBM_2076:V7000-ctr-10:superuser>chcontroller -name "XIV_02_1310114" 0
```

Now you can follow one of the migration strategies described in the next section.

# Data movement strategy overview

Three possible data movement strategies are described in this and subsequent sections. For more information about SVC data movement (migration), see these two IBM Redbooks® publications:

- *Data Migration to IBM Disk Storage Systems*, SG24-7432
- *Implementing the IBM System Storage SAN Volume Controller V6.3*, SG24-7933

## Using SVC migration to move data

You can use standard SVC and Storwize V7000 migration to move volumes from MDisk groups with MDisks from a non-XIV disk controller to MDisk groups with MDisks from XIV. This process does *not* require a host outage, but it does *not* allow the MDisk group extent size to be changed.

The process requires the following steps:

1. Start with existing volumes (VDisks) in an existing pool (MDisk group) that uses MDisks from an old storage device. Confirm the extent size of that MDisk group. This is called the *source* MDisk group.

2. Create new volumes on the XIV, and map these to the SVC or Storwize V7000 by using volumes that are sized according to the principles described in “XIV volume sizes” on page 12.
3. On the SVC or Storwize V7000, detect these new MDisks and use them to create a new MDisk group. This is called the *target* MDisk group. The target MDisk group must use the same extent size as the source MDisk group.
4. Migrate each VDisk from the source MDisk group to the target MDisk group.
5. When all of the VDisks have been migrated, you can choose to delete the source MDisk group (in preparation for removing the non-XIV storage), which puts the MDisks in this group in Unmanaged mode.

For more information, see “Using SVC or Storwize V7000 migration to move data to XIV” on page 26.

## Using VDisk mirroring to move the data

You can use the VDisk copy (mirror) function in SVC firmware Version 4.3 and later to create two copies of the data, one in the source MDisk group and one in the target MDisk group. Remove the VDisk copy in the source MDisk group (pool) and retain the VDisk copy present in the target MDisk group. This process does *not* require a host outage, and it enables you to move to a larger MDisk group extent size. However, it also uses more SVC or Storwize V7000 cluster memory and processor power while the multiple copies are managed by the SVC or Storwize V7000.

The process requires the following steps:

1. Start with existing VDisks in an existing MDisk group. The extent size of that MDisk group is not relevant. This is called the *source* MDisk group.
2. Create new volumes on the XIV, and map these to the SVC or Storwize V7000 by using volumes that are sized according to the principles described in, “XIV volume sizes” on page 12.
3. Detect these XIV MDisks, and create an MDisk group by using an extent size of 1024 MB. This is called the *target* MDisk group.
4. For each VDisk in the source MDisk group, create a VDisk copy in the target MDisk group.
5. When the two copies are in sync, remove the VDisk copy from the source MDisk group.
6. When all of the VDisks are copied from the source MDisk group to the target MDisk group, either delete the source MDisk group (in preparation for removing the non-XIV storage) or split the VDisk copies and retain the copy from the source MDisk group for as long as necessary.

For more information, see “Using VDisk mirroring to move the data” on page 30.

## Using SVC migration with Image mode

This migration method is used in the following situations:

- The extent size must be changed, but VDisk mirroring cannot be used.
- You want to move the VDisks from one SVC cluster to a different one.
- You want to move the data away from the SVC without using XIV migration.

If you take the Image mode VDisk from the SVC or Storwize V7000 to move it to another system, there will be a host outage, although you can keep it short (potentially only seconds or minutes).

Removing an Image mode VDisk requires the following steps:

1. Start with VDisks in an existing MDisk group. The extent size of this MDisk group might be small (for example, 16 MB). This is the *source* MDisk group.
2. Create XIV volumes that are the same size (or larger) than the existing VDisks. This might require extra steps, because the XIV volumes must be created by using 512-byte blocks. Map these specially sized volumes to the SVC.
3. Migrate each VDisk to Image mode by using these new volumes (presented as unmanaged MDisks). The new volumes move into the source MDisk group as Image mode MDisks, and the VDisks become Image mode VDisks.
4. Delete all of the Image mode VDisks from the source MDisk group. This is the disruptive part of this process. These are now unmanaged MDisks, but the data on these volumes is intact. You can map these volumes to a different SVC or Storwize V7000 cluster, or you can remove them from the SVC (in that case, the process is complete).

The following steps outline the process for bringing in an Image mode VDisk:

1. Create another managed disk group by using free space on the XIV. Use an extent size of 1024 MB. This is the *target* MDisk group.
2. Import the Image mode MDisks to managed mode VDisks by importing the MDisks into a *transition* MDisk group (which is automatically created by the import), and then move the data from the transition MDisk group to the *target* MDisk group. The MDisks are already on the XIV.
3. When the process is complete, delete the transition MDisk group, unmap the unmanaged MDisk volumes on the XIV, and delete these XIV volumes so that you can reuse the space. Use the space to create more volumes to present to the SVC. You can add these to the existing MDisk group or use them to create a new one.

For more information, see “Using SVC migration with Image mode” on page 37.

# Using SVC or Storwize V7000 migration to move data to XIV

This process migrates data from a source MDisk group (pool) to a target MDisk group (pool) by using the same extent size. This process has no interruption to the host I/O.

## Determine the required extent size and VDisk candidates

You must determine the extent size of the source MDisk group. Click **Pools** → **MDisks by Pools**, select **mdisk4**, right-click it, and select **Properties**. Figure 16 on page 27 shows the properties of mdisk4, and the stripe size (extent size) is 256 (MB).

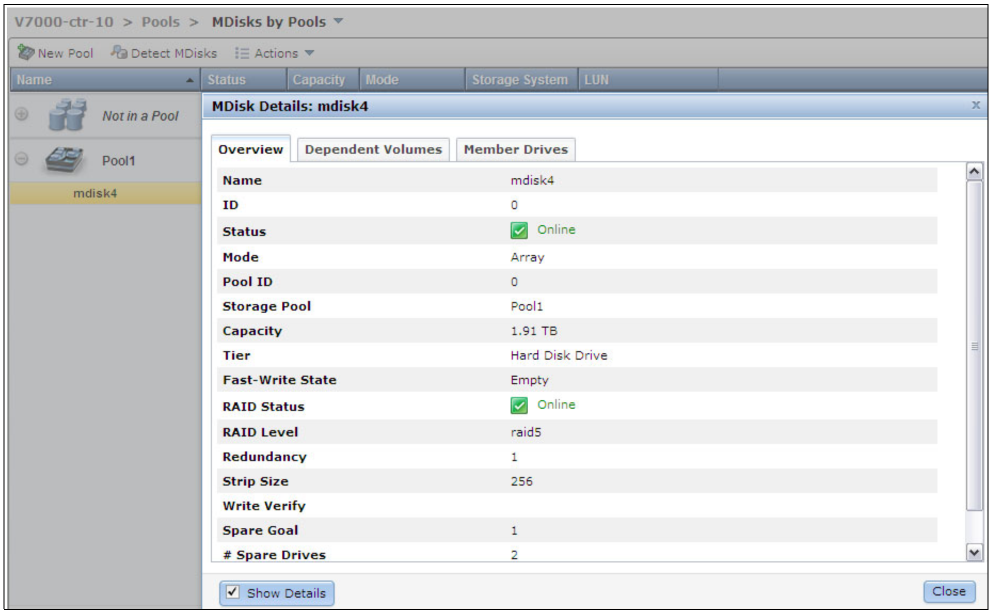


*Figure 16 List MDisk properties*

In Example 17, MDisk group ID 0 is the source group and has an extent size of 256 (MB).

*Example 17 Command to list MDisk groups*

```
IBM_2076:V7000-ctr-10:superuser>lsmdiskgrp
id name  status mdisk_count vdisk_count capacity extent_size ....
0  Pool1 online 1           1           1.91TB   256 ....
```

To identify the VDisk to migrate, click **Pools** → **Volumes by Pool**. In Pool1, VDisk Vol1 is displayed, as shown in Figure 17.


*Figure 17 Volumes by Pool window*

To get VDisks that belong to one pool (MDisk group), filter by MDisk group ID, as shown in Example 18, where only one VDisk must be migrated.

*Example 18 Command to list VDisks filtered by the MDisk group*

```
IBM_2076:V7000-ctr-10:superuser>lsvdisk -filtervalue mdisk_grp_id=0
id name IO_group_id IO_group_name status mdisk_grp_id mdisk_grp_name capacity ....
0  Vol1 0           io_grp0       online 0            Pool1          5.00GB ....
```

## Create the MDisk group

You must create volumes on the XIV and map them to the Storwize V7000 cluster. After you do this, you can detect them under **Pools** → **MDisks by Pools** by clicking **Detect MDisks**, as shown in Figure 18, where three MDisks are displayed.

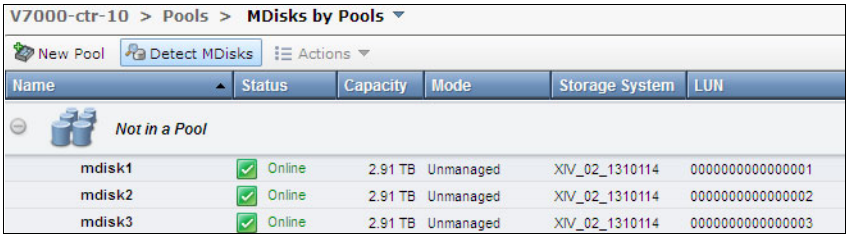


*Figure 18 Detect MDisks*

Example 19 shows the accompanying CLI commands.

*Example 19 Commands for detecting new MDisks*

```
IBM_2076:V7000-ctr-10:superuser>detectmdisk
IBM_2076:V7000-ctr-10:superuser>lsmdisk -filtervalue mode=unmanaged
id name   status mode      mdisk_grp_id mdisk_grp_name capacity ctrl_LUN_# ....
1  mdisk1 online unmanaged                             2.9TB    0000000000000001 ....
2  mdisk2 online unmanaged                             2.9TB    0000000000000002 ....
3  mdisk3 online unmanaged                             2.9TB    0000000000000003 ....
```

Then, create a pool called `XIV02` by clicking **New Pool** and using the new XIV MDisks, with the same default extent size 256 MB as the source group, as shown in Figure 19.

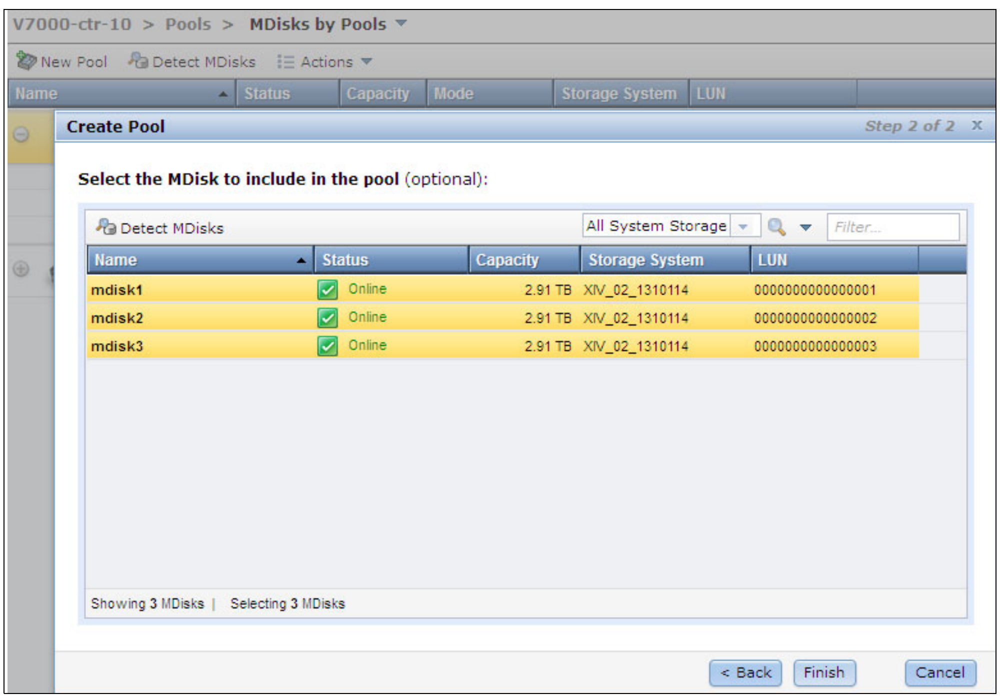


*Figure 19 Create a new pool*

The CLI command that the GUI invokes to create the pool is shown in Example 20.

*Example 20 Command to create an MDisk group*

```
svctask mkmdiskgrp -easytier auto -ext 256 -guiid 0 -mdisk 3:2:1 -name XIV02 -warning 80%
The pool (ID 1) was successfully created.
```

## Migration

Now you are ready to migrate the VDisk.

Click **Pools** → **Volumes by Pool**, right-click **Vol1**, and select **Migrate to Another Pool**. Then, select pool **XIV02** as illustrated in Figure 20.

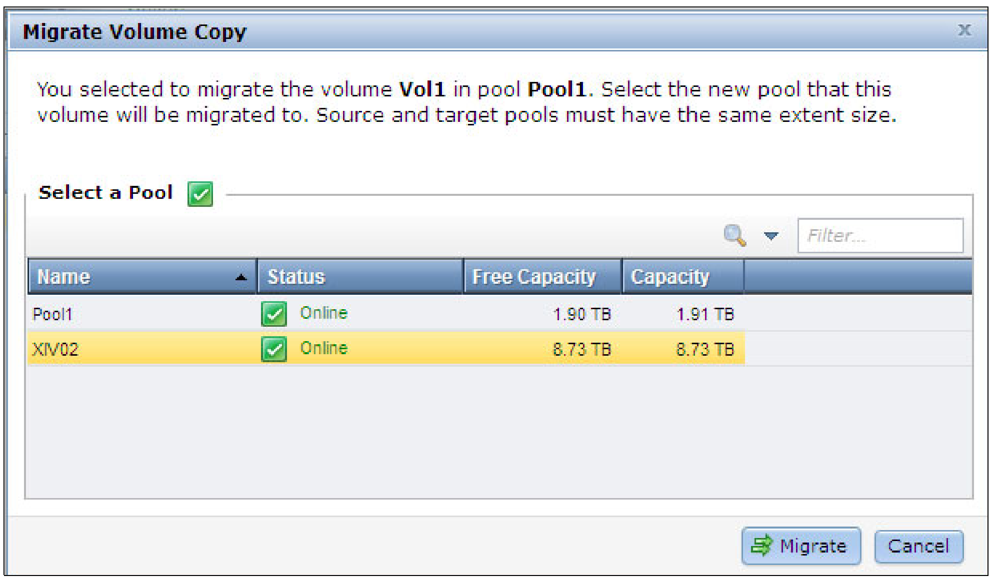


*Figure 20 Migrate the VDisk to another pool*

Example 21 shows the CLI command that the GUI invokes to migrate the VDisk to another pool.

*Example 21 Command for migrating a VDisk*

```
Svctask migratevdisk -copy 0 -mdiskgrp XIV02 -vdisk Vol1
```

After all VDisks are migrated from the source MDisk group, you can delete the source MDisk group by clicking **Pools** → **MDisks by Pools,** right-clicking **Pool1**, and selecting **Delete Pool**, which opens the Delete Pool window shown in Figure 21 on page 30.

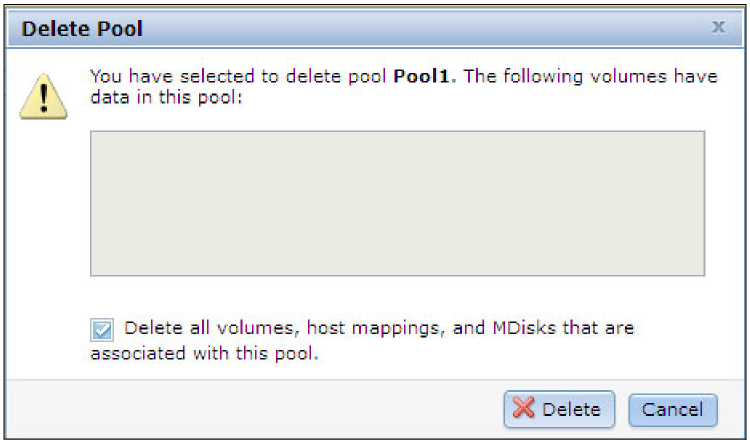

*Figure 21 Delete Pool window*

The corresponding CLI command is shown in Example 22.

*Example 22 Command for removing non-XIV MDisks and the MDisk group*

```
Svctask rmmdiskgrp -force 0
```

**Important:** Scripts that use VDisk names or IDs are not affected by the use of VDisk migration because the VDisk names and IDs do not change.

# Using VDisk mirroring to move the data

This process mirrors data from a source MDisk group to a target MDisk group by using a different extent size, with no interruption to the host.

## Determine the required extent size and VDisk candidates

The determination of the source MDisk group is as shown in Figure 16 on page 27 and in Example 17 on page 27.

The identification of the VDisks that you are migrating is as shown in Figure 17 on page 27 and in Example 18 on page 27.

## Create the MDisk group

Create an MDisk group by using an extent size of 1024 MB with the three free MDisks that are shown in Figure 18 on page 28 and in Example 19 on page 28. In Example 23, the MDisk group with ID of one is created (id [1]). With SVC or Storwize V7000 firmware Versions before 7.1, you must use the CLI to create an MDisk group with an extent size different from the default size of 256 MB.

*Example 23 Command for creating an MDisk group*

```
IBM_2076:V7000-ctr-10:superuser>mkmdiskgrp -name XIV02 -mdisk 1:2:3 -ext 1024
MDisk Group, id [1], successfully created
```

## Set up the I/O group for mirroring

The I/O group requires reserved memory for mirroring. First, determine whether the memory is reserved. You can do that only by using the CLI. In Example 24, the output of the **lsiogrp** command shows that enough space is assigned for mirroring.

*Example 24 Checking the I/O group for mirroring*

```
IBM_2076:V7000-ctr-10:superuser>lsiogrp 0
id 0
name io_grp0
node_count 2
vdisk_count 1
host_count 0
flash_copy_total_memory 20.0MB
flash_copy_free_memory 20.0MB
remote_copy_total_memory 20.0MB
remote_copy_free_memory 20.0MB
mirroring_total_memory 20.0MB
mirroring_free_memory 20.0MB
raid_total_memory 40.0MB
raid_free_memory 39.6MB
maintenance no
compression_active no
accessible_vdisk_count 1
compression_supported yes
```

## Create the mirror

To create the VDisk mirror, click **Pools** → **Volumes by Pool**, right-click **Vol1**, and select **Volume Copy Actions** → **Add Mirrored Copy**.

When the Add Volume Copy window opens, select **XIV02** as the pool, and click **Add Copy**, as shown in Figure 22 on page 32.

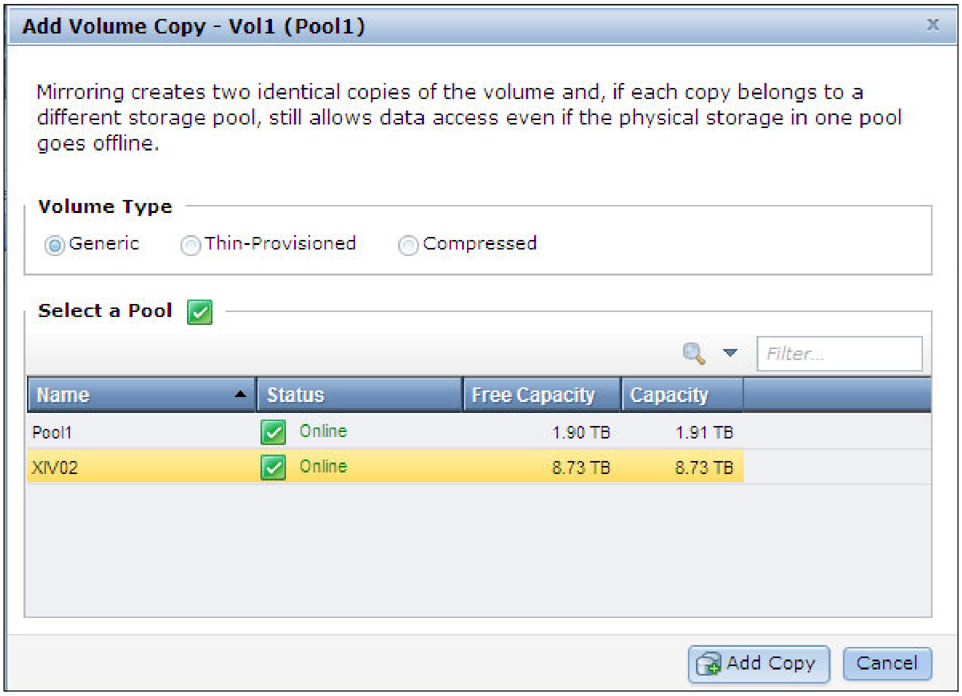


*Figure 22 Add Volume Copy window*

Example 25 shows the CLI command that the GUI invokes.

*Example 25 Command for adding the volume copy*

```
svctask addvdiskcopy -mdiskgrp 1 0
```

In Figure 23, the second volume copy is visible.

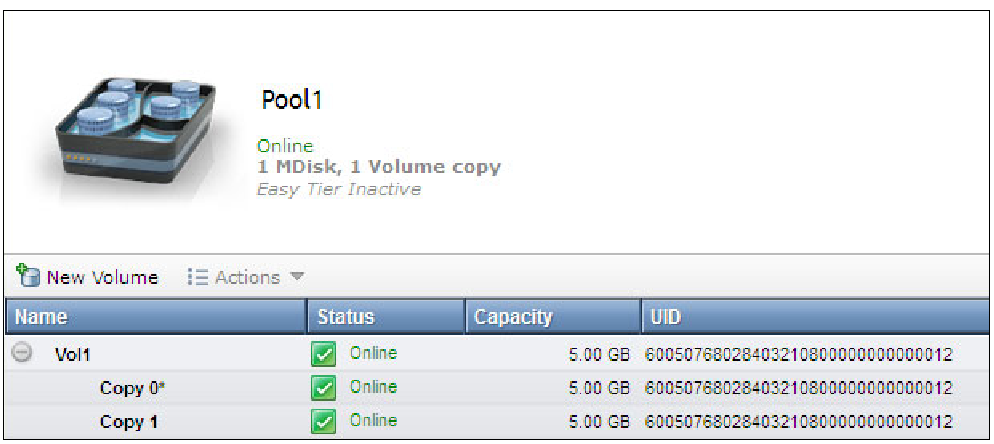


*Figure 23 Volume copies listed*

In Example 26, you can see the two copies (they are not yet in sync).

*Example 26 Command for monitoring mirroring progress*

```
IBM_2076:V7000-ctr-10:superuser>lsvdiskcopy
vdisk_id vdisk_name copy_id status sync primary mdisk_grp_id mdisk_grp_name capacity ....
0        Vol1       0       online yes  yes     0            Pool1          5.00GB ....
0        Vol1       1       online no   no      1            XIV02          5.00GB ....
```

To display the progress of the VDisk copy process, click **Running Tasks** at the bottom of the GUI, and select **Volume Synchronization**, as shown in Figure 24.

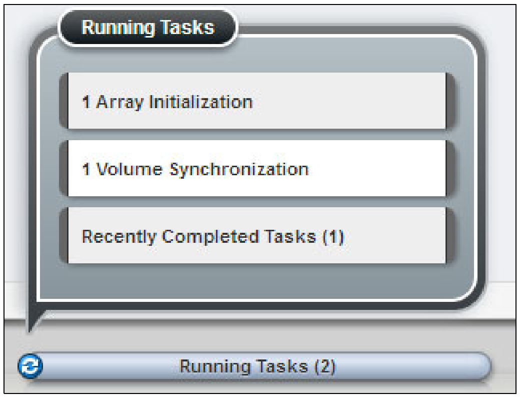


*Figure 24 Running Tasks display*

Figure 25 displays the progress.

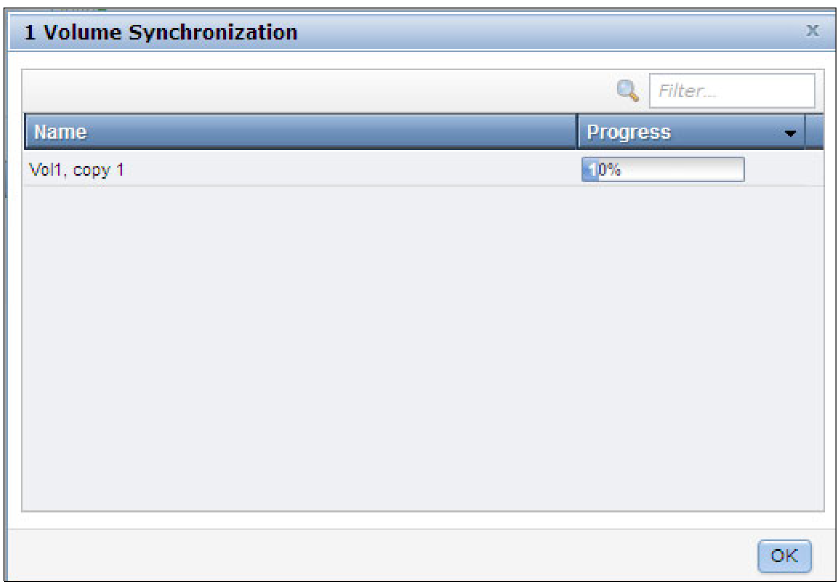


*Figure 25 Volume Synchronization progress display*

Example 27 shows the CLI command to monitor the VDisk sync progress.

*Example 27 Command for checking the VDisk sync*

```
IBM_2076:V7000-ctr-10:superuser>lsvdisksyncprogress
vdisk_id vdisk_name copy_id progress estimated_completion_time
0        Vol1       1       11       131007134415
```

If copying proceeds too slowly, you can set a higher sync rate when you create the copy or at any time while the copy process is running. The default value is 50 (2 MBps), and the maximum value is 100 (64 MBps). This change affects the VDisk and any future copies.

Right-click **Vol1**, and select **Properties**.

Change the value of Mirror sync rate by using the Volume Details dialog window that is shown in Figure 26 on page 34.

Click **Save** after you change the sync rate.

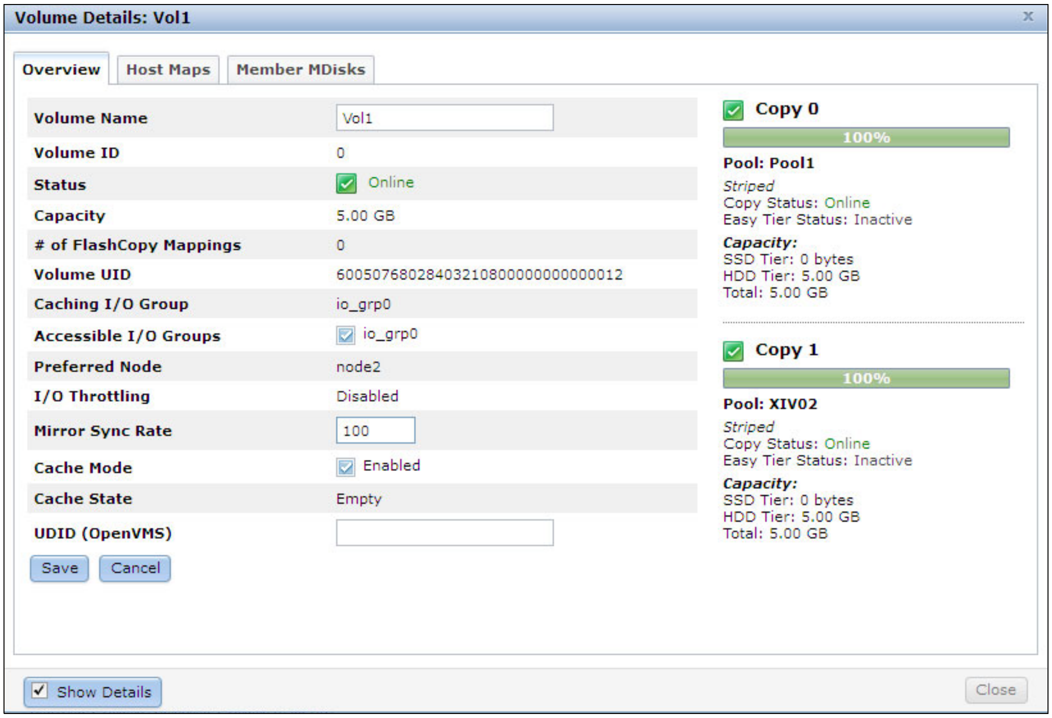


*Figure 26 Volume Details window*

Example 28 shows the corresponding CLI command.

*Example 28 Command for changing the VDisk sync rate*

```
svctask chvdisk -syncrate 100 Vol1
```

There are several possible sync rates, as listed in Table 9 on page 35. A value of zero prevents synchronization.

*Table 9 VDisk copy sync rates*

| VDisk sync rate | Actual copy speed per second |
|---|---|
| 0 | Prevents synchronization |
| 1 – 10 | 128 KBps |
| 11 – 20 | 256 KBps |
| 21 – 30 | 512 KBps |
| 31 – 40 | 1 MBps |
| 41 – 50 | 2 MBps |
| 51 – 60 | 4 MBps |
| 61 – 70 | 8 MBps |
| 71 – 80 | 16 MBps |
| 81 – 90 | 32 MBps |
| 91 – 100 | 64 MBps |

If you want to display the sync rates of all defined VDisks, paste the entire command that is shown in Example 29 into an SSH session.

*Example 29 Command to display VDisk sync rates on all VDisks*

```
svcinfo lsvdisk -nohdr |while read id name IO_group_id;do svcinfo lsvdisk $id
|while read id value;do if [[ $id == "sync_rate" ]];then echo $value"
"$name;fi;done;done
```

If you want to change the sync rate on all VDisks at the same time, paste the entire command that is shown in Example 30 into an SSH session. This example command sets the **syncrate** to 50 (2 MBps, which is the default). To set the **syncrate** on every VDisk to another value, change the value in the command from 50 to another number.

*Example 30 Command for changing the VDisk syncrate on all VDisks at the same time*

```
svcinfo lsvdisk -nohdr |while read id name IO_group_id;do svctask chvdisk
-syncrate 50 $id;done
```

## Removing or splitting the VDisk copy

Now that the synchronization is finished, you can remove Copy 0 from the VDisk so that the VDisk uses only Copy 1 (which must be on the XIV). You can do this either by removing one copy or by splitting the copies.

### Remove the VDisk copy

To delete Copy 0, right-click the copy, and select **Delete this Copy**, as shown in Figure 27 on page 36. This discards the VDisk copy in the MDisk source group. This is simple and quick, but it has one disadvantage, which is that you must mirror the data if you decide to reverse the change.

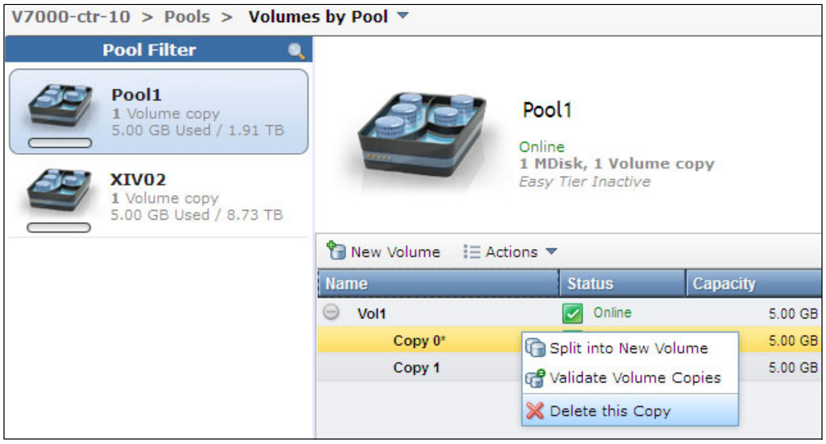


*Figure 27 Delete the volume copy*

Example 31 shows the corresponding CLI command.

*Example 31 Command for removing the VDisk copy*

```
svctask rmvdiskcopy -copy 0 0
```

## Split the VDisk copies

Figure 28 and Figure 29 on page 37 show splitting the VDisk copies and moving Copy 0 (on the Pool1 MDisk group) to become a new, unmapped VDisk. This means that the host continues to access Copy 1 on the XIV02 MDisk group as VDisk 0. The advantage of doing this is that the original VDisk copy remains available if you decide to reverse the action (although it might no longer be in sync after you split the copies). Another step is necessary to specifically delete the new VDisk that was created when you split the two.

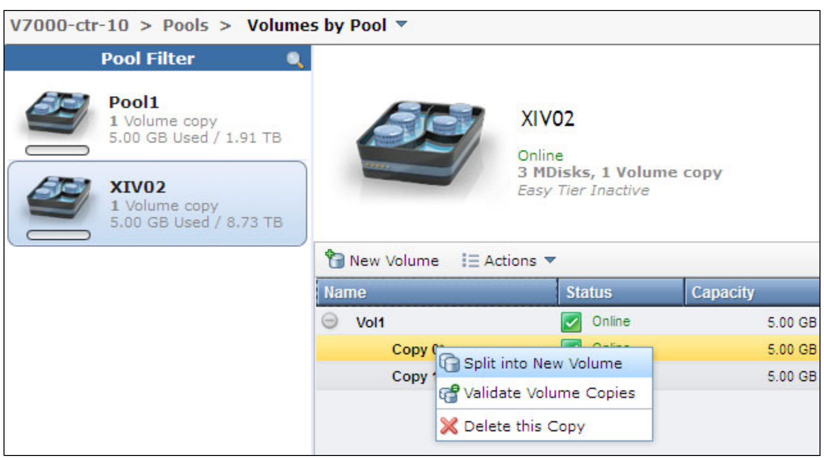


*Figure 28 Split into New Volume selected in Volumes by Pool window*

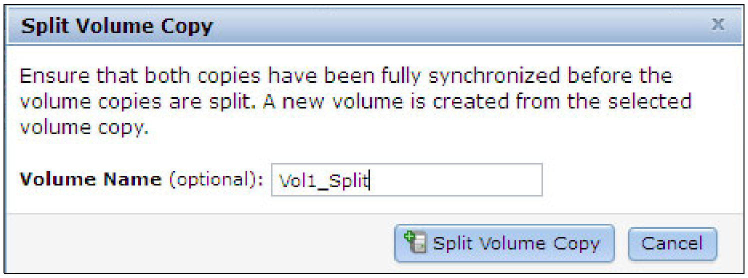


*Figure 29 Split the volume copy and enter a new volume name*

Example 32 shows the corresponding CLI command.

*Example 32 Command for splitting the VDisk copies*

```
svctask splitvdiskcopy -copy 0 -name Vol1_Split 0
```

**Important:** Scripts that use VDisk names or IDs are not affected by the use of VDisk mirroring because the VDisk names and IDs do not change. However, if you split the VDisk copies and continue to use Copy 0, it will be a new VDisk with a new name and a new ID.

# Using SVC migration with Image mode

This process converts SVC or Storwize V7000 VDisks to Image mode MDisks on the XIV. You can reassign the XIV to a different SVC or release it from the SVC (you can use that method as a way of migrating from SVC to XIV). Because of this extra step, the XIV might require sufficient space to hold both the transitional volumes (for Image mode MDisks) and the final destination volumes (for managed mode MDisks if you migrate to managed mode).

## Create Image mode destination volumes on the XIV

On the XIV, you must create one new volume for each SVC VDisk that you are migrating (it must be the exact size as the source VDisk or larger). These are to allow transition of the VDisk to Image mode. To do this, you must determine the size of the VDisk so that you can create a matching XIV volume.

To identify the VDisk to be migrated, click **Pools** → **Volumes by Pool**. In Pool1, VDisk Vol2 is displayed as shown in Figure 30 on page 38.

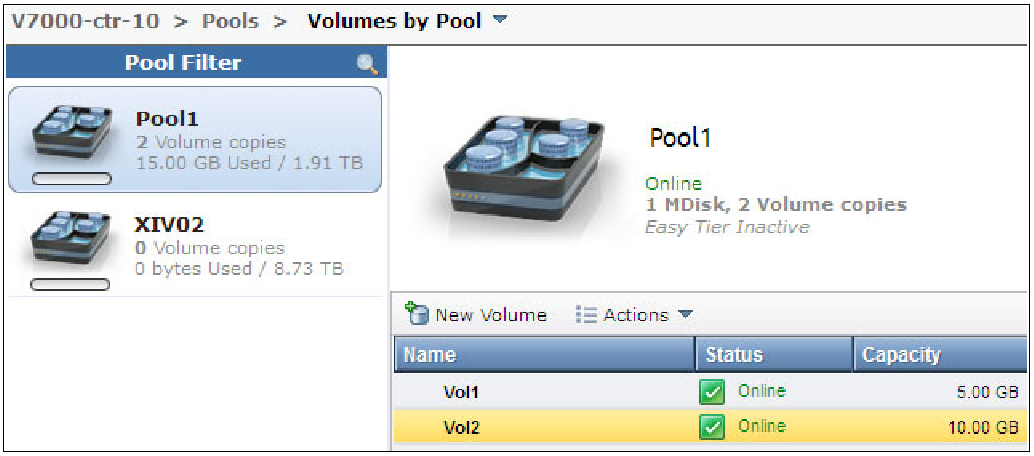


*Figure 30 Volumes by Pool window*

To ensure that the VDisks belong to one pool (MDisk group), by using the CLI, filter by MDisk group ID. In Example 33, Vol2 is the VDisk that must be migrated.

*Example 33 Command for listing VDisks*

```
IBM_2076:V7000-ctr-10:superuser>lsvdisk -filtervalue mdisk_grp_id=0
id name IO_group_id IO_group_name status mdisk_grp_id mdisk_grp_name capacity ....
0  Vol1 0           io_grp0       online 0            Pool1          5.00GB ....
1  Vol2 0           io_grp0       online 0            Pool1          10.00GB ....
```

## Option 1. Create a larger volume

Now, you can make an XIV volume that is either larger than the source VDisk or that is exactly the same size. The easier solution is to create a larger volume. Because the XIV creates volumes in rounded-down 16 GiB portions (those display in the GUI as rounded decimal 17 GB chunks), you can create a 17 GB LUN. Navigate to **Pools** → **Volumes by Pools**, and click **Create Volumes** on the XIV GUI, as shown in Figure 31.

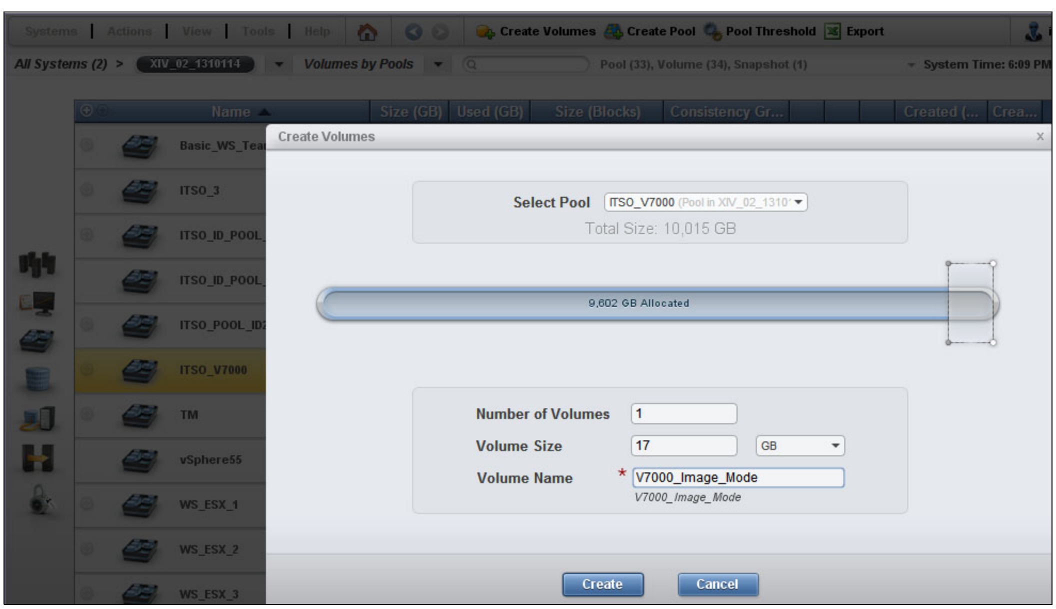


*Figure 31 Create a volume on the XIV*

Map it to the Storwize V7000 by right-clicking the **V7000_Image_Mode** volume and selecting **Map selected volumes**, as shown in Figure 32. In this example, the Storwize V7000 host is defined by the XIV as *ITSO_V7000*.

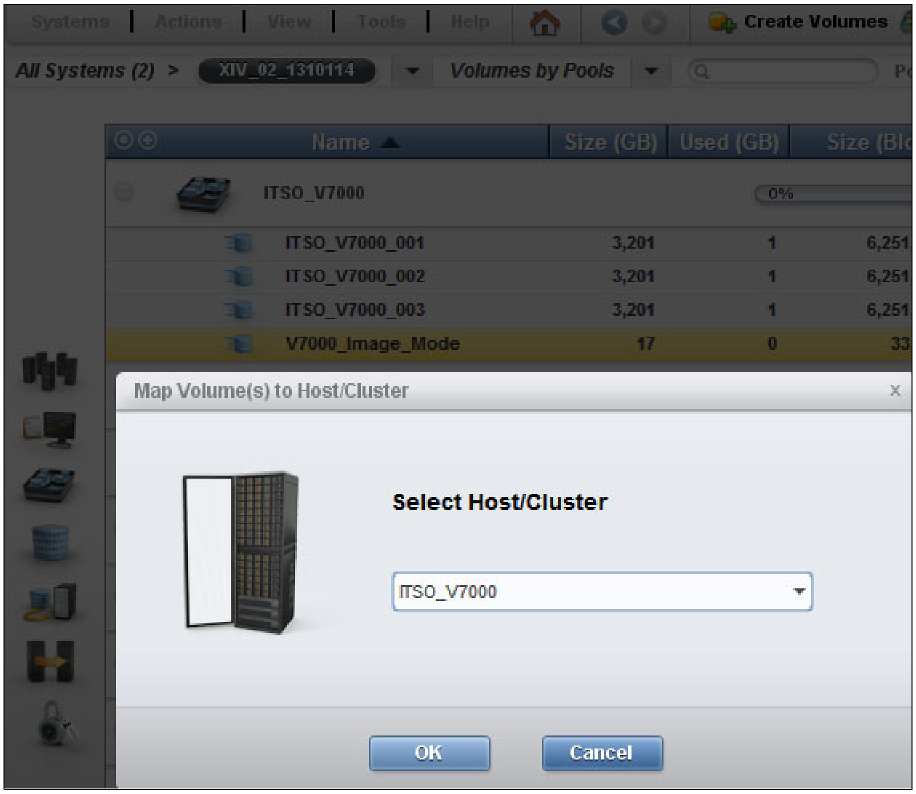


*Figure 32 Map volumes on the XIV*

Example 34 shows the corresponding XIV CLI commands.

*Example 34 Commands to create transitional volumes*

```
vol_create size=17 pool="ITSO_V7000" vol="V7000_Image_Mode"
map_vol cluster="ITSO_V7000" override=yes vol="V7000_Image_Mode" lun="4"
```

## Option 2. Create a volume that is exactly the same size

The drawback of using a larger volume size is that you use extra space. Therefore, it is better to create a volume that is the same size. To do this, you must know the size of the VDisk of Vol2, which is 10 GiB as shown in Figure 30 on page 38 and in Example 33 on page 38.

Now that you know the size of the source VDisk in bytes, you can divide this value by 512 to get the size in blocks (there are always 512 bytes in a standard SCSI block). So, 10,737,418,240 bytes divided by 512 bytes per block is 20,971,520 blocks. This is the size that you use on the XIV to create your Image mode transitional volume, as shown in Figure 33 on page 40.

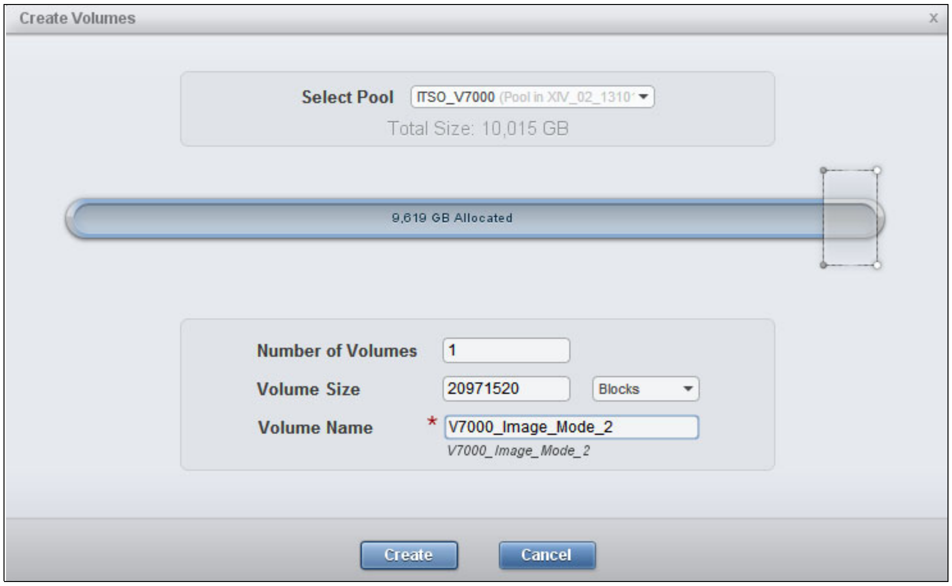

Figure 33 Create a volume in blocks on the XIV

Example 35 shows the XCLI command to create a volume by using blocks.

*Example 35 Command for creating an XIV volume by using blocks*

```
vol_create size_blocks=20971520 pool="ITSO_V7000" vol="V7000_Image_Mode_2"
```

After creating the volume, map the XIV to the Storwize V7000 (by using the XIV GUI or XCLI). Then, on the Storwize V7000, you can detect it as an unmanaged MDisk by using the **`svctask detectmdisk`** command or in the GUI under **Pools** → **MDisks by Pools** by clicking **Detect MDisks**, as shown in Figure 18 on page 28.

## Migrate the VDisk to Image mode

Now, migrate the source VDisk to Image mode by using the MDisk that you created for transition. These examples show an MDisk identified as mdisk4 that is 16.03 GiB (17 GB rounded down on the XIV GUI). This example also shows what eventually happens if you do not match sizes exactly.

Figure 34 shows the MDisks. They are not in a pool, so they are unmanaged.

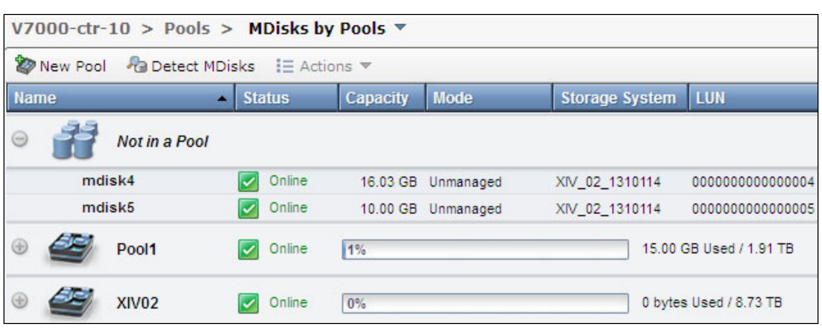

Figure 34 Unmanaged MDisks

In Example 36 on page 41, the unmanaged MDisks are listed by using the CLI.

*Example 36 Command for listing unmanaged MDisks*

```
IBM_2076:V7000-ctr-10:superuser>lsmdisk -filtervalue mode=unmanaged
id name   status mode      mdisk_grp_id mdisk_grp_name capacity ....
4  mdisk4 online unmanaged                             16.0GB ....
5  mdisk5 online unmanaged                             10.0GB ....
```

In Example 33 on page 38 and in Figure 30 on page 38, a source VDisk (Vol2) of 10 GiB and a target MDisk (mdisk4) of 16.03 GiB are identified.

Now, migrate the VDisk into Image mode without changing pools (stay in Pool1, which is where the source VDisk is located). The target MDisk must be unmanaged to be able to do this. If you migrate to a different MDisk group, the extent size of the target group must be the same as the source group. The advantage of using the same group is simplicity, but the disadvantage is that the MDisk group contains MDisks from two different controllers. That is not the best option for normal operations.

Navigate to **Pools** → **Volumes by Pool**, right-click **Vol2**, and select **Export to Image Mode**, as shown in Figure 35.

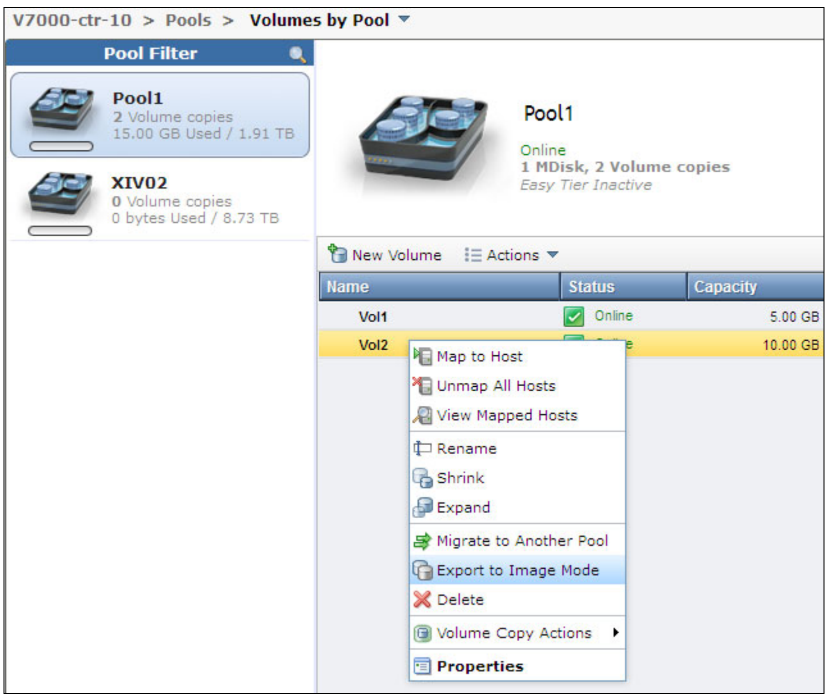


*Figure 35 Export to Image mode*

Select the MDisk (**mdisk4** in this example) target for Image mode migration, as shown in Figure 36 on page 42.

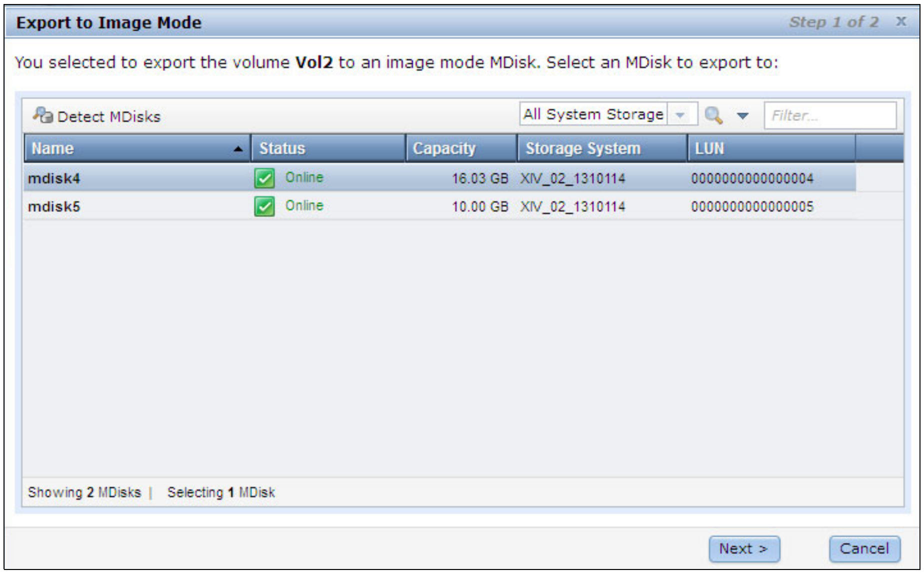


*Figure 36 Export to Image mode MDisk selection*

Select the target pool for the Image mode volume (**Pool1** in this example), as shown in Figure 37.

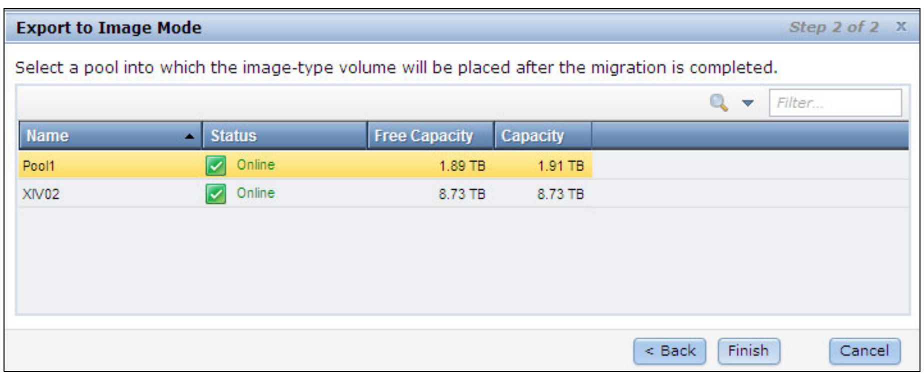


*Figure 37 Export to Image mode pool selection*

Example 37 shows the corresponding CLI command to start the migration.

*Example 37 Command for migrating a VDisk to Image mode*

```
svctask migratetoimage -mdisk mdisk4 -mdiskgrp Pool1 -vdisk Vol2
```

To monitor the migration and check for completion, click **Running Tasks**, as shown in Figure 24 on page 33 at the bottom of the GUI, and then select **Migration**. In the CLI, use the `lsmigrate` command shown in Example 38 (no response means that migration is complete).

*Example 38 Command for monitoring the migration*

```
IBM_2076:V7000-ctr-10:superuser>lsmigrate
IBM_2076:V7000-ctr-10:superuser>
```

**Important:** You must confirm that the VDisk is in Image mode or data loss will occur in the next step.

Right-click the volume, and select **Properties** to open the Properties window shown in Figure 38.

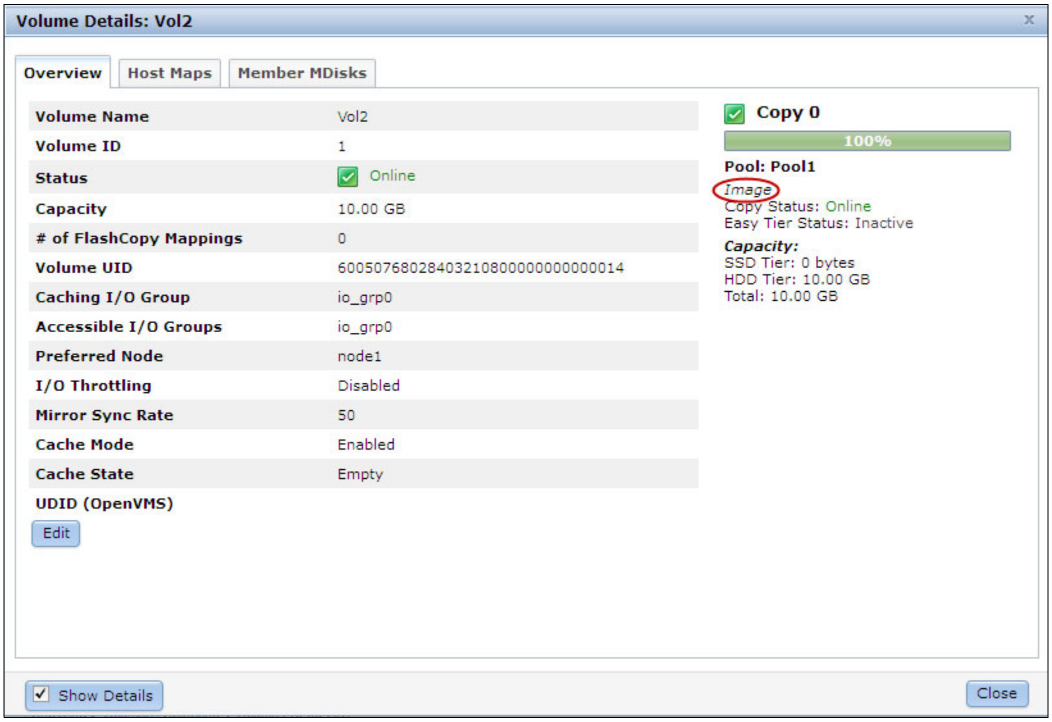


*Figure 38 Confirm that the volume is in Image mode*

Use the **lsvdisk** command in Example 39 to confirm that the volume is in Image mode.

*Example 39 Command for using the lsvdisk command to verify that the volume is in Image mode*

```
IBM_2076:V7000-ctr-10:superuser>lsvdisk
id name IO_group_id IO_group_name status mdisk_grp_id mdisk_grp_name capacity type
0  Vol1 0           io_grp0       online 0            Pool1          5.00GB   striped
1  Vol2 0           io_grp0       online 0            Pool1          10.00GB  image
```

At this point, a system downtime outage is required.

## Remove the Image mode VDisk

At the SVC or Storwize V7000, unmap the volume (which disrupts the host) and then remove the VDisk. At the host, you must have unmounted the volume (or shut down the host) to ensure that any data cached at the host has been flushed to the SVC. However, if write data in still in the cache for this VDisk at the SVC, you do not get an `Empty` message. You can check whether this is the case by displaying the `fast_write_state` for the VDisk by using an **svcinfo lsvdisk** command. You must wait for the data to flush from the cache, which might take several minutes.

To unmap the volume, right-click it and select **Unmap All Hosts**, as shown in Figure 39 on page 44.

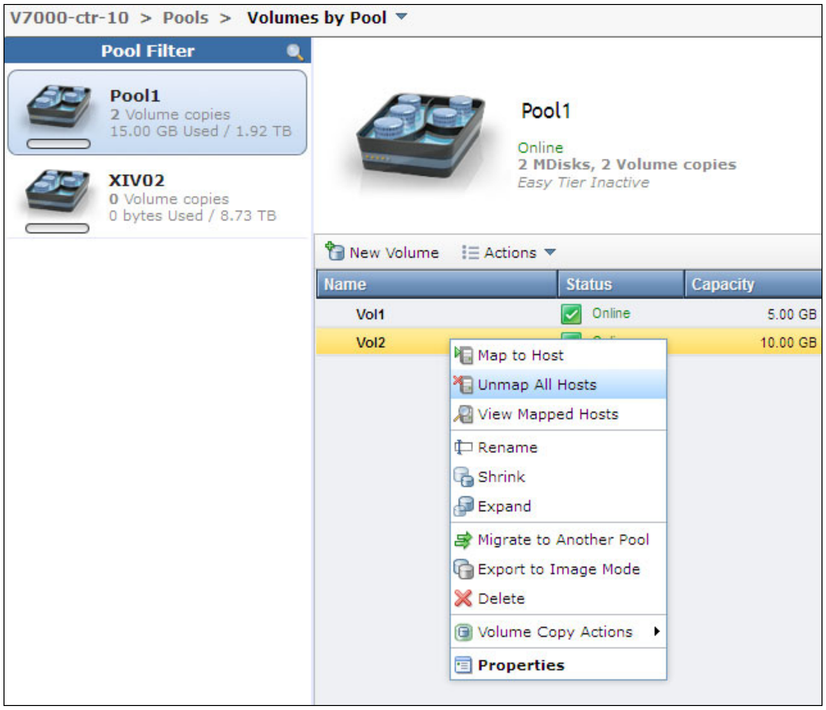


*Figure 39 Unmap All Hosts selected*

The corresponding CLI command is displayed in Example 40.

*Example 40 Command for unmapping the host*

```
rmvdiskhostmap -host Blade4 Vol2
```

To delete the volume, right-click it and select **Delete**, enter 1 for one volume, and click **Delete**, as shown in Figure 40.

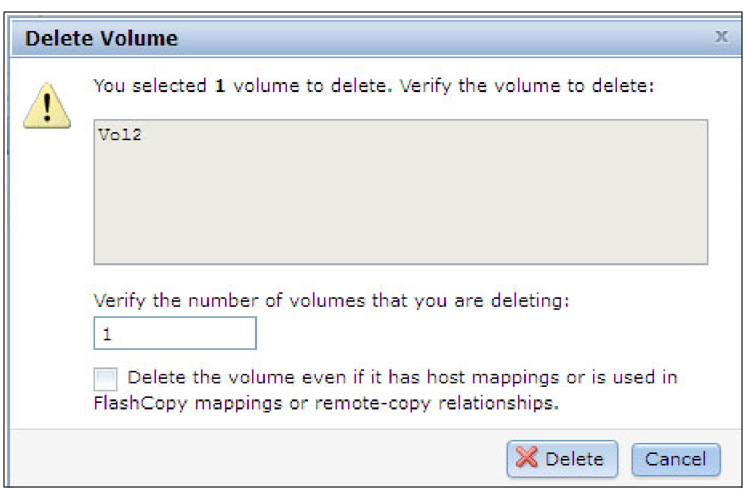


*Figure 40 Delete Volume dialog window*

Example 41 shows the corresponding CLI command.

*Example 41 Command for deleting a volume*

```
rmvdisk Vol2
```

The MDisk is now unmanaged (even though it contains data). From the XIV, remap that volume to a different SVC or Storwize V7000 cluster, or map the volume directly to a host (to convert that volume to a native XIV volume).

## Migration from Image mode to managed mode

Now, migrate the VDisks from Image mode on individual Image mode MDisks to striped mode VDisks in a managed mode MDisk group. (These instructions are based on the assumption that the volume is already unmapped from the host on the XIV and mapped to SVC.) Use the Image mode MDisk from “Migrate the VDisk to Image mode” on page 40 for the following examples.

First, create a new pool (MDisk group) by using volumes on the XIV that you intend to use as the destination. Detect the MDisks under **Pools** → **MDisks by Pools** by clicking **Detect MDisks**, as shown in Figure 11 on page 21. Four MDisks are displayed in Figure 41. MDisk4 is the Image mode MDisk. Create a pool for the other three MDisks (MDisk group).

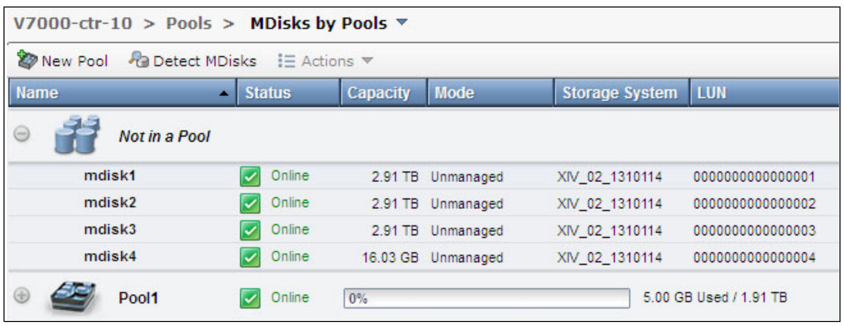


*Figure 41 Detect MDisks by pools*

Example 42 shows the accompanying CLI command.

*Example 42 Command for listing the free MDisks*

```
IBM_2076:V7000-ctr-10:superuser>lsmdisk -filtervalue mode=unmanaged
id name   status mode      mdisk_grp_id mdisk_grp_name capacity ....
1  mdisk1 online unmanaged                             2.9TB ....
2  mdisk2 online unmanaged                             2.9TB ....
3  mdisk3 online unmanaged                             2.9TB ....
4  mdisk4 online unmanaged                             16.0GB ....
```

Create an MDisk group by using an extent size of 1024 MB with the three free MDisks. In Example 43, an MDisk group with an ID of 1 (id [1]) is created. In SVC or Storwize V7000 firmware versions before 7.1, you can create an MDisk group with a different extent size than the default of 256 MB only through the CLI.

*Example 43 Command for creating the target MDisk group*

```
IBM_2076:V7000-ctr-10:superuser>mkmdiskgrp -name XIV02 -mdisk 1:2:3 -ext 1024
MDisk Group, id [1], successfully created
```

Now, import the Image mode MDisk, **mdisk4**, to an Image mode VDisk (volume) by right-clicking the MDisk and selecting **Import** (see Figure 42 on page 46).

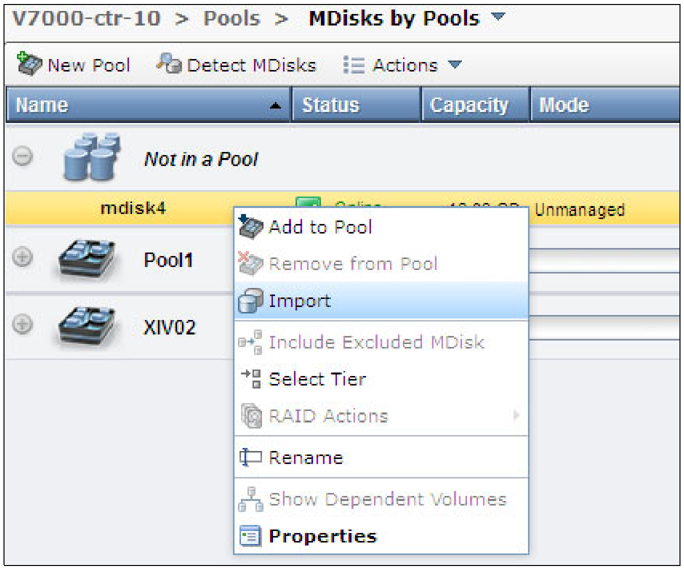


*Figure 42   Import selected to import the MDisk*

When the Import Wizard window shown in Figure 43 opens, click **Next**.

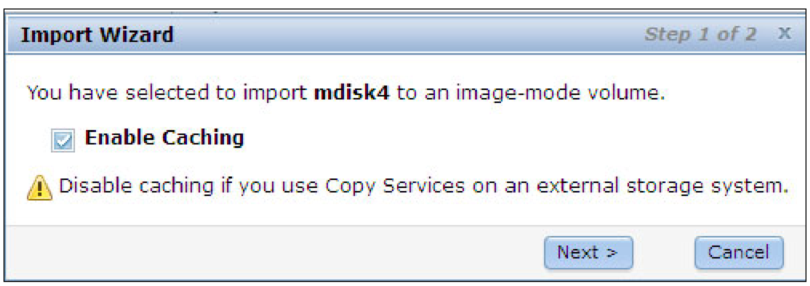


*Figure 43   Import Wizard, Step 1*

Select pool **XIV02**, and then click **Finish**, as shown in Figure 44.

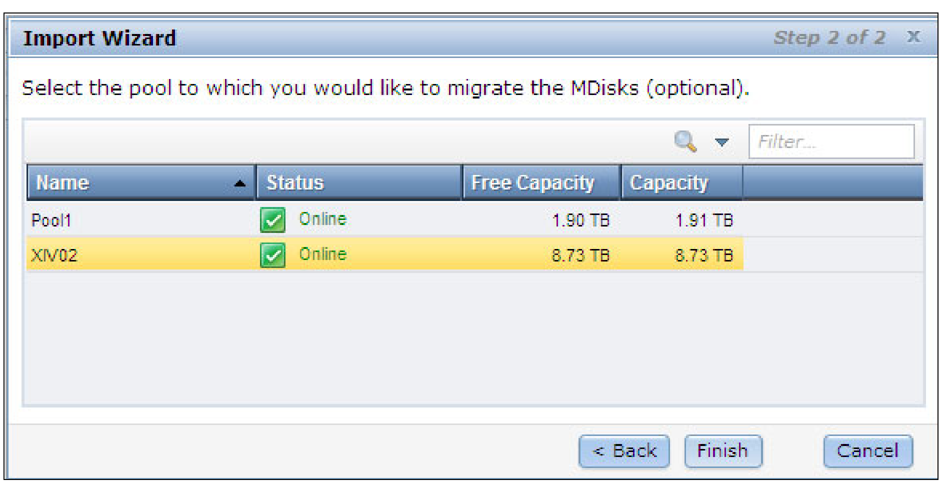


*Figure 44   Import Wizard, Step 2*

The corresponding CLI commands to import and migrate the VDisk are shown in Example 44 on page 47.

*Example 44   Commands to import and migrate the VDisk*

```
svctask mkvdisk -iogrp io_grp0 -mdisk mdisk4 -mdiskgrp MigrationPool_1024 -name
XIV_02_1310114_0000000000000004 -syncrate 80 -vtype image
svctask migratevdisk -mdiskgrp XIV02 -vdisk 1
```

For this example, we created an Image mode VDisk from the MDisk in the automatically created pool, MigrationPool_1024. In the second step, we migrated the VDisk to pool XIV02. Notice in Figure 45 that it is now 16.03 GB rather than 10 GB. This is because we initially migrated it to a 16 GB Image mode MDisk. It is better to use the 10 GB Image mode MDisk instead.


*Figure 45   Volumes by Pool window*

Use the **lsvdisk** command to confirm the size of the volume, as shown in Example 45.

*Example 45   Command for confirming VDisk space use*

```
IBM_2076:V7000-ctr-10:superuser>lsvdisk
id name .............................. mdisk_grp_name capacity ....
0  Vol1 .............................. Pool1          5.00GB ....
1  XIV_02_1310114_0000000000000004 .... XIV02          16.03GB ....
```

Under **Running Tasks**, check whether the migration has finished, as shown in Figure 24 on page 33 at the bottom of the GUI. Then, either select **Migration** or run the **lsmigrate** command (no response means that the migration is finished) shown in Example 38 on page 42.

## Remove transitional MDisk group

Now, under **Pools** → **MDisks by Pools**, you can delete MigrationPool_1024 by right-clicking **MigrationPool_1024** and selecting **Delete Pool** to open the Delete Pool window that is shown in Figure 46 on page 48. The deletion of MigrationPool_1024 also places mdisk4 in unmanaged mode.

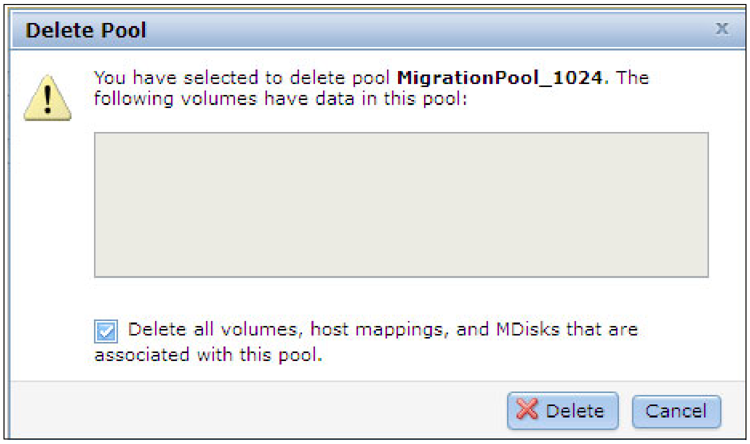


*Figure 46 Delete Pool window to delete the migration pool*

Example 46 shows the accompanying CLI command.

*Example 46 Command to delete the MDisk migration group*

```
svctask rmmdiskgrp -force 2
```

You can then unmap and delete the transition volume on the XIV to free the space so you can reuse that space for other migrations.

### Use transitional space as managed space

If all volumes are migrated from non-XIV disks to XIV disks, you can now use the space on the XIV that you reserved for the transitional Image mode MDisks to create new volumes to assign to the SVC. These volumes can be put into either the existing MDisk group or a new MDisk group.

### Remove non-XIV MDisks

The non-XIV disk controller's MDisks still exist. You can remove the MDisk group. Then, by using the non-XIV disk interface, you can unmap these LUNS from the SVC and reuse or remove the non-XIV disk controller.

## Future configuration tasks

This section documents other tasks that might be necessary after installation and migration are finished.

### Adding capacity to the XIV

When more capacity is added to a partially populated XIV, perform the following steps:

1. IBM adds the additional modules as a hardware upgrade (known as a miscellaneous equipment specification, or MES). The additional capacity appears as free space after the IBM service support representative completes the process to equip these modules.

**Note:** If the XIV has the Capacity on Demand (CoD) feature, no hardware change or license key is necessary to use available capacity that you have not purchased yet. You simply start using the additional capacity until all available usable space is allocated. The billing process to purchase this capacity occurs afterward.

2. From the Pools section of the XIV GUI, right-click the relevant pool and resize it, depending on how you plan to split the new capacity between any pools. If all of the space on the XIV is dedicated to a single SVC, there might be only one pool.
3. From the Volumes by Pools section of the XIV GUI, add new, identically sized volumes until you cannot create more volumes. (There is space left over that you can use as scratch space for testing and for non-SVC hosts.)
4. From the Host section of the XIV GUI, map these new volumes to the relevant SVC cluster. This completes the XIV portion of the upgrade.
5. From the SVC, detect and then add the new MDisks to the existing managed disk group. Alternatively, you can create a new managed disk group. Every MDisk uses a different XIV host port. Ideally, a new MDisk group contains several MDisks to spread the Fibre Channel traffic.
6. If new volumes are added to an existing managed disk group, “Why an extent rebalance might be desirable” explains when it might be desirable to rebalance the existing extents across the new space.

### Why an extent rebalance might be desirable

The SVC or Storwize V7000 before firmware Version 6.3 uses one XIV host port as a preferred port for each MDisk. If a VDisk is striped across eight MDisks, I/O from that VDisk is potentially striped across eight separate I/O ports on the XIV. If the space on these eight MDisks is fully allocated, when new capacity is added to the MDisk group, new VDisks are striped only across the new MDisks. If additional capacity that supplies only two new MDisks is added, I/O for VDisks striped across just those two MDisks is directed to only two host ports on the XIV. Therefore, the performance characteristics of these VDisks might be slightly different even though all XIV volumes have effectively the same back-end disk performance. You can find the script to do an extent rebalance in the IBM developerWorks® Community for Scripting Tools for SAN Volume Controller:

http://ibm.co/1c9hOGr

## SVC with XIV implementation checklist

Table 10 contains a checklist that you can use when you are implementing XIV behind SVC. It is based on the assumption that the XIV has been installed already by the IBM service support representative.

*Table 10 XIV implementation checklist*

| Task number | Completed | Where to perform | Task |
|---|---|---|---|
| 1 | | SVC | Increase SVC virtualization license if required. |
| 2 | | XIV | Get XIV WWPNs. |
| 3 | | SVC | Get SVC WWPNs. |
| 4 | | Fabric | Zone XIV to SVC (one big zone). |
| 5 | | XIV | Define the SVC cluster as a cluster. |
| 6 | | XIV | Define the SVC nodes as hosts. |
| 7 | | XIV | Add the SVC ports to the hosts. |
| 8 | | XIV | Create a storage pool. |
| 9 | | XIV | Create volumes in the pool. |
| 10 | | XIV | Map the volumes to the SVC cluster. |
| 11 | | XIV | Rename the XIV. |
| 12 | | SVC | Detect the MDisk. |
| 13 | | SVC | Rename the XIV controller. |
| 14 | | SVC | Rename the XIV MDisks. |
| 15 | | SVC | Create an MDisk group. |
| 16 | | SVC | Relocate the quorum disks if necessary. |
| 17 | | SVC | Identify VDisks to migrate. |
| 18 | | SVC | Mirror or migrate your data to XIV. |
| 19 | | SVC | Monitor migration. |
| 20 | | SVC | Remove non-XIV MDisks. |
| 21 | | SVC | Remove non-XIV MDisk group. |
| 22 | | Non-XIV storage | Unmap LUNs from SVC. |
| 23 | | SAN | Remove zone that connects SVC to non-XIV disk. |
| 24 | | SVC | Clear 1630 error that was generated by Task 23 (unzoning non-XIV disk from SVC). |

# Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

**Roger Eriksson** is an STG Lab Services consultant based in Stockholm, Sweden. He is a Senior Accredited IBM Product Service Professional. Roger has more than 20 years of experience working on IBM servers and storage, including enterprise and midrange disk, network-attached storage (NAS), storage area network (SAN), IBM System x®, POWER®, and BladeCenter®. Since 2008, he has consulted, produced proofs of concepts, and taught, mainly with the XIV product line, working with clients and various IBM teams. He holds a degree in Mechanical Engineering.

**Markus Oscheka** is an IT Specialist for proofs of concepts and benchmarks on the Disk Solution Europe team in Mainz, Germany. His areas of expertise include the setup and demonstration of IBM System Storage solutions in various environments, such as IBM AIX®, Linux, Microsoft Windows, VMware ESX, and Solaris. He has worked at IBM for 12 years. He has performed many proofs of concepts with Copy Services on IBM DS8000®, XIV, SVC, and V7000, as well as performance benchmarks with DS8000, XIV, SVC, V7000. Markus has contributed and acted as project co-lead for DS8000 and XIV Redbooks. He has spoken at several System Technical Universities. He holds a degree in Electrical Engineering from the Technical University in Darmstadt.

**Brian Sherman** is an Executive IT Specialist and IBM Distinguished Engineer with more than 25 years of experience as an IT Specialist since joining IBM from McMaster University in 1985 with a Mathematics and Computer Science degree. Brian has held various storage-related roles since then, including Level 2 software support, storage implementation services, and branch systems engineer in the public and financial sectors. He currently provides storage hardware and software technical sales support and competitive products across the Americas and Growth Markets Unit as part of the Advanced Technical Skills organization. Brian also develops and provides worldwide technical education for new storage hardware and software product launches and participates in several storage product development teams.

**Stephen Solewin** is an IBM XIV Corporate Solutions Architect in Tucson, Arizona. He has 16 years of experience in working on IBM storage, including enterprise and midrange disk, Linear Tape-Open (LTO) drives and libraries, SAN, storage virtualization, and storage software. Steve has been working on the XIV product line since March of 2008. He holds a BS degree in Electrical Engineering from the University of Arizona, where he graduated with honors.

# Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

**ibm.com**/redbooks/residencies.html

## Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document REDP-5063-00 was created or updated on January 8, 2014.

Send us your comments in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:
  **ibm.com**/redbooks
- Send your comments in an email to:
  redbooks@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

IBM®

Redpaper™

# Trademarks