Redbooks
ibm.com/redbooks

IBM FlashSystem A9000 and A9000R Replication Solutions

Markus Oscheka
Roger Eriksson
Detlef Helmbrecht
Bruce Spell

Storage

IBM

Redpaper

International Technical Support Organization

IBM FlashSystem A9000 and A9000R Replication Solutions

January 2017

REDP-5401-00

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (January 2017)

This edition applies to the IBM FlashSystem A9000 and IBM FlashSystem A9000R, with System software Version 12.0.2.b and Hyper-Scale Manager Version 5.1.

This document was created or updated on January 27, 2017.

Contents

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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. The sample programs are provided “AS IS”, without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®
IBM®
IBM FlashSystem®
IBM Spectrum™
IBM Spectrum Accelerate™
IBM Spectrum Storage™
IBM Spectrum Virtualize™
Jazz™
Redbooks®
Redpaper™
Redbooks (logo)  ®
System Storage®
XIV®

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

Preface

IBM® FlashSystem A9000 and IBM FlashSystem® A9000R provide copy functions suited for various data protection scenarios that enable you to enhance your business continuance, disaster recovery, data migration, and backup solutions. These functions allow point-in-time copies, known as snapshots, and also include remote copy capabilities in either synchronous or asynchronous mode.

Furthermore, support for IBM Hyper-Scale Mobility enables a seamless migration of volumes from one IBM FlashSystem A9000 to another with no interference to the host.

This IBM Redpaper™ publication is intended for anyone who needs a detailed and practical understanding of the IBM FlashSystem A9000 and IBM FlashSystem A9000R replication functions.

This publication describes replication only for the IBM FlashSystem A9000 and A9000R; it does not describe replication for the other members of the IBM Spectrum™ Accelerate family.

Authors

This paper was produced by a team of specialists from around the world.

Markus Oscheka is an IT Specialist for Proof of Concepts and Benchmarks in the Disk Solution Europe team in Germany. He has worked at IBM for fourteen years. He has performed many Proof of Concepts with Copy Services on IBM Spectrum Virtualize™ and IBM Spectrum Storage™, and Performance-Benchmarks with IBM Spectrum Virtualize and IBM Spectrum Storage. He has written extensively and acted as project lead for various IBM Redbooks® publications. He has also spoken on several System Technical Universities. He holds a degree in Electrical Engineering from the Technical University in Darmstadt.

Roger Eriksson works at IBM Systems Lab Services Nordic, based in Stockholm, Sweden. He is a Senior Accredited IBM Product Service Professional. Roger has over 20 years of experience working on IBM servers and storage. He has done consulting, proof of concepts, and education, mainly with the IBM Spectrum Accelerate™ product line, since December 2008. Working with both clients and various IBM teams worldwide is a normal day's work for Roger. He holds a Technical College Graduation in Mechanical Engineering.

Detlef Helmbrecht is an Advanced Technical Skills (ATS) IT Specialist working for IBM Systems at the European Storage Competence Center (ESCC), Germany. Detlef has over 30 years of experience in IT performing numerous roles, including software design, sales, and solution architect. He joined IBM in 2013. His current areas of expertise include high-performance computing (HPC), application and database tuning, and IBM products, such as IBM Spectrum Virtualize, IBM Spectrum Accelerate, and the IBM FlashSystem family. Detlef holds a degree in Mathematics from the Ruhr-Universität Bochum in Germany.

Bruce Spell is a Storage Technical Advisor located in Tucson, Arizona. As a Technical Advisor he delivers proactive and comprehensive consulting services for IBM storage products to ensure the highest availability for vital IT solutions. He is a customer advocate providing focus for problem management, change management, crisis management, and best practices. Bruce has over 32 years of experience and a broad background spanning across development, test, field support, and program management. He has been a Technical Advisor since 2007 and recognized as a Notable Advocate over the last several years. Bruce holds a BS degree in Electrical Engineering from UCLA.

Bert Dufrasne is an IBM Certified Consulting IT Specialist and Project Leader for IBM System Storage® disk products at the International Technical Support Organization (ITSO), San Jose Center. Bert has worked at IBM in various IT areas. He has authored many IBM Redbooks publications, and he has also developed and taught technical workshops. Before Bert joined the ITSO, he worked for IBM Global Services as an Application Architect. He holds a Master's degree in Electrical Engineering.

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

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications 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

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

1

Remote connectivity with IBM FlashSystem A9000 and IBM FlashSystem A9000R

Remote system connectivity is the connection between two systems for mirroring or data migration. This chapter provides the information required for defining and managing remote target connectivity using IBM Hyper-Scale Manager and the command-line interface (XCLI).

This chapter includes the following sections:

- IBM FlashSystem A9000 and A9000R remote system actions
- Planning
- Using the Hyper-Scale Manager or XCLI for remote mirror

1.1 IBM FlashSystem A9000 and A9000R remote system actions

Remote system actions described in this section are the fundamental building blocks of IBM FlashSystem A9000 and A9000R remote mirroring solutions and usage scenarios.

1.1.1 Defining the mirroring target

To connect two IBM FlashSystem A9000 and A9000R systems for remote mirroring, each system must be defined to be a mirroring target of the other. This connection and definition allows for either IBM FlashSystem A9000 or A9000R to become the source, which is needed for switching roles. A mirroring target is a IBM FlashSystem A9000 or A9000R with volumes that receive data copied through IBM FlashSystem remote mirroring.

Defining a IBM FlashSystem A9000 or A9000R mirroring target for a IBM FlashSystem A9000 or A9000R involves giving the target a name and specifying whether Fibre Channel (FC) or iSCSI protocol is used to copy the data. For a practical illustration, see 1.3.2, “Remote mirror target configuration through Hyper-Scale Manager” on page 16.

Remote mirroring copies data from a volume on one IBM FlashSystem A9000 or A9000R to a peer on another IBM FlashSystem A9000 or A9000R (that is, the mirroring target system). Whereas the basic underlying mirroring relationship is a one-to-one relationship between two peers, IBM FlashSystem A9000 and A9000R systems can be connected in several ways:

- IBM FlashSystem A9000 and A9000R target configuration: One-to-one

  The most typical IBM FlashSystem A9000 and A9000R remote mirroring configuration is a one-to-one relationship between a local system (production system) and a remote system (DR system), as shown in Figure 1-1. This configuration is typical where there is a single production site and a single disaster recovery (DR) site.

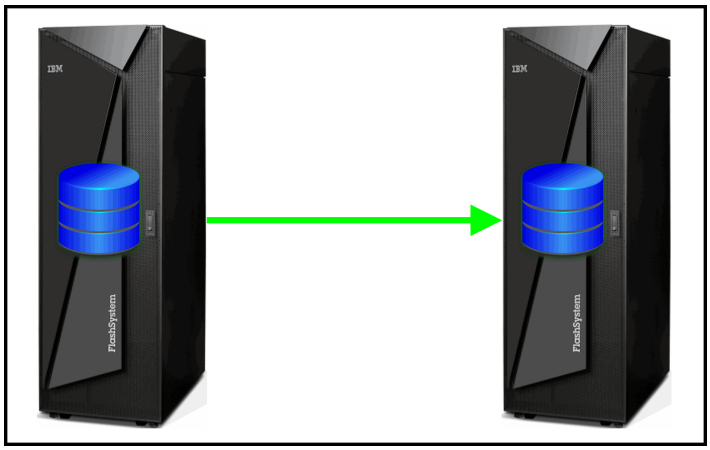

Figure 1-1 One-to-one target configuration

During normal remote mirroring operation, one IBM FlashSystem A9000 or A9000R (at the DR site) is active as a mirroring target. The other IBM FlashSystem A9000 or A9000R (at the local production site) is active as a mirroring target only when it becomes available again after an outage and change of roles between the production and the DR site. Data changes made when production is running on the remote (DR) site are copied back to the original production site, as shown in Figure 1-2 on page 3.

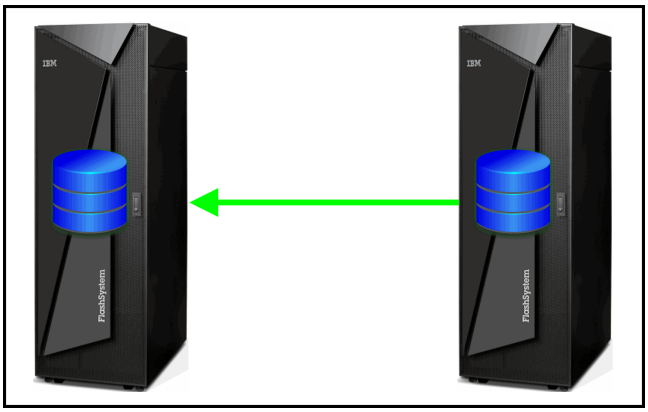

Figure 1-2 Copying changes back to production

In a configuration with two identically provisioned sites, production might be periodically switched from one site to another as part of normal operation. The IBM FlashSystem A9000 or A9000R system that is the active mirroring target is switched at the same time. The **`mirror_switch_roles`** command allows for switching roles in both synchronous and asynchronous mirroring. There are special requirements for doing so with asynchronous mirroring.

- IBM FlashSystem A9000 and A9000R target configuration: Synchronous and asynchronous one-to-one

  IBM FlashSystem A9000 and A9000R support both synchronous and asynchronous mirroring (for different mirror couplings) on the same IBM FlashSystem A9000 or A9000R. Thus, a single local IBM FlashSystem can have certain volumes synchronously mirrored to a remote IBM FlashSystem; whereas other volumes are asynchronously mirrored to the same remote IBM FlashSystem A9000 or A9000R, as shown in Figure 1-3. Highly response-time-sensitive volumes can be asynchronously mirrored and less response-time-sensitive volumes can be synchronously mirrored to a single remote IBM FlashSystem A9000 or A9000R.

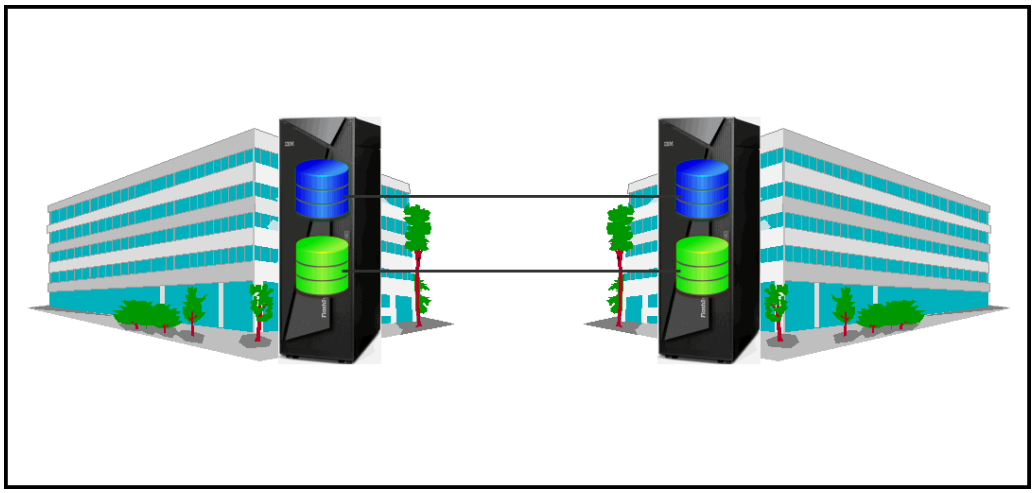

Figure 1-3 Synchronous and asynchronous peers

- IBM FlashSystem A9000 and A9000R target configuration: Fan-out

  A single local (production) IBM FlashSystem A9000 or A9000R can be connected to two remote (DR) IBM FlashSystem A9000s or A9000Rs in a fan-out configuration, as shown in Figure 1-4. Both remote systems can be at the same location, or each of the targets can be at a different location. Certain volumes on the local system are copied to one remote system, and other volumes on the same local IBM FlashSystem A9000 or A9000R are copied to a different remote IBM FlashSystem A9000 or A9000R. This configuration can be used when each system at the DR site has less available capacity than the system at the local site.

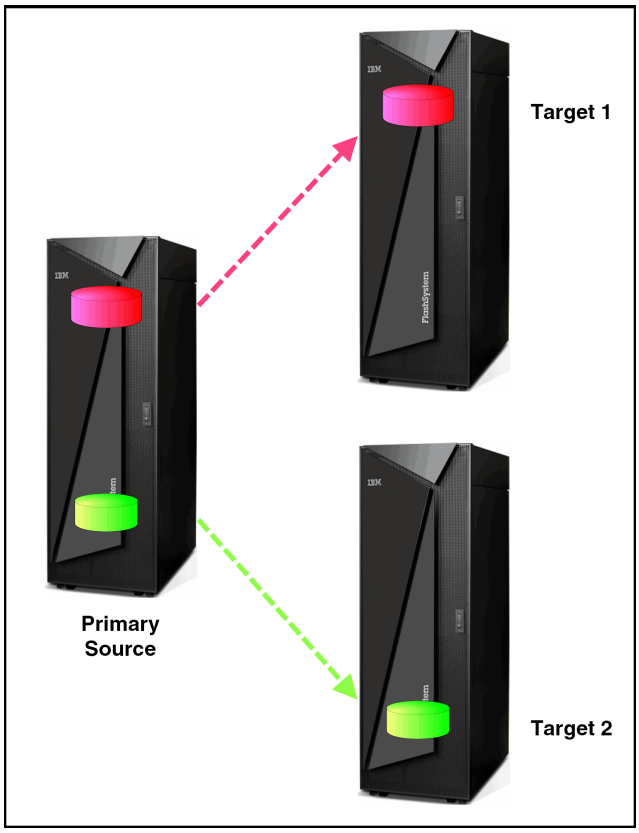

Figure 1-4 Fan-out target configuration

- IBM FlashSystem A9000 and A9000R target configuration: Synchronous and asynchronous fan-out

  IBM FlashSystem A9000 and A9000R support both synchronous and asynchronous mirroring (for different mirror couplings) on the same IBM FlashSystem A9000 or A9000R. Thus, a single local IBM FlashSystem A9000 or A9000R can have certain volumes synchronously mirrored to a remote IBM FlashSystem A9000 or A9000R at a metro distance; whereas other volumes are asynchronously mirrored to a remote system at a global distance, as shown in Figure 1-5 on page 5. This configuration can be used when higher priority data is synchronously mirrored to another system within the metro area, and lower priority data is asynchronously mirrored to a system within or outside the metro area.

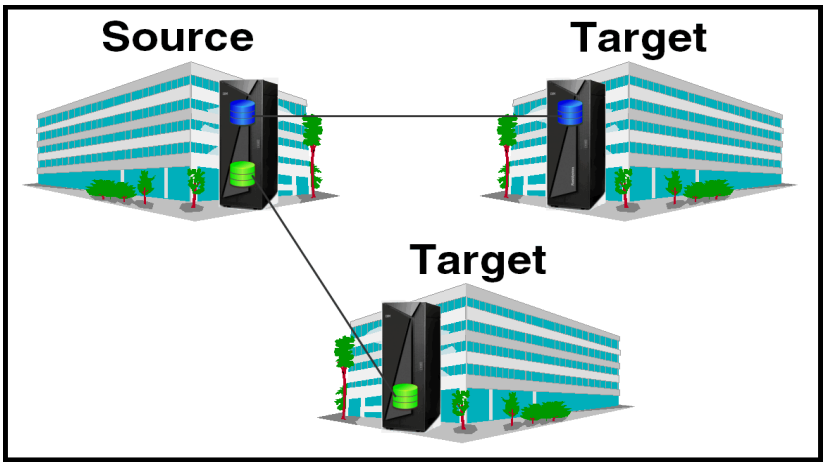

Figure 1-5 Synchronous and asynchronous fan-out

- IBM FlashSystem A9000s or A9000R target configuration: Fan-in

  Two, or more, local IBM FlashSystem A9000s or A9000Rs can have peers mirrored to a single remote IBM FlashSystem A9000 or A9000R in a fan-in configuration, as shown in Figure 1-6. This configuration must be evaluated carefully and used with caution because there is a risk of overloading the single remote system, affecting performance.

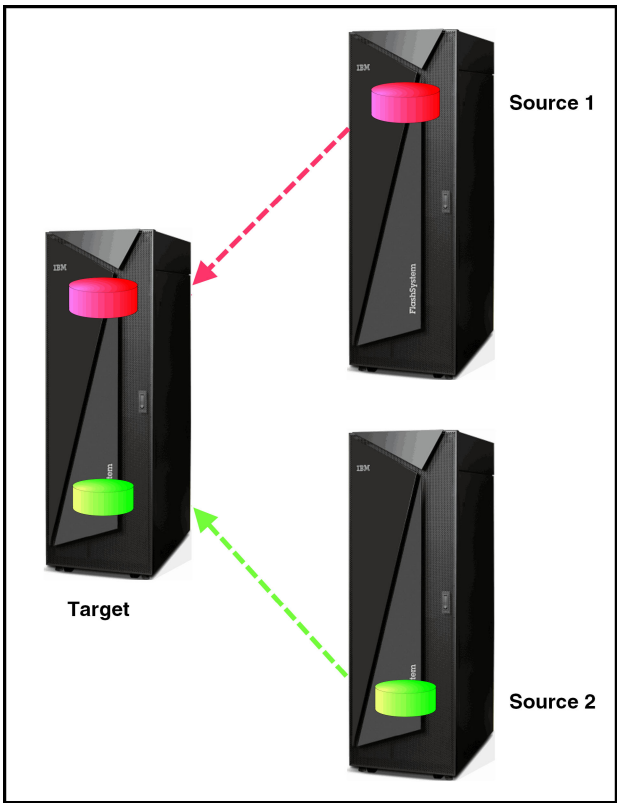

Figure 1-6 Fan-in configuration

This configuration can be used in situations where there is a single DR data center supporting multiple production data centers. It can also be used when multiple IBM FlashSystem A9000s or A9000Rs are mirrored to a single IBM FlashSystem A9000 or A9000R at a service provider.

- IBM FlashSystem A9000 and A9000R target configuration: Bidirectional

  Two different IBM FlashSystem A9000 or A9000Rs can have different volumes mirrored in a bidirectional configuration, as shown in Figure 1-7. This configuration can be used for situations where there are two active production sites and each site provides a DR solution for the other. Each IBM FlashSystem A9000 or A9000R is active as a production system for certain peers and as a mirroring target for other peers.

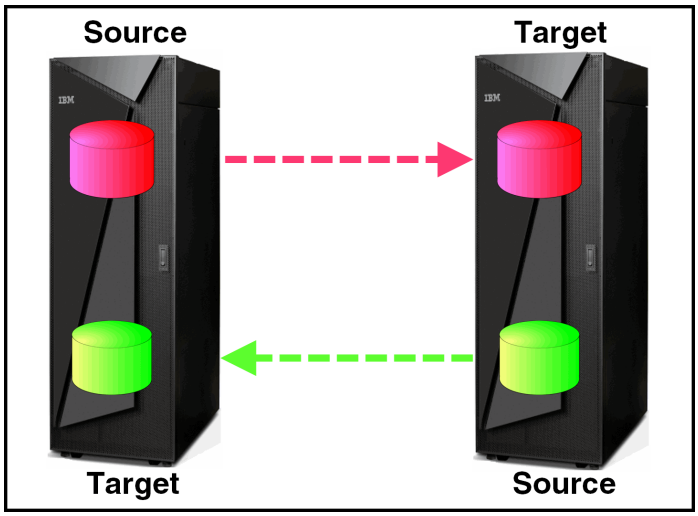

Figure 1-7 Bidirectional configuration

1.1.2 Connecting mirroring ports

After defining remote mirroring targets, one-to-one connections must be made between ports on each IBM FlashSystem A9000 or A9000R. For a description of these actions using the Hyper-Scale Manager or the XCLI, see 1.3, “Using the Hyper-Scale Manager or XCLI for remote mirror” on page 9.

- FC ports

  For the IBM FlashSystem A9000 and A9000R FC ports, connections are unidirectional. An example is a connection from an initiator port (Grid Controller Port 4 configured as an FC initiator by default) on the source IBM FlashSystem to a target port (typically Grid Controller Port 2) on the target IBM FlashSystem. Use a minimum of four connections (two connections in each direction, from ports in different grid controllers, using a total of eight ports) to provide availability protection. See Figure 1-8 on page 7.

  In Figure 1-8, the solid lines represent mirroring connections that are used during normal operation (the mirroring target system is on the right). The dotted lines represent mirroring connections that are used when production is running at the disaster recovery site and changes are being copied back to the original production site (the mirroring target is on the left). IBM FlashSystem A9000 and A9000R FC ports can be easily and dynamically configured as initiator or target ports.

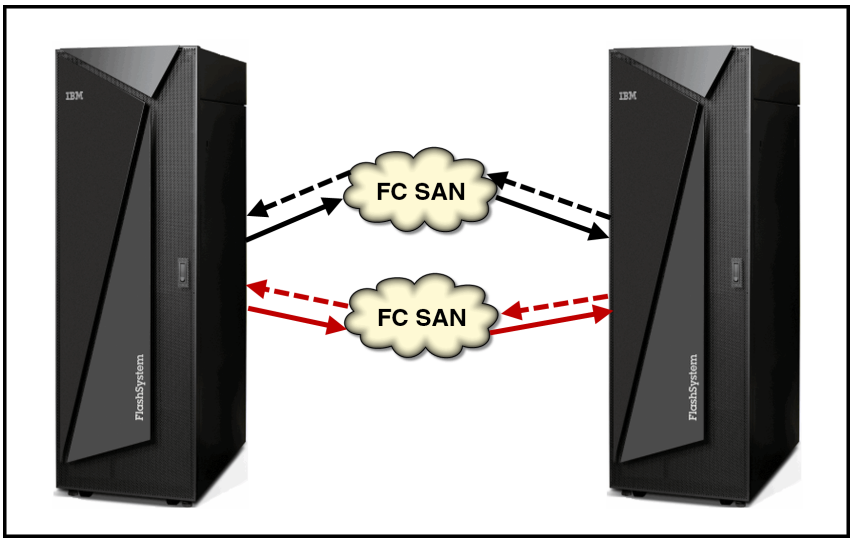

Figure 1-8 Connecting A9000 and A9000R mirroring ports (FC connections)

Note: Each Grid Controller in the IBM FlashSystem A9000 and A9000R is configured with interface ports. Figure 1-9 shows the rear view of a Grid Controller with feature code 5001 (4 x 16 Gb FC ports + 2 x 10 GbE iSCSI ports. All Grid Controllers in a IBM FlashSystem A9000 or A9000R have to be the same feature code.

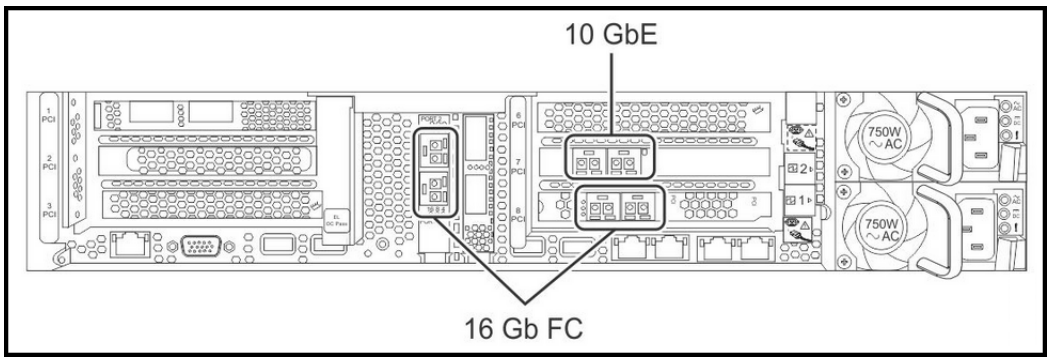

Figure 1-9 Grid Controller with feature code 5001

- iSCSI ports

  For iSCSI ports, connections are bidirectional.

Important: If the IP network includes firewalls between the mirrored IBM FlashSystem A9000s and A9000Rs, TCP port 3260 (iSCSI) must be open within firewalls so that iSCSI replication can work.

Use a minimum of two connections, with each of these ports in a different grid controller, using a total of four ports to provide redundancy. In Figure 1-10 on page 8, during normal operation, the data flow starts from the production system (on the left) and goes towards the mirroring target system (on the right). The data flow is reversed when production is running at the disaster recovery site (on the right) and changes are being copied back to the original production site (the mirroring target is on the left).

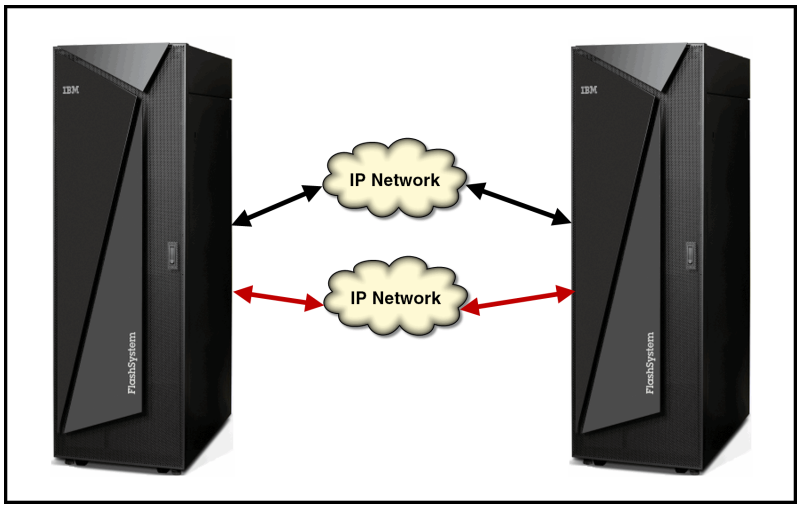

Figure 1-10 Connecting A9000 and A9000R mirroring ports (Ethernet connections)

Note: For asynchronous mirroring over iSCSI links, a reliable, dedicated network must be available. It requires consistent network bandwidth and a non-shared link.

Figure 1-11 shows the rear view of a Grid Controller with feature code 5002 (4 x 10 GbE iSCSI ports). All Grid Controllers in a IBM FlashSystem A9000 or A9000R must be the same feature code.

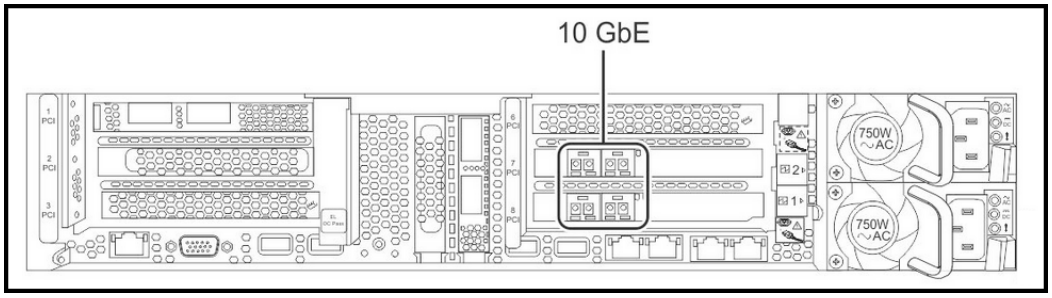

Figure 1-11 Grid Controller with feature code 5002

1.2 Planning

The most important planning considerations for IBM FlashSystem A9000 and A9000R remote mirroring are those related to ensuring availability and performance of the mirroring or online mobility connections between the IBM FlashSystems. Planning for snapshot capacity usage is also important.

To optimize availability, IBM FlashSystem A9000 and A9000R remote mirroring connections must be spread across multiple ports, on different system cards, in different grid controllers, and must be connected to different networks. Minimum network bandwidth requirements must be maintained to ensure a stable environment. Adequate bandwidth must also be allocated to ensure that the anticipated amount of changed data can be transported across the network between IBM FlashSystem A9000s and A9000Rs within the wanted RPO.

Important: Network bandwidths are typically expressed in megabits per second (Mbps), and disk array bandwidths are expressed in megabytes per second (MBps). Although not exact, a factor of eight between the two gives an acceptable approximation.

1.3 Using the Hyper-Scale Manager or XCLI for remote mirror

This section illustrates remote-mirroring definition actions through the Hyper-Scale Manager and the XCLI.

1.3.1 Initial setup

When preparing to set up remote mirroring, consider the following questions:

- Will the paths be configured for FC or for iSCSI?
- Is the port that you want to use configured as an initiator or a target?
  - Port 4 default configuration is initiator.
  - Port 2 is suggested as the target port for remote mirror links.
  - Port designations can be changed if needed.
- How many mirror pairs will be established?

  The answer is related to the bandwidth needed between sites.
- How many secondary systems will be used for a single primary?

  Remote mirroring can be set up on paths that are SAN-attached FC or iSCSI protocols.

  For most disaster recovery solutions, the secondary system is at a geographically remote site.

  The sites are connected using either SAN connectivity with Fibre Channel Protocol (FCP) or Ethernet with iSCSI.

Reminder: If the IP network includes firewalls between the mirrored IBM FlashSystem A9000s and A9000Rs, TCP port 3260 must be open within firewalls so that iSCSI replication can work.

Bandwidth considerations must be taken into account when planning the infrastructure to support the remote mirroring implementation. Knowing when the peak write rate occurs for systems attached to the storage helps with the planning for the number of paths needed to support the remote mirroring function and any future growth plans.

When the protocol is selected, it is time to determine which ports on the IBM FlashSystem A9000 or A9000R are used. The port settings are easily displayed using the XCLI session environment and the **fc_port_list** command for FC or the **ipinterface_list** command for iSCSI.

There must always be a minimum of two paths configured within remote mirroring for FC connections, and these paths must be dedicated to remote mirroring. These two paths also must be considered a set. Generally, FC port 4 and FC port 2 in the desired grid controller are used for this purpose. For redundancy, extra sets of paths are recommended to be configured in different grid controllers.

FC paths for remote mirroring have slightly more requirements for setup (versus iSCSI), which is the method that is described here first.

As Example 1-1 shows in the Role column, each FC port is identified as a target or an initiator. A target in a remote mirror configuration is the port that receives data from the other system, whereas an initiator is the port that sends the data. In this example, there are five initiators configured. Initiators, by default, are configured on 1:FC_Port:X:4 (X is the grid controller module number). In this example, port 4 on all four grid controller modules is configured as an initiator, as is port 3 in grid controller module 4.

Example 1-1 The fc_port_list output command

```
A9000R>>fc_port_list
Component ID    Status   Currently Functioning   WWPN              Port ID    Role
1:FC_Port:1:1   OK       yes                     5001738051A60110  007EE680   Target
1:FC_Port:1:2   OK       yes                     5001738051A60111  00011000   Target
1:FC_Port:1:3   OK       yes                     5001738051A60112  00011400   Target
1:FC_Port:1:4   OK       yes                     5001738051A60113  00042400   Initiator
1:FC_Port:2:1   OK       yes                     5001738051A60120  007FC400   Target
1:FC_Port:2:2   OK       yes                     5001738051A60121  00011100   Target
1:FC_Port:2:3   OK       yes                     5001738051A60122  00011500   Target
1:FC_Port:2:4   OK       yes                     5001738051A60123  00042500   Initiator
1:FC_Port:3:1   OK       yes                     5001738051A60130  00042200   Target
1:FC_Port:3:2   OK       yes                     5001738051A60131  00011200   Target
1:FC_Port:3:3   OK       yes                     5001738051A60132  007FC440   Target
1:FC_Port:3:4   OK       yes                     5001738051A60133  00011600   Initiator
1:FC_Port:4:1   OK       yes                     5001738051A60140  00042300   Target
1:FC_Port:4:2   OK       yes                     5001738051A60141  00011300   Target
1:FC_Port:4:3   OK       yes                     5001738051A60142  007EE6C0   Initiator
1:FC_Port:4:4   OK       yes                     5001738051A60143  00011700   Initiator
```

The iSCSI connections that are shown in Example 1-2 use the **ipinterface_list** command. The output is truncated to show only the iSCSI connections that are of interest here; the command also displays all other Ethernet connections and settings. In this example, four connections are displayed for iSCSI: one connection in each grid controller module.

Example 1-2 The ipinterface_list command

```
A9000R>>ipinterface_list
Name          Type         IP Address      Network Mask    Default Gateway MTU    Module       Port
iSCSI_M1_P1   iSCSI        9.155.116.205   255.255.240.0   9.155.112.1     9000   1:Module:1   1
iSCSI_M2_P1   iSCSI        9.155.116.206   255.255.240.0   9.155.112.1     9000   1:Module:2   1
iSCSI_M3_P1   iSCSI        9.155.116.207   255.255.240.0   9.155.112.1     9000   1:Module:3   1
iSCSI_M4_P1   iSCSI        9.155.116.208   255.255.240.0   9.155.112.1     9000   1:Module:4   1
```

Alternatively, you can query for the existing port settings among the managed IBM FlashSystem A9000s and A9000Rs by selecting a system in the Hyper-Scale Manager, followed by selecting the **P.Panel** spoke in the Hub view on the right (Figure 1-12).

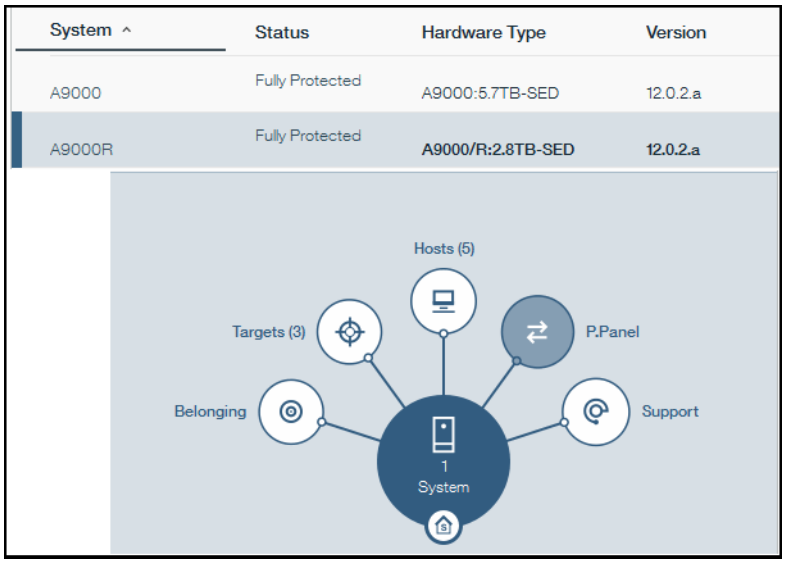

Figure 1-12 Selecting P.Panel spoke from the Hyper-Scale Manager Hub View

After selecting the **P.Panel** spoke, the ports are displayed below the Hub. Figure 1-13 on page 12 shows some of the ports for a IBM FlashSystem A9000R (rack system).

Figure 1-13 Port listing from the P.Panel spoke in the Hub view (A9000R)

Click the action menu for a specific port and select View/Edit; the output is shown in Figure 1-14 on page 13. This FC port is configured as a target, as indicated by the radio button on the right.

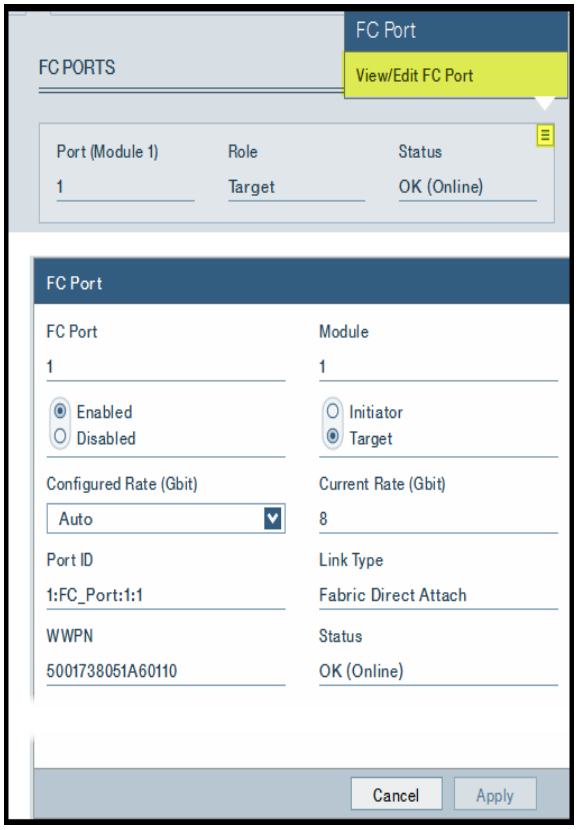

Figure 1-14 FC Port properties displayed with Hyper-Scale Manager

Similar information can be displayed for the iSCSI connections using the Hyper-Scale Manager, as shown in Figure 1-15 on page 14.

This view can be seen by clicking the action menu for an IP Interface and selecting **View/Update IP Interface** (similar to the FC port shown in Figure 1-14). This sequence displays the same iSCSI definitions that are shown with the **`XCLI`** command.

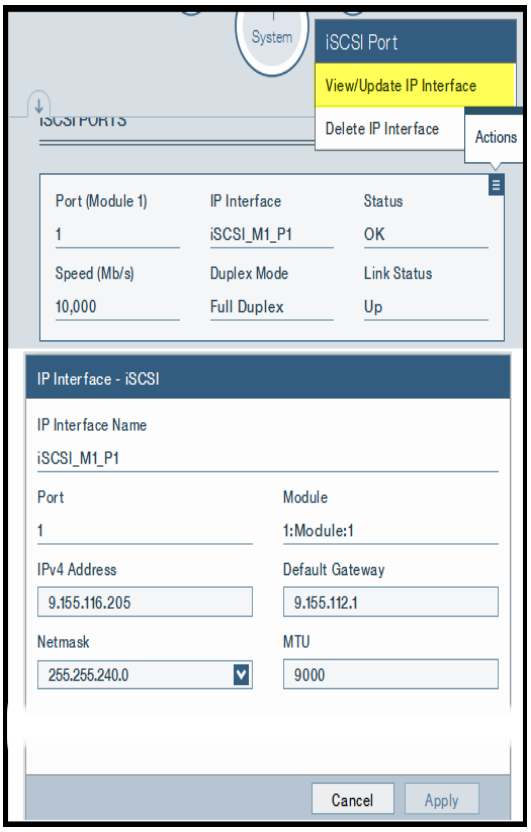

Figure 1-15 iSCSI Port properties displayed with Hyper-Scale Manager

Most commonly, Fibre Channel ports 2 and 4 (target and initiator) from every grid controller are used for remote mirroring. For example, port 4 in grid controller 3 (initiator) on the local system is connected to port 2 in grid controller 3 (target) on the remote system.

When setting up a new system, plan for any remote mirroring and reserve these ports for that purpose. However, different ports can be used if desired or as needed.

If a port role does need changing, you can change the port role with either the XCLI or the Hyper-Scale Manager.

Use the XCLI **`fc_port_config`** command to change a port, as shown in Example 1-3. Using the output from fc_port_list, you can get the fc_port name to be used in the command, changing the port role to be either initiator or target, as needed.

Example 1-3 XCLI command to configure a port

```
A9000R>>fc_port_config fc_port=1:FC_Port:4:3 role=initiator
Command executed successfully.
A9000R>>fc_port_list fcport=1:FC_Port:4:3
Component ID    Status   Currently Functioning   WWPN               Port ID    Role
1:FC_Port:4:3   OK       yes                     5001738051A60142   007EE6C0   Initiator
```

To perform the same function with the Hyper-Scale Manager, select the desired FC port action menu as shown in Figure 1-14 on page 13, then click **View/Edit FC Port**, which displays the FC port properties as shown in Figure 1-16.

The settings are displayed and the port can be enabled (or disabled), its role defined as target or initiator, and the speed configured. The options are Auto, 2, 4, 8, and 16 Gbps.

After any changes, click **Apply**.

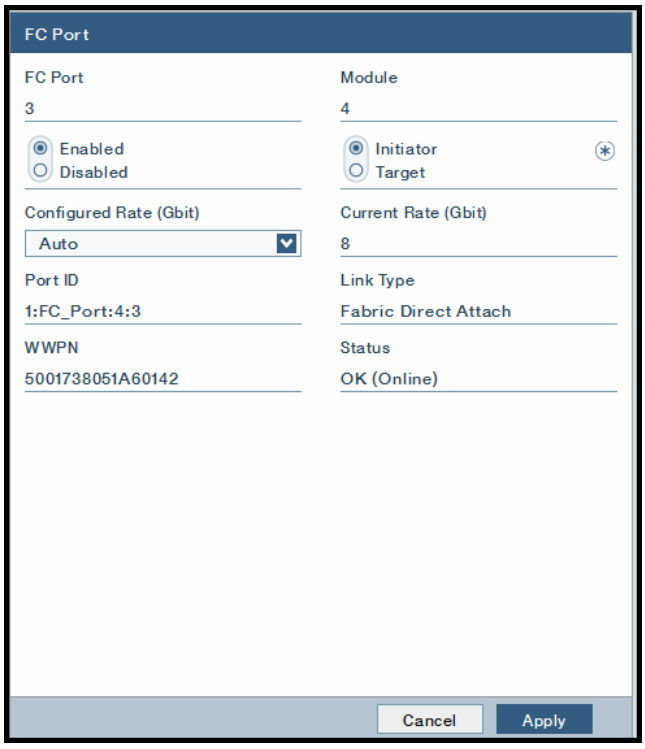

Figure 1-16 Configure FC port with the Hyper-Scale Manager

Planning for remote mirroring is important when determining how many mirror pairs will exist. All volumes defined in the system can be mirrored. A single primary system is limited to a maximum of 10 secondary systems.

Volumes cannot be part of a IBM FlashSystem A9000 Hyper-Scale Mobility online migration and a remote mirror volume at the same time. Chapter 3, “IBM Hyper-Scale Mobility with IBM FlashSystem A9000” on page 71 describes Hyper-Scale Mobility.

Similarly, volumes cannot be part of a IBM FlashSystem A9000 or A9000R data migration and a remote mirror volume at the same time. Data migration information is in Chapter 4, “Data migration with IBM FlashSystem A9000 and IBM FlashSystem A9000R” on page 103.

1.3.2 Remote mirror target configuration through Hyper-Scale Manager

The connections to the target (secondary) IBM FlashSystem A9000 or A9000R must be defined. The assumption here is that the physical connections and zoning (if using FC) have already been established. Target configuration is done from the Targets spoke of the Hub display using the following steps:

1. Add the target system by selecting the **Targets** spoke from the System Hub display and clicking the plus sign (**+)** to ADD a new target, as shown in Figure 1-17.

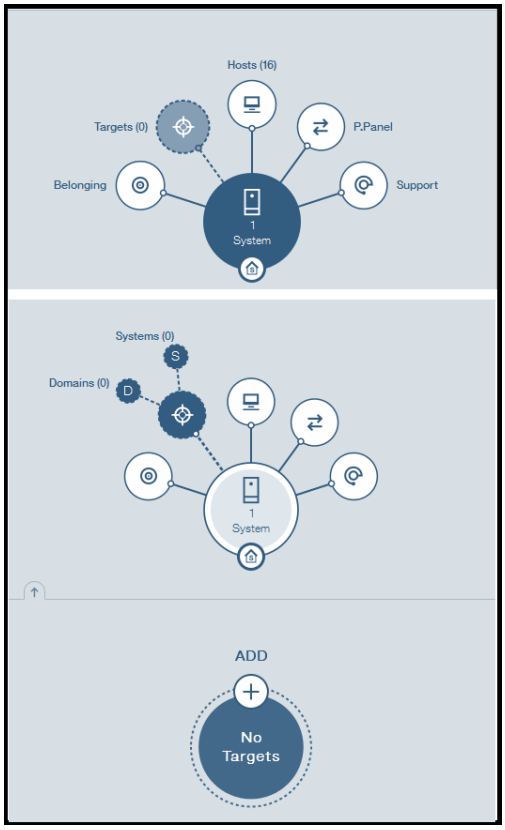

Figure 1-17 Create target

Important: IBM FlashSystem A9000 does not support using more than one mirroring target between two systems in a mirroring relation, or in IBM Hyper-Scale Mobility, because it can compromise the data on the destination.

The Hyper-Scale Manager (assuming it has a connection to the systems involved) prevents you from defining more than one target between two systems.

Depending on the version that you use, the XCLI might not prevent that action or give you a warning.

2. Choose the Remote System from the drop-down box, select the Remote Domain, define the type of connectivity to be used (mirroring or migration) and the protocol (iSCSI or FC), as shown in Figure 1-18 on page 17.

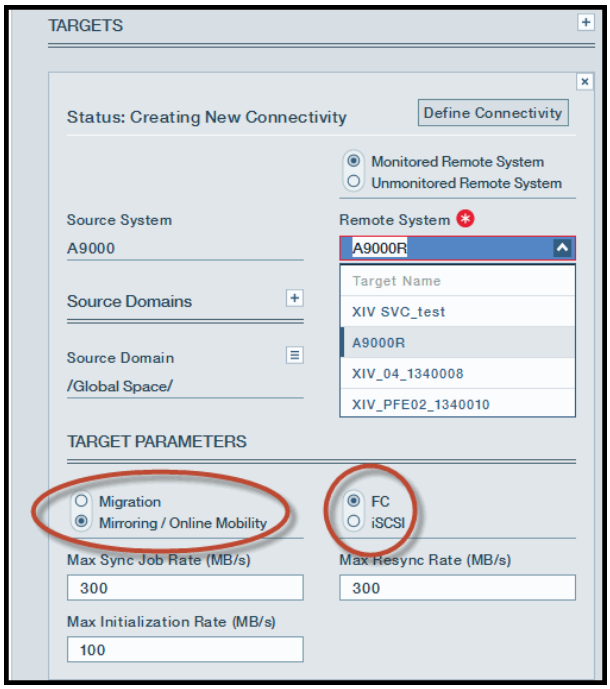

Figure 1-18 Target type and protocol

If choosing iSCSI for the target connection, also provide the iSCSI Name. If the target is monitored by the Hyper-Scale Manager, the iSCSI Name is filled in automatically, as shown in Figure 1-19.

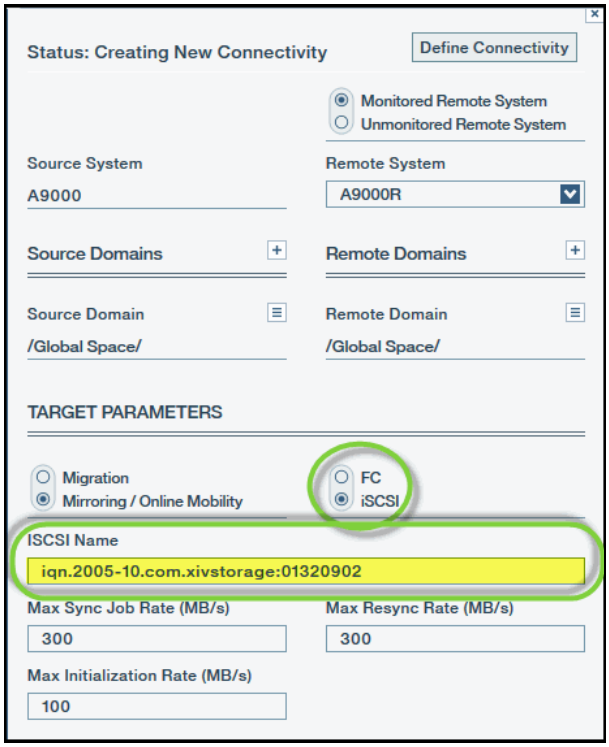

Figure 1-19 iSCSI Target Parameters

3. Optionally, modify the Max Sync Job Rate, Max Resync Rate, and Max Initialization Rate.

   IBM FlashSystem A9000 and A9000R allows a user-specifiable maximum rate (in MBps) for remote mirroring coupling initialization, a different user-specifiable maximum rate for normal sync jobs, and another for resynchronization.

   The initialization rate, sync job rate, and resynchronization rate are specified for each mirroring target using the XCLI **`target_config_sync_rates`** command or through the Hyper-Scale Manager, as shown in Figure 1-20.

   The actual effective initialization or synchronization rate also depends on the number and speed of connections between the A9000s and A9000Rs. The maximum initialization rate must be less than or equal to the maximum sync job rate (asynchronous mirroring only), which must be less than or equal to the maximum resynchronization rate.

   Important: In normal mirror operations, the rates are cumulative. For example, if initialization, synchronous, and asynchronous operations are all active, the amount of data the A9000 or A9000R attempts to send is the sum of those three values.

   The defaults are as follows:

   – Maximum initialization rate: 100 MBps
   – Maximum sync job: 300 MBps
   – Maximum resync rate: 300 MBps

4. Click **Apply,** as shown in Figure 1-20, to create the new Target.

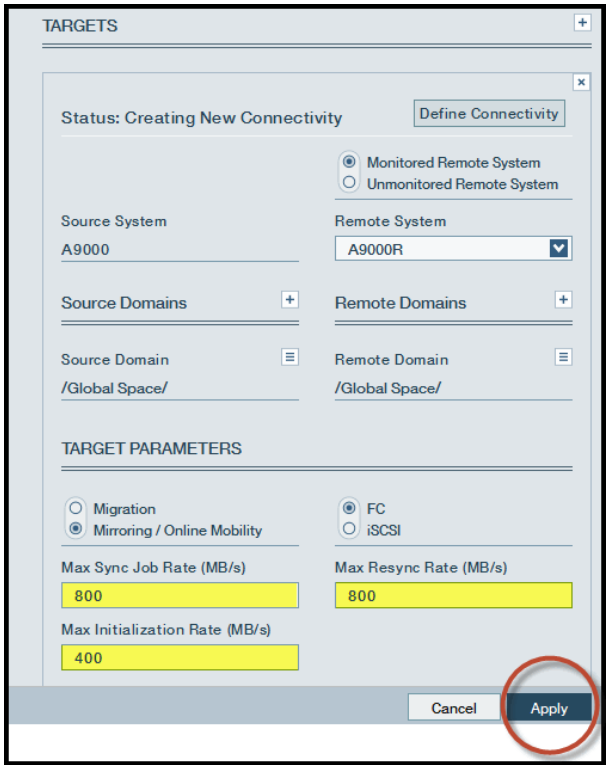

Figure 1-20 Click Apply to create Target

5. As shown in Figure 1-21, define connections by clicking **Define Connectivity** in the Targets panel to open the detail window, as shown in Figure 1-22.

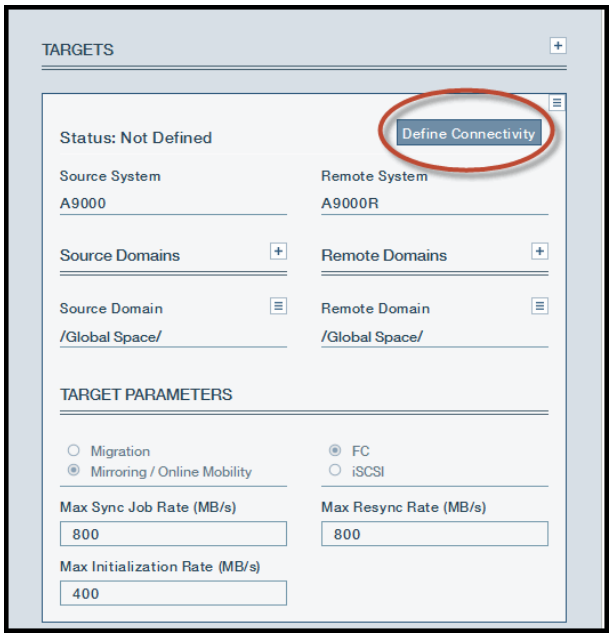

Figure 1-21 Define connectivity

Connections are easily defined by clicking **Auto Connect**. Remember that for FC ports an initiator must be connected to a target and the proper zoning must be established for the connections to be successful.

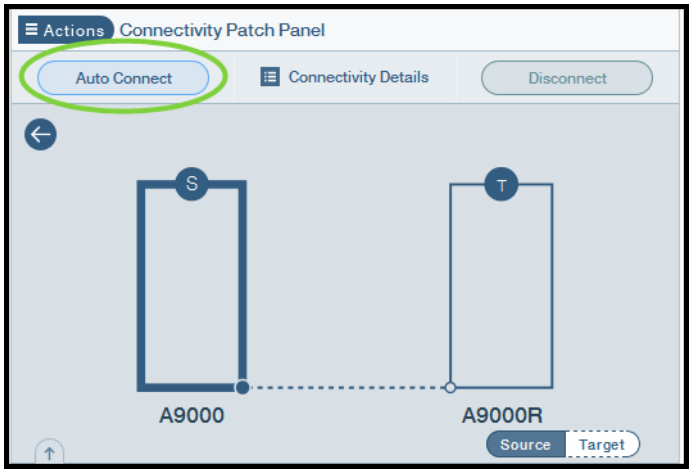

Figure 1-22 Define Connectivity Detail Display

When completed, the Source to Target connection in the diagram displayed changes from a dashed line to a solid line with a check mark, as depicted in Figure 1-23.

Figure 1-23 Connectivity Patch Panel after Auto Connect

6. Click **Connectivity Details** to open the Connectivity Patch Panel view that shows the active connections between Source and Target, as shown in Figure 1-24 (FC connections).

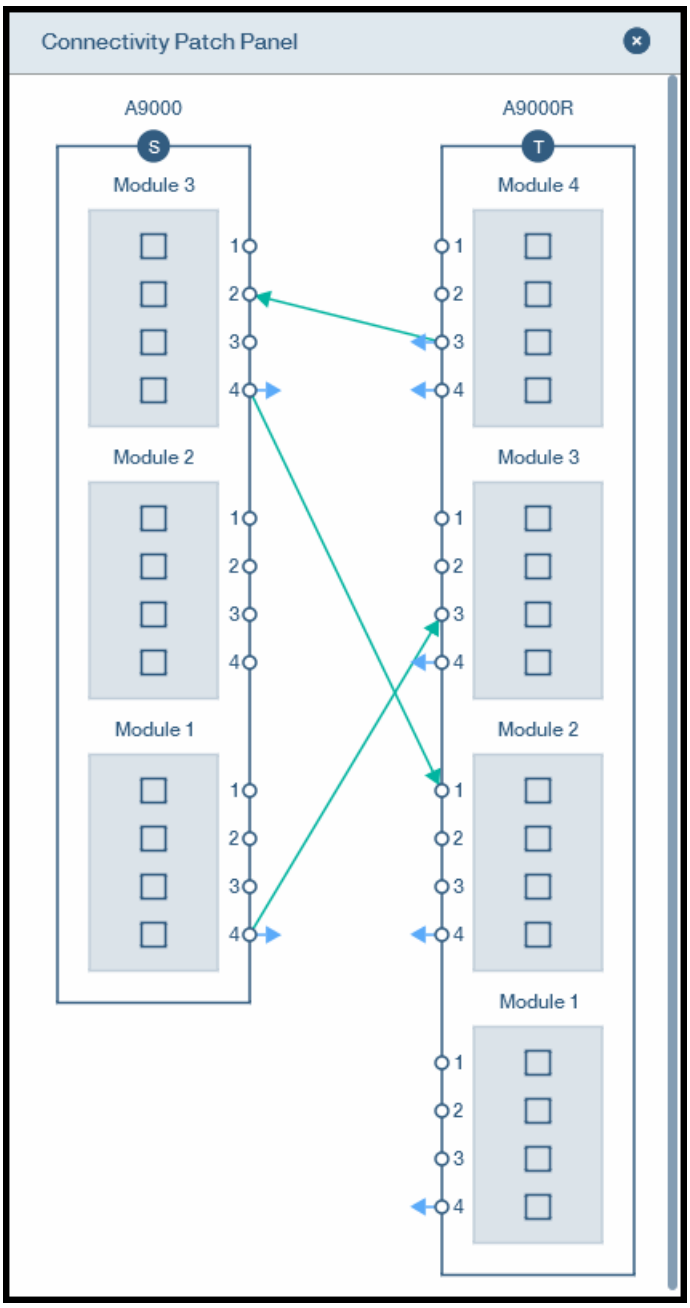

Figure 1-24 Connectivity Patch Panel (FC) in Hyper-Scale Manager

Note: The connectivity shown in Figure 1-24 is an example of FC connections and is *not* intended to reflect connectivity best practice.

Similarly, Figure 1-25 shows the auto-connected Source and Target connections for an iSCSI-connected mirroring target system.

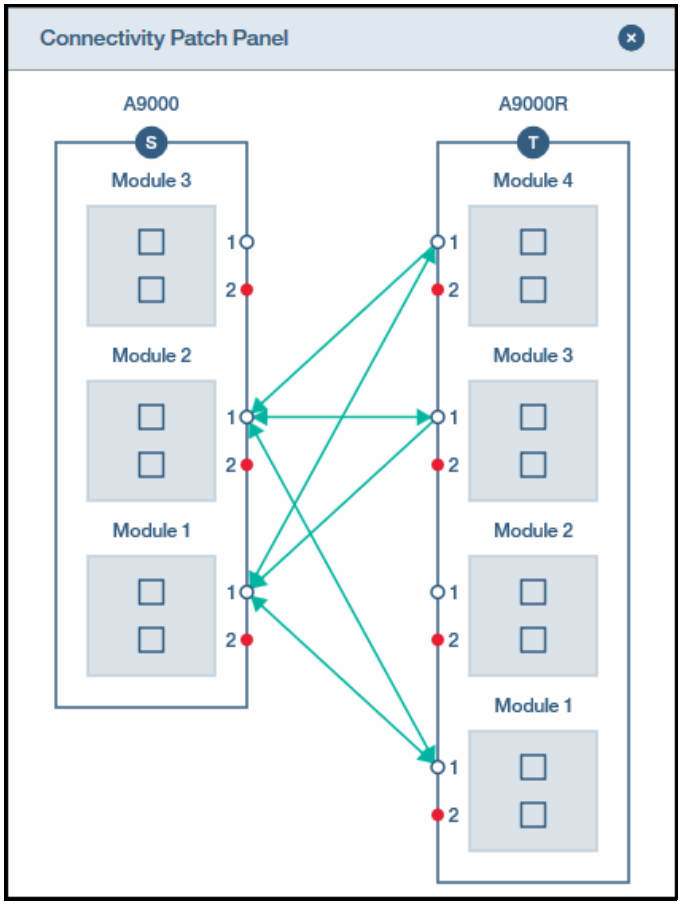

Figure 1-25 Connectivity Patch Panel (iSCSI) in Hyper-Scale Manager

Note: The connectivity shown in Figure 1-25 is an example of iSCSI connections and is *not* intended to reflect best practice connectivity.

To delete the connections between two IBM FlashSystem A9000s or A9000Rs, complete the following steps:

1. Click **Disconnect** in the Connectivity Patch Panel display, as shown in Figure 1-26, and click **Apply** in the resulting confirmation display.

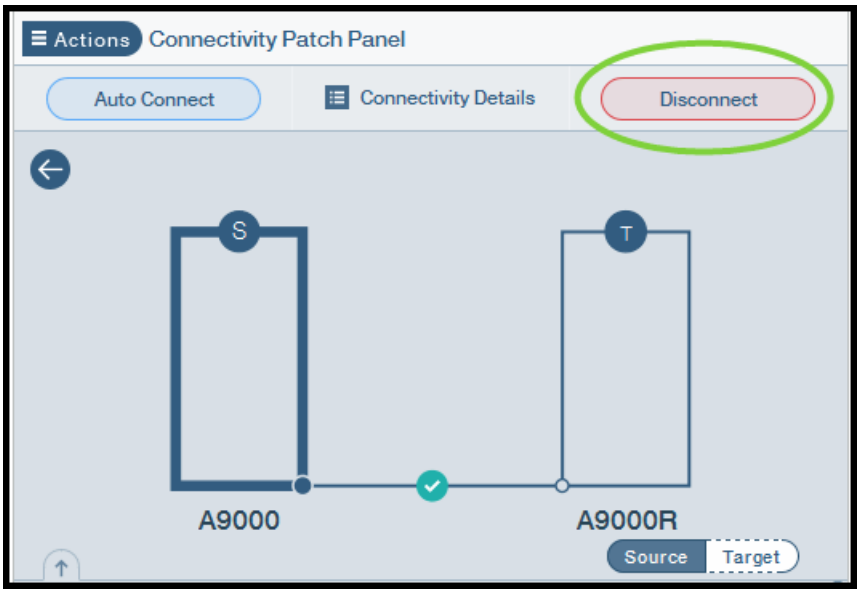

Figure 1-26 Target Connectivity Disconnect

2. From the Targets spoke in the Systems Hub display, click the **Actions** menu and select Remove Target, as shown in Figure 1-27. Click **Apply** in the resulting confirmation display.

Figure 1-27 Remove Target; note links have already been removed

1.3.3 XCLI examples

The following **XCLI** commands can be used to configure connectivity between the primary IBM FlashSystem A9000 or A9000R and the target or secondary IBM FlashSystem A9000 or A9000R (Example 1-4):

- **target_define**
- **target_list**
- **target_mirroring_allow**
- **target_config_sync_rates**
- **target_port_add**
- **target_port_list**
- **target_connectivity_define**
- **target_connectivity_list**

Example 1-4 Define Target XCLI commands

Source:

```
A9000>>target_define target="A9000R" protocol=FC
A9000>>target_list
Name     SCSI Type   Connected   Max Initialization Rate   Max Resync Rate   Max Syncjob Rate
A9000R   FC          no          100                       300               300
A9000>>target_mirroring_allow target=A9000R
A9000>>target_config_sync_rates target="A9000R" max_initialization_rate=400 max_resync_rate=800
max_syncjob_rate=800
A9000>>target_list
Name     SCSI Type   Connected   Max Initialization Rate   Max Resync Rate   Max Syncjob Rate
A9000R   FC          no          400                       800               800
A9000>>target_port_add target=A9000R fcaddress=5001738051A60142
A9000>>target_port_add target=A9000R fcaddress=5001738051A60120
A9000>>target_port_add target=A9000R fcaddress=5001738051A60132
A9000>>target_port_list
Target Name   Port Type   Active   WWPN               iSCSI Address   iSCSI Port
A9000R        FC          yes      5001738051A60142                   0
A9000R        FC          yes      5001738051A60120                   0
A9000R        FC          yes      5001738051A60132                   0
A9000>>target_connectivity_define target="A9000R" fcaddress=5001738051A60132
local_port=1:FC_Port:1:4
A9000>>target_connectivity_define target="A9000R" fcaddress=5001738051A60120
local_port=1:FC_Port:3:4
A9000>>target_connectivity_list
Target Name   Remote Port        FC Port         IP Interface   Active   Up
A9000R        5001738051A60132   1:FC_Port:1:4                  yes      yes
A9000R        5001738051A60120   1:FC_Port:3:4                  yes      yes
```

Target:

```
A9000R>>target_define target="A9000" protocol=FC system_id=22131

Warning:   Defining more than one target to the same remote system is not supported, and may
compromise the data on the slave system. Are you sure the remote system is not already defined
as a target? y/n: y
A9000R>>target_list
Name         SCSI Type   Connected   Max Initialization Rate   Max Resync Rate   Max Syncjob
Rate
```

```
A9000        FC          no          100                       300              300
A9000R>>target_mirroring_allow target=A9000
A9000R>>target_config_sync_rates target="A9000" max_initialization_rate=400 max_resync_rate=800
max_syncjob_rate=800
A9000R>>target_list
Name         SCSI Type   Connected   Max Initialization Rate   Max Resync Rate   Max Syncjob
Rate
A9000        FC          no          400                       800              800
A9000R>>target_mirroring_allow target=A9000
A9000R>>target_port_add target="A9000" fcaddress=5001738056730133
A9000R>>target_port_add target="A9000" fcaddress=5001738056730131
A9000R>>target_port_add target="A9000" fcaddress=5001738056730113
A9000R>>target_port_list
Target Name   Port Type   Active   WWPN               iSCSI Address   iSCSI Port
A9000         FC          yes      5001738056730133                   0
A9000         FC          yes      5001738056730131                   0
A9000         FC          yes      5001738056730113                   0
A9000R>>target_connectivity_define target="A9000" fcaddress=5001738056730131
local_port=1:FC_Port:4:3
A9000R>>target_connectivity_list
Target Name   Remote Port        FC Port         IP Interface   Active   Up
A9000         5001738056730131   1:FC_Port:4:3                  yes      yes
```

Similarly, the XCLI **target_connectivity_delete**, **target_port_delete**, and **target_delete** commands can be used to delete the connectivity between the primary IBM FlashSystem A9000 or A9000R and the secondary IBM FlashSystem A9000 or A9000R, as shown in Example 1-5.

Example 1-5 Delete Target XCLI commands

Source:

```
A9000>>target_connectivity_delete target="A9000R" fcaddress=5001738051A60132 local_port=1:FC_Port:1:4
A9000>>target_connectivity_delete target="A9000R" fcaddress=5001738051A60120 local_port=1:FC_Port:3:4
A9000>>target_port_delete target="A9000R" fcaddress=5001738051A60142
A9000>>target_port_delete target="A9000R" fcaddress=5001738051A60120
A9000>>target_port_delete target="A9000R" fcaddress=5001738051A60132
A9000>>target_delete target="A9000R"
```

Target:

```
A9000R>>target_connectivity_delete target="A9000" fcaddress=5001738056730131 local_port=1:FC_Port:4:3
A9000R>>target_port_delete target="A9000" fcaddress=5001738056730133
A9000R>>target_port_delete target="A9000" fcaddress=5001738056730131
A9000R>>target_port_delete target="A9000" fcaddress=5001738056730113
A9000R>>target_delete target="A9000"
```

2

Remote mirroring with IBM FlashSystem A9000 and IBM FlashSystem A9000R

The remote mirroring function provides a real-time copy between two or more IBM FlashSystem A9000 and A9000R systems. This feature provides a method to protect data from site failures.

This chapter includes the following sections:

- Mirroring Overview
- Boundaries
- Mirroring schemes
- Remote mirroring usage
- Remote mirroring scenarios
- Defining volume and consistency group mirroring
- Migrating data using remote mirroring
- Recovering from a site failure
- Deleting mirror coupling definitions
- Recovery from a temporary site outage
- Temporary deactivation of mirroring

2.1 Mirroring Overview

Mirroring creates a set of consistent data that can be used when there are problems with the production volumes. It can also be used for other purposes such as testing and backup on the remote site using snapshots of consistent data.

Remote mirroring is independent of applications and operating systems, and does not require host-processor cycle usage.

Remote mirroring can be a synchronous copy solution where a write operation is completed on both copies (local and remote sites) before an acknowledgment is returned to the host that issued the write. This type of remote mirroring is typically used for geographically close sites to minimize the effect of I/O delays, which are proportional to the distance between the sites.

Remote mirroring can also be an asynchronous solution. This is where consistent sets of data are copied to the remote location at predefined intervals at the same time the host I/O operations are acknowledged directly after they are written on the primary site alone. This configuration is typically used for longer distances between sites.

A reliable, dedicated network is preferred for mirroring. Links can be shared, but require available and consistent network bandwidth. The specified minimum bandwidth (50 Mbps for iSCSI) is a functional minimum and might not meet the replication speed that is required for a customer environment and workload.

Also, minimum bandwidths are *not* time-averaged, as typically reported by network monitoring packages; they are instantaneous, constant requirements, typically achievable only through network quality of service (QoS).

Unless otherwise noted, this chapter describes the basic concepts, functions, and terms that are common to both synchronous and asynchronous mirroring available with IBM FlashSystem A9000 and A9000R.

The following list defines the mirroring-related terms used in this document:

- Local site: This site consists of the primary storage system and the servers running applications that are stored on that storage system.
- Remote site: This site holds the mirror copy of the data and, usually, also has standby servers. A remote site can become the active production site using a consistent data copy.
- Primary: This denotes the IBM FlashSystem A9000 and A9000R used for production, during typical business routines, to serve hosts and have its data replicated to a secondary IBM FlashSystem A9000 and A9000R. You might also refer to it as the source system.
- Secondary. This denotes the system used during normal circumstances to act as the mirror (backup) for the primary. You might also refer to it as the destination system.
- Consistency group (CG): This is a set of related volumes on a single system that are treated as one logical unit. Thus, all CG data reflects correctly ordered writes across all respective volumes within the CGs. Consistency groups are supported within remote mirroring.
- Coupling: This is the pairing of volumes or CGs to form a mirror relationship between the source of a replication and its destination (target).

- Peer: This is one side of a coupling. It can be either a volume or a consistency group. However, peers must be of the same type (that is, both volumes or both CGs). Whenever a coupling is defined, a role is specified for each peer. One peer is designated as the source and the other peer is designated as the destination (target).
- Role: This denotes the actual role that the peer is fulfilling:
  - Source: A role that indicates that the peer serves host requests and acts as the source for replication.
  - Destination: A role that indicates that the peer does not serve host write requests (it can be used in read-only mode) and acts as the target for replication.

  Changing a peer's role might be necessary after the peer is recovered from a site or system, or due to a link failure or disruption.
- Sync job: This applies to asynchronous mirroring only. It denotes a synchronization procedure run by the source at user-configured intervals corresponding to the asynchronous mirroring definition; or, it is run upon manual execution of the XCLI command **`mirror_create_snapshot`**, which is also used for synchronous mirroring, but not as part of a scheduled job. The resulting job is referred to as snapshot mirror sync job, ad hoc sync job, or manual sync job in contrast with a scheduled sync job. The sync job entails synchronization of data updates recorded on the source since the creation time of the most-recent snapshot that was successfully synchronized.
- Offline initialization (offline init): A mechanism whereby IBM FlashSystem A9000 and A9000R, using HASH values, compares respective source and target 64 KB data blocks and copies over only the parts that have different data. Offline initialization expedites the synchronization of mirror pairs that are known to be inherently similar (for example, when an asynchronous pair is changed to a synchronous pair).

  This feature of IBM FlashSystem A9000 and A9000R can be used when the data links do not have adequate speed or capacity to transmit the entire volume in a timely fashion. In that case, the pair is first created when the systems are at close proximity and can use fast links. Then, when the storage system that hosts the remote mirror is placed at its final physical destination, only the changed data since those volumes were identical must be copied over the wire.
- Asynchronous schedule interval: This applies only to asynchronous mirroring. It represents how often the source automatically runs a new sync job. The default interval and the minimum possible is 20 seconds.
- Recovery point objective (RPO): The RPO is a setting that is applicable to asynchronous mirroring. It represents an objective set by the user, implying the maximal time difference considered acceptable between the mirror peers (the actual difference between mirror peers can be shorter or longer than the RPO set).

  An RPO of zero indicates that no difference between the mirror peers can be tolerated, and that implies that sync mirroring is required. An RPO that is greater than zero indicates that the replicated volume is less current or lags somewhat behind the source volume. In this case, there is a potential for certain transactions that have been run against the production volume to be rerun when applications start to use the replicated volume.

  For asynchronous mirroring, the required RPO is user-specified. The storage system then reports effective RPO and compares it to the required RPO.

  Connectivity, bandwidth, and distance between the systems directly impact RPO. More connectivity, greater bandwidth, and less distance typically enable a lower RPO.

2.1.1 Remote mirroring modes

As mentioned earlier, IBM FlashSystem A9000 and A9000R supports both synchronous mirroring and asynchronous mirroring:

- Synchronous mirroring

  Synchronous mirroring is designed to accommodate a requirement for zero RPO.

  To ensure that data is also written to the secondary system (destination role), an acknowledgment of the write operation to the host is issued only after the data is written to both storage systems. This ensures that mirroring peers always contain the same data. A write acknowledgment is returned to the host only after the write data has been cached by the IBM FlashSystem A9000 and A9000R systems at each site. This is depicted in Figure 2-1.

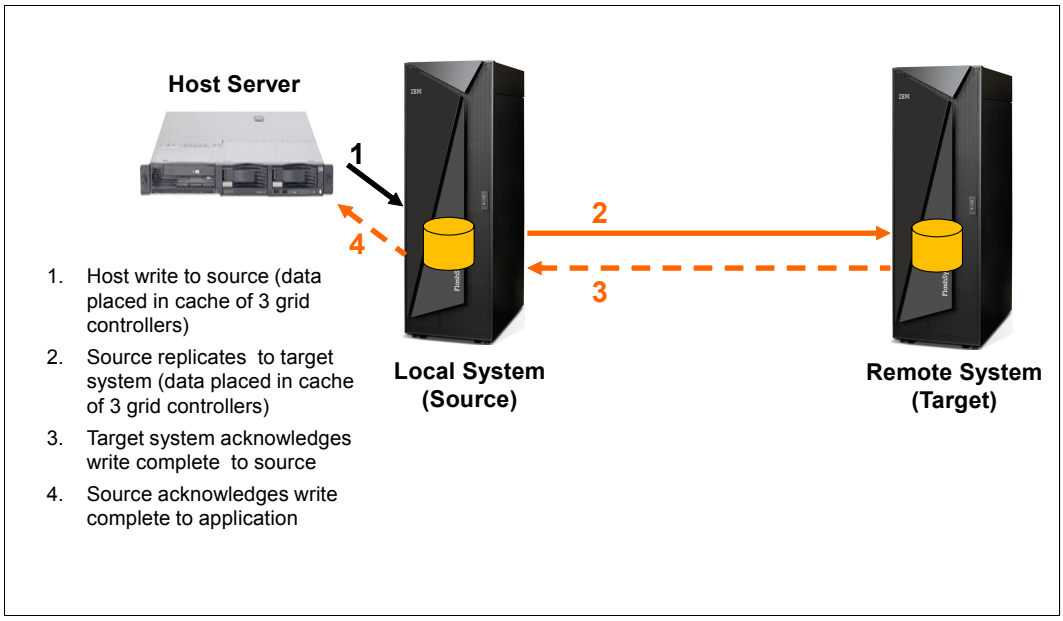

Figure 2-1 Synchronous mirroring

Host read operations are provisioned by the primary (source role), whereas writing is handled at the primary (source role) and replicated to the secondary systems.

- Asynchronous mirroring

  Asynchronous mirroring is designed to provide a consistent replica of data on a target peer through timely replication of data changes recorded on a source peer.

  Asynchronous mirroring uses the snapshot function, which creates a point-in-time image. In asynchronous mirroring, successive snapshots (point-in-time images) are made and used to create consistent data on the destination peers. The system sync job copies the data corresponding to the differences between two designated snapshots on the source (most-recent and last-replicated).

For asynchronous mirroring, acknowledgment of *write complete* is returned to the application as soon as the write data is received at the local storage system, as shown in Figure 2-2.

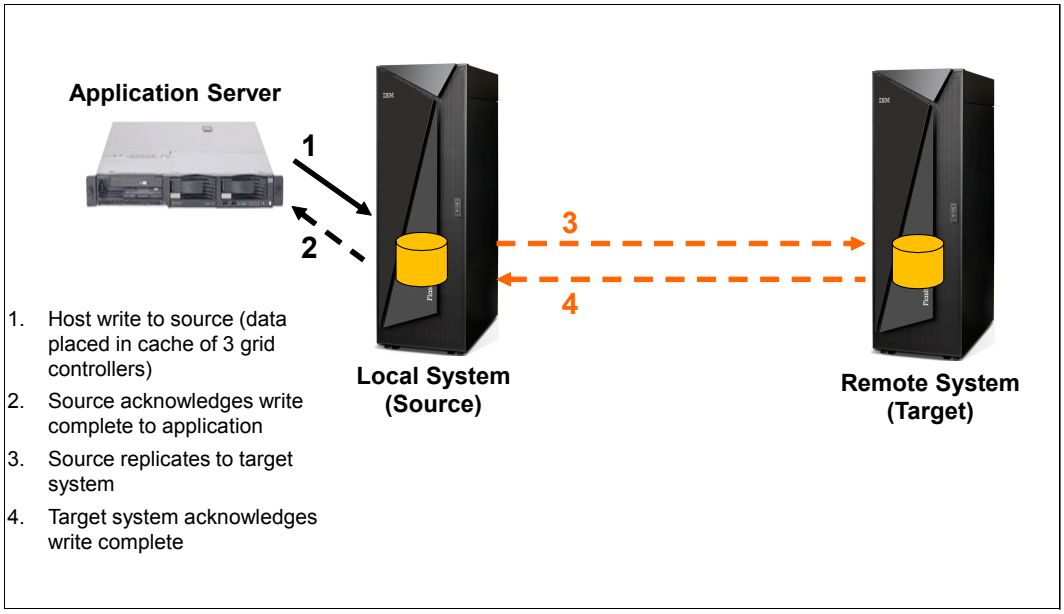

Figure 2-2 Asynchronous mirroring

The data replicated between two IBM FlashSystem A9000 and A9000R systems is non-data reduced. IBM FlashSystem A9000 and A9000R data reduction is inside one system only.

2.2 Boundaries

IBM FlashSystem A9000 and A9000R have the following boundaries or limitations:

- Maximum remote systems: The maximum number of remote systems that can be attached to a single primary is 10, with a maximum number of 32 ports on the target.
- Number of remote mirrors: The combined number of source and destination volumes (including in mirrored CG) cannot exceed 1536.
- Distance: Distance is limited only by the response time of the medium used. Use asynchronous mirroring when the distance causes unacceptable delays to the host I/O in synchronous mode.
- Consistency groups are supported within remote mirroring. The maximum number of consistency groups is 512.
- Snapshots: Snapshots are allowed with either the primary or secondary volumes without stopping the mirror. There are also special-purpose snapshots that are used in the mirroring process. Space must be available in the storage pool for snapshots.
- Source and destination peers cannot be the target of a copy operation and cannot be restored from a snapshot. Peers cannot be deleted or formatted without deleting the coupling first.
- Asynchronous volumes cannot be resized when mirroring is active.

2.3 Mirroring schemes

Mirroring, whether synchronous or asynchronous, requires two or more IBM FlashSystem A9000 and A9000R systems. The source system can be IBM FlashSystem A9000 or A9000R. The target system can be IBM FlashSystem A9000 or A9000R. Any combination of IBM FlashSystem A9000 and IBM FlashSystem A9000R is possible for mirroring.

The source and target of the asynchronous mirroring can be at the same site and form a local mirroring, or they can be at different sites and facilitate a disaster recovery plan. Figure 2-3 shows how peers can be spread across multiple storage systems and sites. Any system that is depicted can be an IBM FlashSystem A9000 and A9000R.

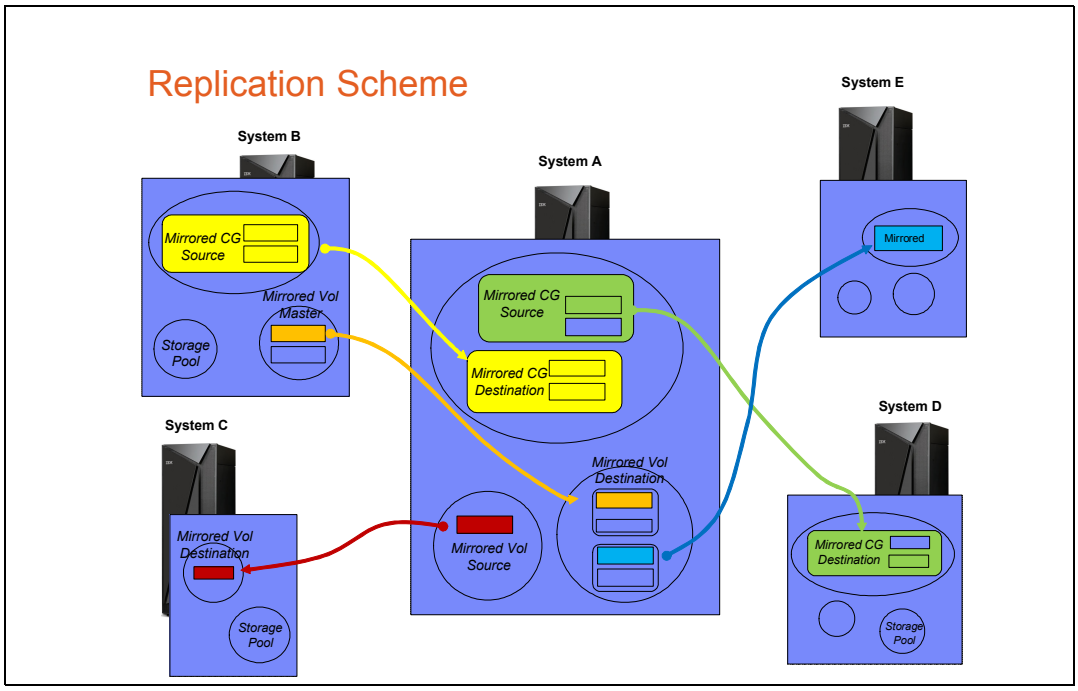

Figure 2-3 Mirroring replication schemes

A system can be connected to up to 10 other IBM FlashSystem A9000 and A9000R targets for mirroring purposes. Any system can be used simultaneously as replication source and replication target (for different volumes).

In a bidirectional configuration, a system concurrently functions as the replication source (source) for one or more couplings, and as the replication target (destination) for other couplings. Figure 2-3 illustrates possible schemes for how mirroring can be configured.

Figure 2-4 shows connectivity among systems and groups of systems, as shown in the Hyper-Scale Manager GUI.

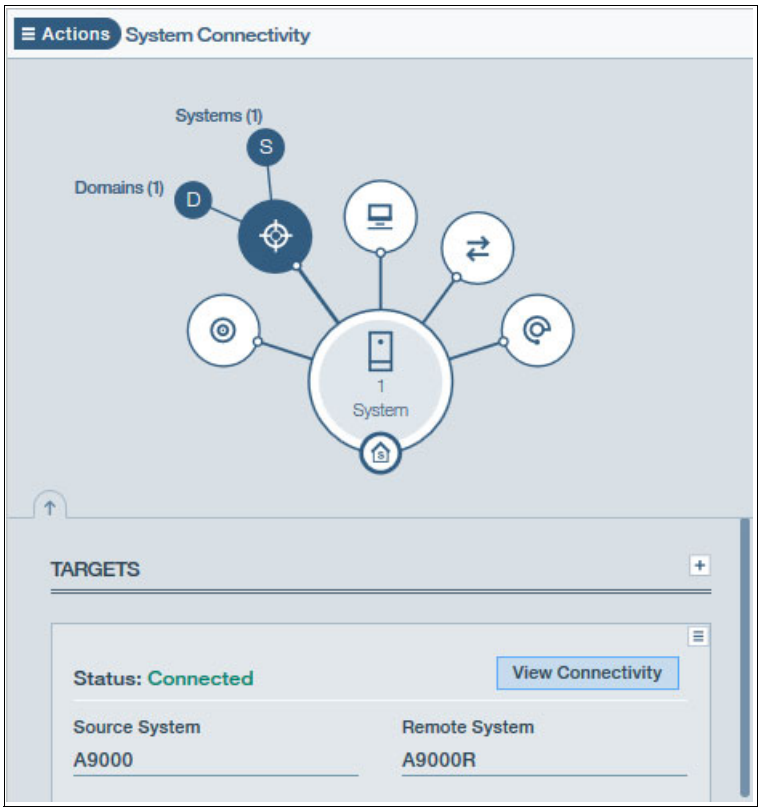

Figure 2-4 Connectivity among systems and group of systems in Hyper-Scale Manager

Figure 2-4 shows the target information. The target systems are listed below the TARGETS header in the Hyper-Scale Manager GUI. In this example, there is one target system. The names of the source system and the target system are shown. The remote system is the target system.

2.3.1 Peer designations and roles

A peer (volume or consistency group) is assigned either a source or a destination role when the mirror is defined. By default, in a new mirror definition, the location of the source denotes the primary system, and the target denotes the secondary system. An active mirror must have exactly one primary and exactly one secondary.

Important: A single system can contain both source volumes and consistency groups (mirroring to another system) and destination volumes and consistency groups (mirroring from another system). Peers in a source role and peers in a destination role on the same system must belong to different mirror couplings.

The various mirroring role status options are as follows:

- Designations:
  - Primary: The designation of the source peer, which is initially assigned the source role.
  - Secondary: The designation of the target peer, which initially plays the destination role.

- Role status:
  - Source: Denotes the peer with the source data in a mirror coupling. Such peers serve host requests and are the source for synchronization updates to the destination peer. Destination and source roles can be switched by using the **mirror_switch_roles** command if the status is synchronized for synchronous mirror and it is in an RPO_OK state for an asynchronous mirror. For both synchronous and asynchronous mirroring, the source can be changed (**mirror_change_role** command) to a destination if the status is inactive.
  - Destination: Denotes the target peer in a mirror. Such peers do not serve host write requests and accept synchronization updates from a corresponding source. A destination logical unit number (LUN) can be accessed in read-only mode by a host.

Consistency group (CG) within a IBM FlashSystem A9000 and A9000R

With mirroring (synchronous or asynchronous), the major reason for consistency groups is to handle many mirror pairs as a *group* (mirrored volumes are consistent). Instead of dealing with many mirror pairs individually, consistency groups simplify the handling of those related pairs.

Important: If your mirrored volumes are in a mirrored consistency group, you cannot do mirroring operations such as *deactivate* or *change_role* on a single volume basis. If you want to do mirroring operations, you must remove the volume from the consistency group.

Mirrored volumes in a consistency group must have the same domain and the same pool.

Consistency groups also play an important role in the recovery process. If mirroring was suspended (for example, because of complete link failure), data on different destination volumes at the remote system are consistent. However, when the links are up again and the resynchronization process is started, data spread across several destination volumes is not consistent until the source state has reached the *synchronized* state.

To preserve the consistent state of the destination volumes, the system automatically creates a snapshot of each destination volume and keeps it until the remote mirror volume pair is synchronized. In other words, the snapshot is kept until all pairs are synchronized to enable restoration to the same consistent point in time. If the remote mirror pairs are in a consistency group, the snapshot is taken for the whole group of destination volumes and the snapshots are preserved until *all* pairs are synchronized. Then the snapshot is deleted automatically.

2.3.2 Operational procedures

Mirroring operations involve configuration, initialization, ongoing operation, handling of communication failures, and role switching activities.

The following list defines the mirroring operation activities:

- Configuration

  Local and remote replication peers are defined by an administrator who specifies the source and destination peers roles. These peers can be volumes or consistency groups. The secondary peer provides a backup of the primary.

- Initialization

  Mirroring operations begin with a source volume that contains data and a formatted destination volume. The first step is to copy the data from the source volume (or CG) to the destination volume (or CG). This process is called *initialization*.

Initialization is performed one time in the lifetime of a mirror. After it is performed, both volumes or CGs are considered to be synchronized to a specific point in time. The completion of initialization marks the first point-in-time that a consistent source replica on the destination is available. Details of the process differ depending on the mirroring mode (synchronous or asynchronous).

- Offline initialization

  Offline initialization operation begins with a source volume that contains data and a destination volume, which also contains data and is related to this same source. At this step, only different blocks of data are copied from the source to its destination. Offline initialization can be run whenever a mirror pair was suspended or when the mirror type changes from asynchronous to synchronous or by restoring an image (or backup) from the source volume.

- Mirror mode switching

  The toggling between synchronous and asynchronous modes implies the deactivation of the current mirror mode, the deletion of the mirror pair and the respective snapshots on both ends, and unlocking of the destination mirror. Then the new mode is selected and a new mirror relationship is created between the peers. Using the offline initialization, only the new data that was written to the primary since the deletion of the original mirror is copied over. Thus, the toggling between the two operational modes does not require a full copy.

- Ongoing operation

  After the initialization process is complete, normal mirroring operations begin.

  In synchronous mirroring, normal ongoing operation means that all data written to the primary volume or CG is first mirrored to the destination volume or CG. At any point in time, the source and destination volumes or CGs are identical except for any unacknowledged (pending) writes.

  In asynchronous mirroring, ongoing operation means that data is written to the source volume or CG, and is replicated to the destination volume or CG at specified intervals.

- Monitoring

  The storage system effectively monitors the mirror activity and places events in the event log for error conditions. Alerts can be set up to notify the administrator of such conditions. You must have set up SNMP trap monitoring tools or email notification to be informed about abnormal mirroring situations.

- Handling of communication failures

  Sometimes the communication between the sites might break down. The source continues to serve host requests as the synchronous mirroring is based on best effort, to minimize any impact to the host write operations. Upon recovery from a link down incident, the changed data is copied to the remote site and mirroring is resumed. Events are generated for link failures.

- Role switching

  If required, mirror peer roles of the destination and source can be switched. Role switching is always initiated at the source site. Usually, this is done for maintenance operations or because of a test drill that verifies the disaster recovery (DR) procedures. Use role switching cautiously, especially with asynchronous mirroring. When roles are switched for an asynchronous mirror, data can be lost for an interval up to the RPO time because the remote site is typically lagging, in time, for a given asynchronous pair.

Role switching in the case of synchronous mirror is designed so that no data loss occurs. Use role switching only for cases such as a catastrophic host failure at the source site when the pairing is intact, but there have been no write operations to the source since the last sync job was completed.

- Role changing

  In a disaster at the primary site, the source peer might fail. To allow read/write access to the volumes at the remote site, the volume's role must be changed from destination to source. A *role change* changes only the role of the system volumes or CGs to which the command was addressed. Remote mirror peer volumes or CGs are not changed automatically. That is why changing roles on both mirror sides if mirroring is to be restored is important (if possible).

2.3.3 Mirroring status

The status of a mirror is affected by several factors, such as the links between the IBM FlashSystem A9000 and A9000R systems or the initialization state.

Link status

The link status reflects the connection from the source to the destination volume or CG. A link has a direction (from local site to remote or vice versa). A failed link or a failed secondary system both result in a link error status. The link state is one of the factors determining the mirror operational status. Links can be in one of the following states:

- Connected: Link is up and is functioning.
- Connectivity Failure

Figure 2-5 depicts how the link status is displayed from the System Connectivity window for a specific Target System. The status *Connected* is shown in green letters.

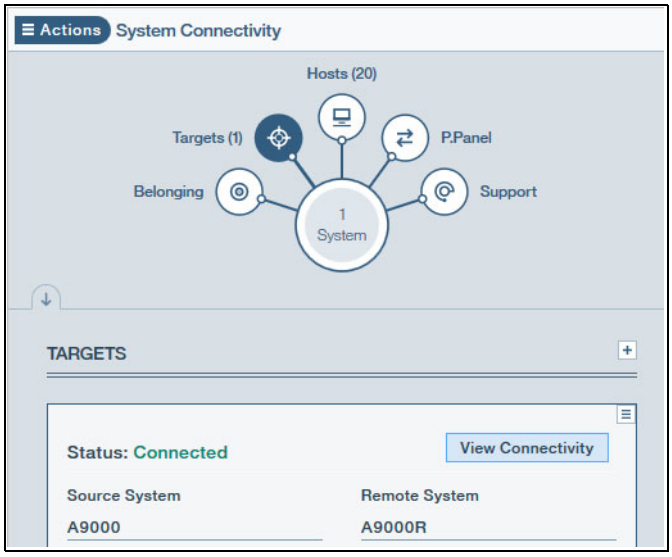

Figure 2-5 Red indicates link down, green indicates link up

Figure 2-6 on page 35 depicts how the link status is reflected in the Hyper-Scale Manager GUI when the Target Connectivity Details is selected. (Click **View Connectivity** for Target Connectivity Details.)

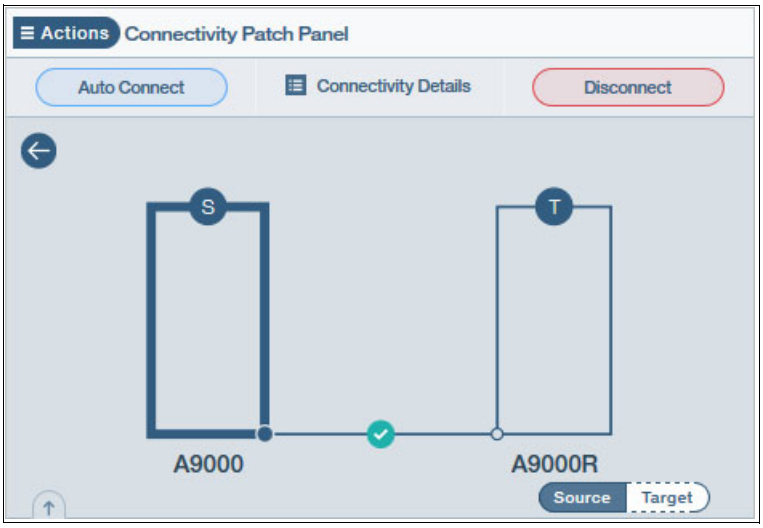

Figure 2-6 Connectivity details

If several links (at least two) are in one direction and one link fails, it usually does not affect mirroring if the bandwidth of the remaining link is high enough to keep up with the data traffic.

Monitoring the link utilization

The mirroring bandwidth of the links must be high enough to cope with the data traffic generated by the changes on the source volumes. During the planning phase, before setting up mirroring, monitor the write activity to the local volumes. The bandwidth of the links for mirroring must be as large as the peak write workload. The physical link between the source and target might also been used for other traffic. This link between the source and target must sustain the maximal peak workload of all workload using the link.

Note: Check the physical link utilisation to assure bandwidth capacity for the peak write workload.

Hyper-Scale Manager Version 5.1 does not show statistics on port level. You can check the physical link utilization using the SAN components, for example an FC switch for FC connections. Using the volume view of the Hyper-Scale Manage gives you the current bandwidth of the selected volumes.

Figure 2-7 on page 36 shows the current write bandwidth of one target volume. The volume source is on the system A9000R and the target volume is on system A9000.

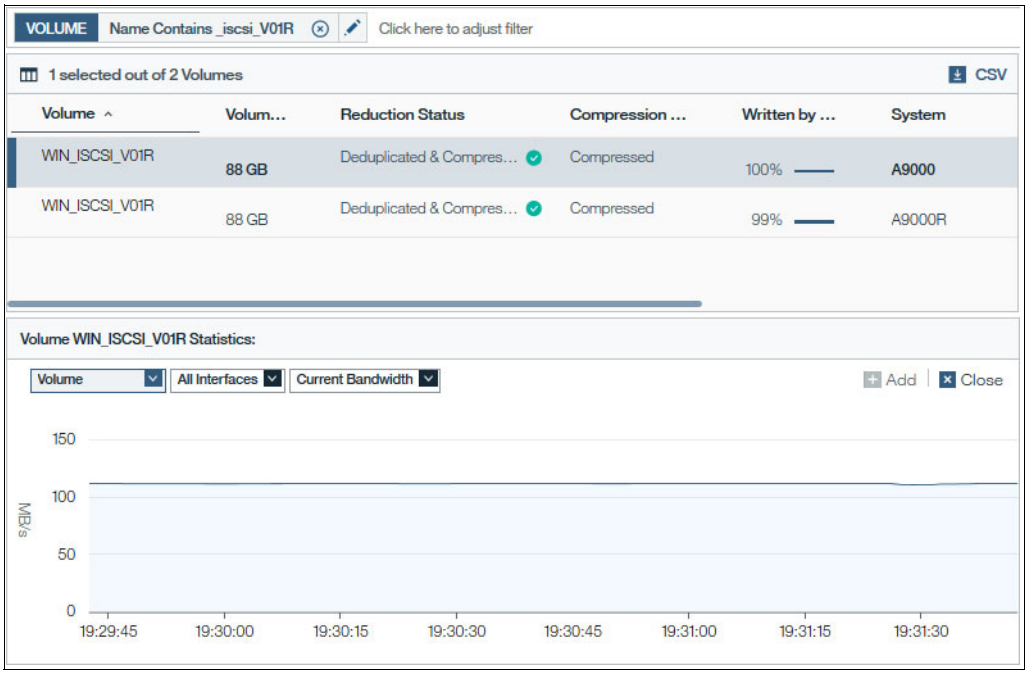

Figure 2-7 Monitoring bandwidth of the target volume

Mirror operational status

In this section, the mirroring status is defined as either *operational* or *non_operational*:

- Mirroring is operational in the following situations:
  - The activation state is *active*.
  - The link is up.
  - Both peers have different roles (source or destination).
  - The mirror is active.
- Mirroring is non_operational in the following situations:
  - The mirror is inactive.
  - The link is in an error state or deactivated (link down).

Synchronous mirroring states

Note: This section applies only to synchronous mirroring.

The synchronization status reflects whether the data of the destination's volume is identical to the source's volume. Because the purpose of the remote mirroring feature is to ensure that the destination's volume is an exact copy of the source's volume, this status indicates whether this objective is being achieved.

The following states or statuses are possible:

- Initializing

  The first step in remote mirroring is to create a copy of all the data from the source volume or CG to the destination volume or CG. During this initial copy phase, the status remains *initializing*.

- Synchronized (source volume or CG) and consistent (destination volume or CG)

  This status indicates that all data that has been written to the source volume or CG has also been written to the destination volume or CG. Ideally, the source and destination volumes or CGs must always be synchronized. However, this does not always indicate that the two volumes are absolutely identical in a disaster. This is because there are situations when there might be a limited amount of data that was written to one volume, but that was not yet written to its peer volume. This means that the write operations have not yet been acknowledged to the respective hosts. Such writes are known as *pending writes* or *data in flight*.

- Unsynchronized (source volume) and inconsistent (destination volume)

  After a volume or CG has completed the initializing stage and achieved the synchronized status, it can become *unsynchronized* (source) and *inconsistent* (destination). This occurs when it is not known whether all the data that has been written to the source volume has also been written to the destination volume. This status can occur in the following cases:

  - The communications link is down and, as a result, certain data might have been written to the source volume, but was not yet written to the destination volume.
  - Secondary system is down. This is similar to communication link errors because in this state the Primary system is updated, whereas the secondary system is not.
  - Remote mirroring is deactivated. As a result, certain data might have been written to the source volume and not to the secondary volume.

The system tracks the partitions that have been modified on the source volumes. When the link is operational again or the remote mirroring is reactivated, these changed partitions are sent to the remote system and applied to the respective destination's volumes.

Asynchronous mirroring states

Note: This section applies only to asynchronous mirroring.

The mirror state can be either inactive or initializing:

- Inactive: The synchronization process is disabled. It is possible to delete a mirror in this state.
- Initializing: The initial copy is not done. Synchronization does not start until the initialization completes. The mirror cannot be deleted during this state.

Important: In cases of an unstable data link, it is possible for the initialization to restart. In this case, the progress bar returns to the left side of the display. This does not mean that the initialization is starting again from the beginning. On the restart of a mirror initialization, the initialization resumes where it left off, and the progress bar displays the percent complete of the remaining data to copy, not the percentage of the full volume.

When initialization is complete, the synchronization process is enabled. Then, it is possible to run sync jobs and copy data between source and destination. Synchronization states are as follows:

- RPO_OK: Synchronization completed within the specified sync job interval time (RPO).
- RPO_Lagging: Synchronization completed but took longer that the specified interval time (RPO).

2.4 Remote mirroring usage

Remote mirroring solutions can be used to address multiple types of failures and planned outages. The failure scenarios vary. They can be a result of events that affect a single system. The failure can stem from events that affect an entire data center or campus. Worse, they can be caused by events that affect a whole geographical region.

Figure 2-8 shows recovery plan solutions that provide protection from single-system failures, local disasters, and regional disasters.

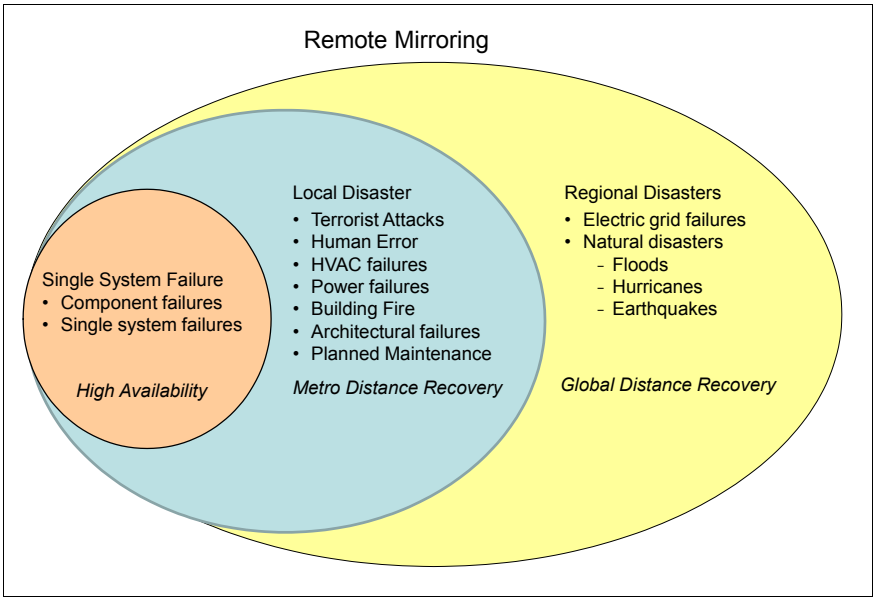

Figure 2-8 Disaster recovery protection levels

Several configurations are possible:

- Single-site high-availability remote mirroring configuration

  Protection for the event of a failure or planned outage of a system (single-system failure) can be provided by a *zero-distance* high-availability (HA) solution including another storage system in the same location (zero distance). Typical usage of this configuration is a synchronous mirroring solution that is part of a high-availability clustering solution that includes both servers and storage systems.

Figure 2-9 shows a single-site high-availability configuration where both systems are in the same data center.

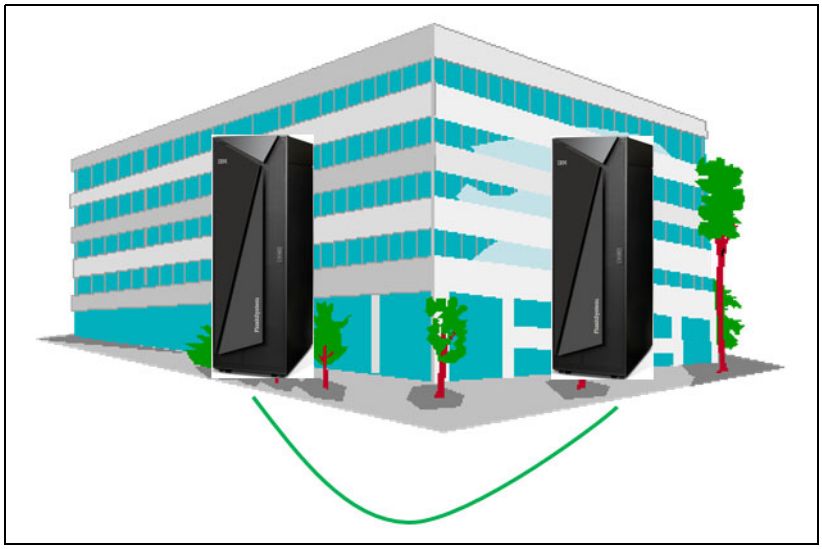

Figure 2-9 Single site HA configuration

- Metro region remote mirroring configuration

  Protection during a failure or planned outage of an entire location (*local disaster*) can be provided by a metro distance disaster recovery solution, including another system in a different location within a metro region. The two systems might be in different buildings on a corporate campus or in different buildings within the same city (typically up to approximately 100 km apart). IBM FlashSystem A9000 and A9000R deliver the lowest response time to the host. The added latency due to the distance of the mirrored systems must be tolerated by the host if using synchronous mirroring.

  Note: The SAN infrastructure must be able to sustain the micro latency of IBM FlashSystem A9000 and a9000R.

  Typical usage of this configuration is a synchronous mirroring solution. Figure 2-10 shows a metro region disaster recovery configuration.

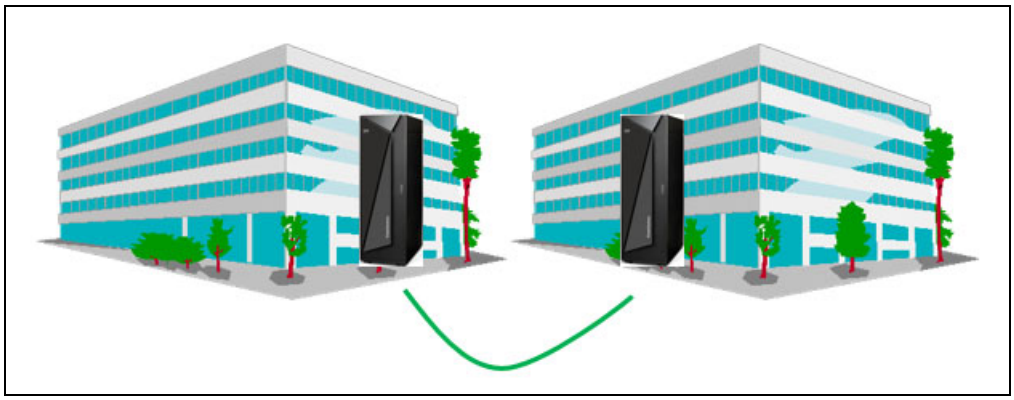

Figure 2-10 Metro region disaster recovery configuration

- Out-of-region remote mirroring configuration

  Protection during a failure or planned outage of an entire geographic region (*regional disaster*) can be provided by a global distance disaster recovery solution including another system in a different location outside the metro region. The two locations might be separated by up to a global distance. Typical usage of this configuration is an asynchronous mirroring solution. Figure 2-11 shows an out-of-region disaster recovery configuration.

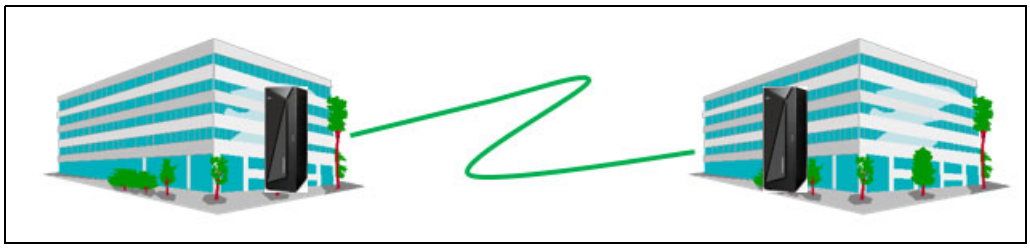

Figure 2-11 Out-of-region disaster recovery configuration

2.4.1 Using snapshots

Snapshots can be used with remote mirroring to provide copies of production data for business or IT purposes. Moreover, when used with remote mirroring, snapshots provide protection against data corruption. Like any continuous or near-continuous remote mirroring solution, remote mirroring cannot protect against software data corruption because the corrupted data is copied as part of the remote mirroring solution. However, the snapshot function provides a point-in-time image that can be used for a rapid recovery if software data corruption occurs after the snapshot was taken. The snapshot can be used in combination with remote mirroring, as illustrated in Figure 2-12.

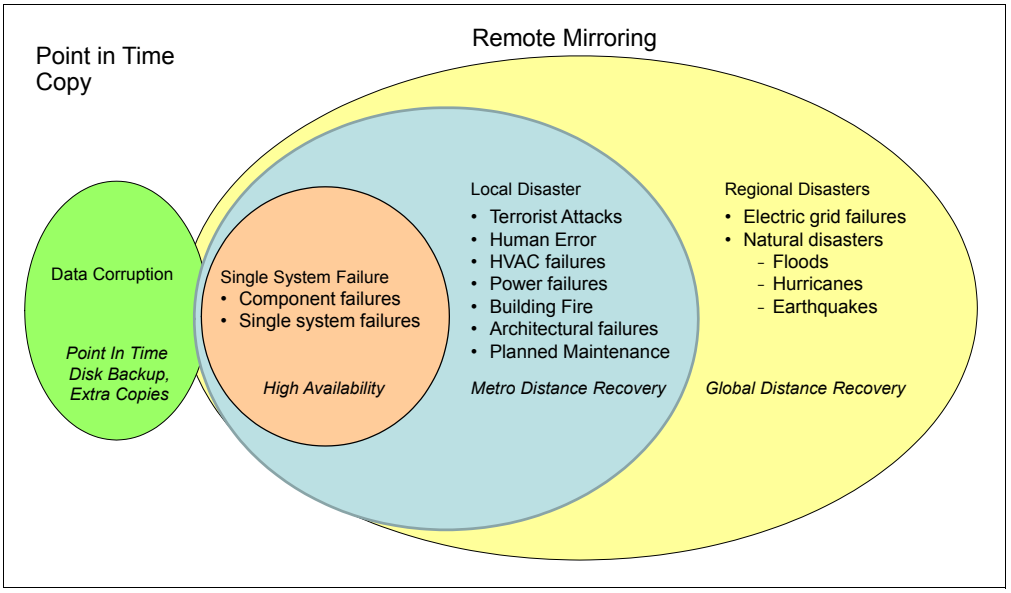

Figure 2-12 Combining snapshots with remote mirroring

Recovery using a snapshot requires deletion and re-creation of the mirror.

- Snapshot (within a single system)

  Protection for the event of software *data corruption* can be provided by restoring the volume to a healthy point-in-time snapshot. The snapshot can be backed up if needed.

- Local snapshot plus remote mirroring configuration

  A snapshot of the production (local) volume can be used in addition to remote mirroring of the production volume when protection from logical data corruption is required in addition to protection against failures and disasters. The extra snapshot of the production volume provides a quick restoration to recover from data corruption. An extra snapshot of the production (local) volume can also be used for other business or IT purposes such as reporting, data mining, development, and testing.

  Figure 2-13 shows a local snapshot plus remote mirroring configuration.

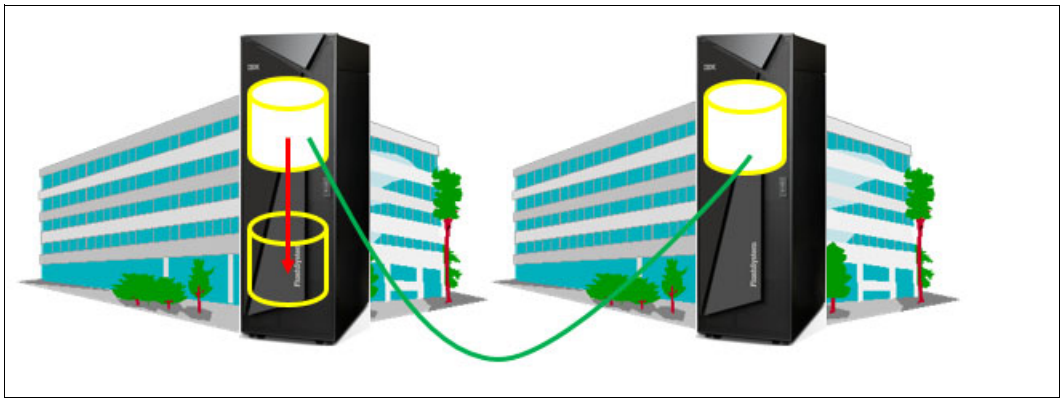

Figure 2-13 Local snapshot plus remote mirroring configuration

- Remote snapshot plus remote mirroring configuration

  A snapshot of the consistent replicated data at the remote site can be used in addition to remote mirroring to provide an extra consistent copy of data. This copy can be used for business purposes such as data mining, reporting, and for IT purposes, such as remote backup to tape or development, test, and quality assurance. Figure 2-14 shows a remote snapshot plus remote mirroring configuration.

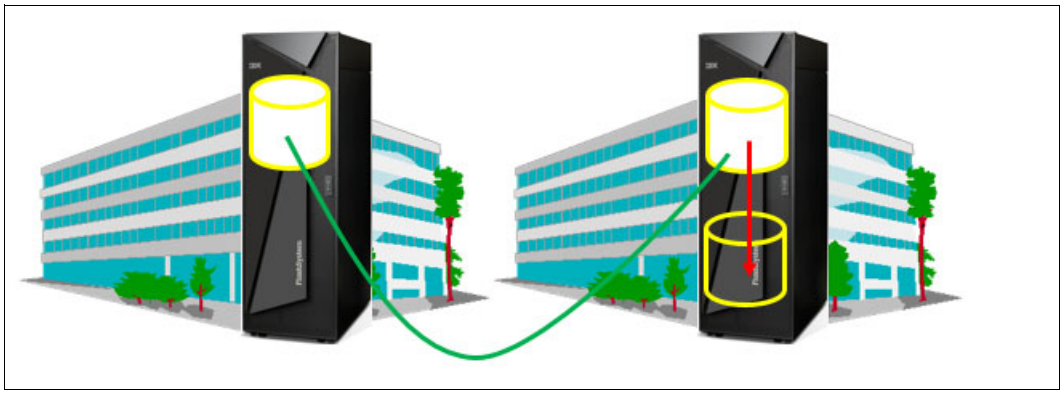

Figure 2-14 Remote snapshot plus remote mirroring configuration

2.5 Remote mirroring scenarios

The following sections show the basic steps used to set up mirrored volumes and mirrored consistency groups. This section includes these topics:

- Defining volume and consistency group mirroring
- Migrating data using remote mirroring
- Recovering from a site failure
- Deleting mirror coupling definitions
- Recovery from a temporary site outage
- Temporary deactivation of mirroring

The remote mirroring setup is described in 1.3, “Using the Hyper-Scale Manager or XCLI for remote mirror” on page 9.

2.6 Defining volume and consistency group mirroring

The two peers in the mirror can be either two volumes (volume peers) or two consistency groups (CG peers).

Each of the two peers in the mirroring relationship is given a designation and a role. The designation indicates the original or *normal* function of each of the two peers: Primary or secondary. The peer designation does not change with operational actions or commands. If necessary, the peer designation can be changed by explicit user command or action.

The role of a peer indicates its current operational function (either source or destination). The operational role of a peer can change as the result of user commands or actions. Peer roles typically change during DR testing or a true disaster recovery and production site switch.

When a mirror coupling is created, the first peer specified is the source for data to be replicated to the target system, so it is given the primary designation and the source role. The first peer can be, for example, the volumes or consistency group (CG) at site 1, as shown in Figure 2-15 on page 43.

Important: A consistency group to be mirrored must not contain any volumes when the CG coupling is defined, or the coupling will not allow it to be defined.

The second peer specified (or automatically created by the storage system) when the mirroring coupling is created is the target of data replication, so it is given the secondary designation and the destination role.

2.6.1 Defining synchronous volume mirroring

Select the volumes in the Volume view and select **Mirror** in the Hub view, as shown in Figure 2-15 on page 43.

Then add the mirroring for the volumes. The target volumes is created automatically when two or more volumes are selected for mirroring. When setting up a mirror for one volume, you can select the destination volume name and decide if a new volume is created or an existing volume is used.

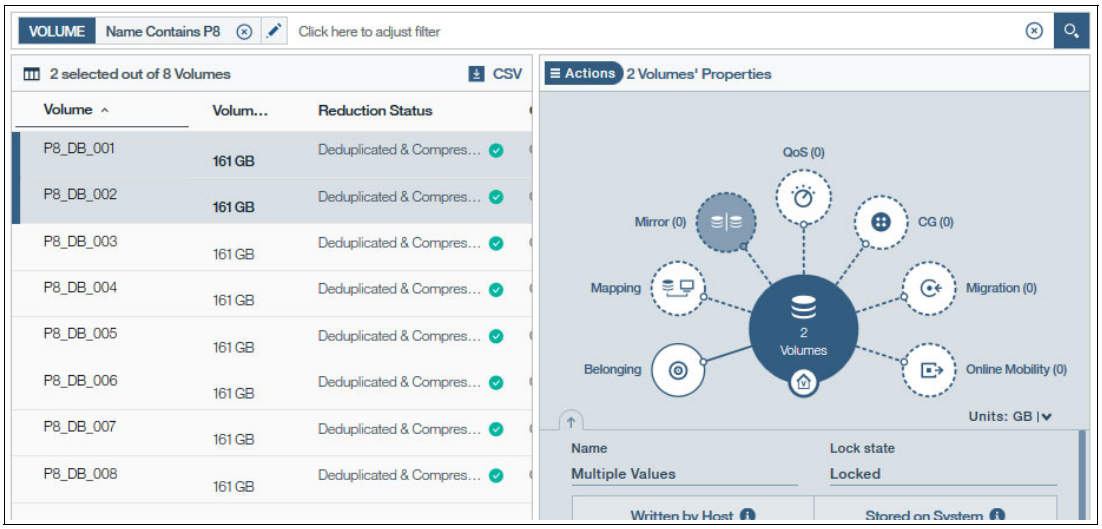

Figure 2-15 Selecting volumes for mirroring

Figure 2-16 shows the parameters for synchronous mirroring.

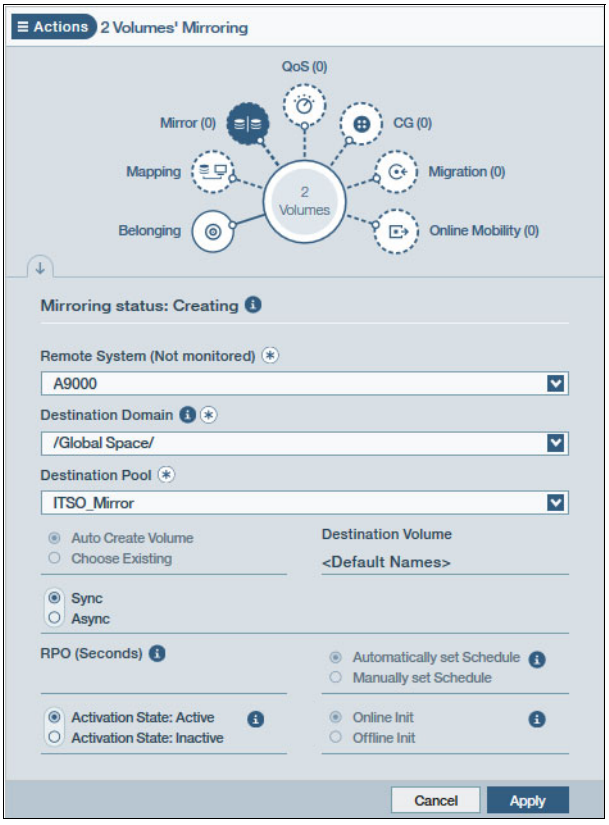

Figure 2-16 Volume mirror parameters

Click **Apply** to establish the volume mirrors. Because **Activation State: Active** is selected, the synchronization starts immediately, as shown in Figure 2-17.

Figure 2-17 Mirroring status

You can check the mirroring initialization progress in the Volume pane. One volume after the other is synchronized. You can change the listing header to show the mirror status. Figure 2-18 shows only the target volumes. Volumes PB_DB_001 to PB_DB_005 are synchronized. Volumes PB_DB_006 and PB_DB_007 are initializing. The target volumes are locked.

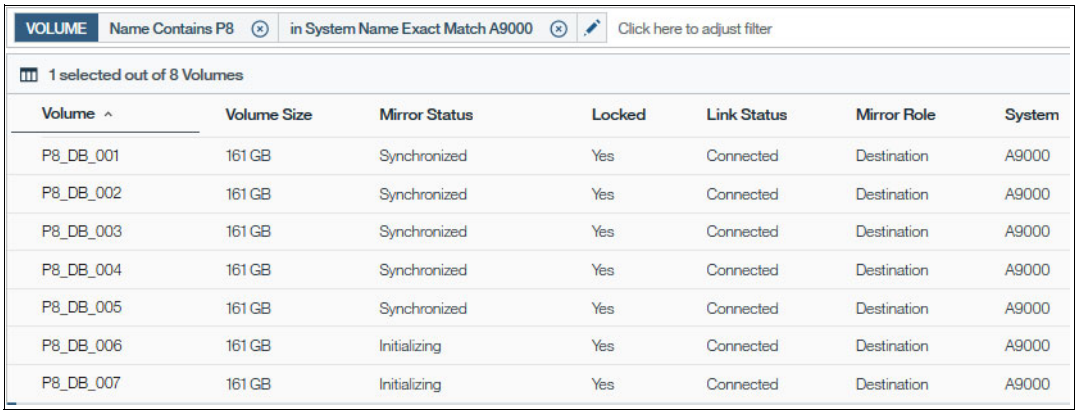

Volume	Volume Size	Mirror Status	Locked	Link Status	Mirror Role	System
P8_DB_001	161 GB	Synchronized	Yes	Connected	Destination	A9000
P8_DB_002	161 GB	Synchronized	Yes	Connected	Destination	A9000
P8_DB_003	161 GB	Synchronized	Yes	Connected	Destination	A9000
P8_DB_004	161 GB	Synchronized	Yes	Connected	Destination	A9000
P8_DB_005	161 GB	Synchronized	Yes	Connected	Destination	A9000
P8_DB_006	161 GB	Initializing	Yes	Connected	Destination	A9000
P8_DB_007	161 GB	Initializing	Yes	Connected	Destination	A9000

Figure 2-18 Mirroring progress

You can check the initializing status of one volume by selecting the mirror information of a volume, as shown in Figure 2-19.

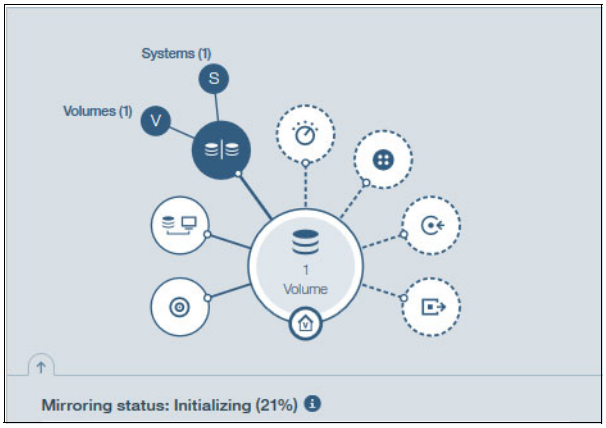

Figure 2-19 Initializing status of one mirrored volume

After synchronizing is completed, the status *Synchronized* is shown, as depicted in Figure 2-20.

Figure 2-20 Volume mirror state synchronized

When a mirror coupling is first activated, all actual data on the source is either copied to the destination (normal initialization) or verified to be on the destination, and only changed data is copied (offline initialization). This process is referred to as *initialization*. Remote mirroring copies volume identification information (that is, physical volume ID/PVID) and any actual data on the volumes. Space that has not been used is not copied.

Initialization might take a significant amount of time if a large amount of data exists on the source when a mirror coupling is activated. As discussed earlier, the rate for this initial copy of data can be specified by the user as described in “Using the Hyper-Scale Manager or XCLI for remote mirror” on page 9. The speed of this initial copy of data is also affected by the connectivity and bandwidth (number of links and link speed) between the primary and secondary systems.

As an option to remove the impact of distance on initialization, mirroring can be initialized with the target system installed locally. The target system can then be disconnected after initialization, shipped to the remote site, reconnected, and the mirroring reactivated.

2.6.2 Defining asynchronous volume mirroring

Select the volumes in the Volume view and select **Mirror** in the Hub view, as shown in Figure 2-21.

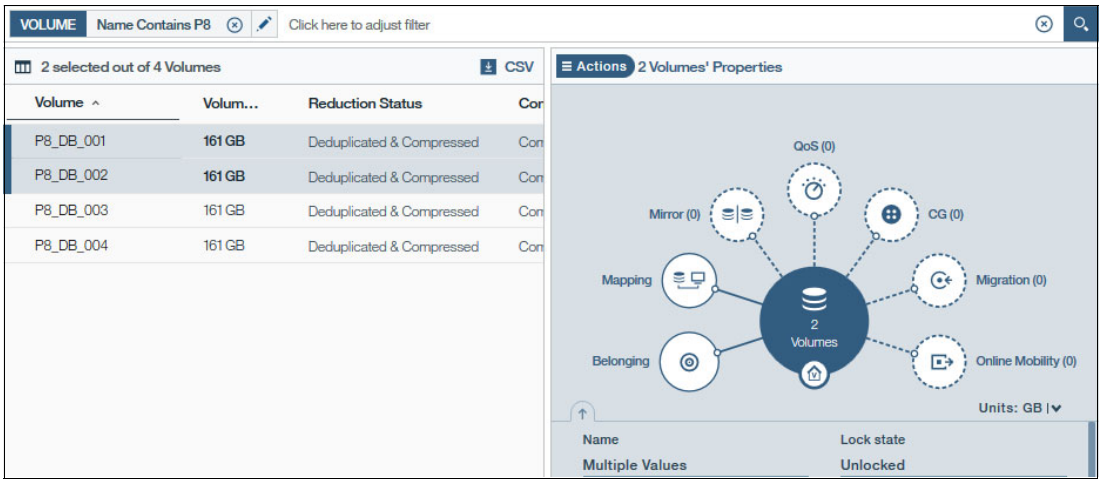

Figure 2-21 Selecting volumes for mirroring

Then add the mirroring for the volumes. The target volumes are created automatically when two or more volumes are selected for mirroring.

When setting up a mirror for one volume, you can select the destination volume name and decide if a new volume is created or an existing volume is used.

Figure 2-22 on page 47 shows the parameters for asynchronous mirroring. For asynchronous replication, select **Async** and enter an RPO time.

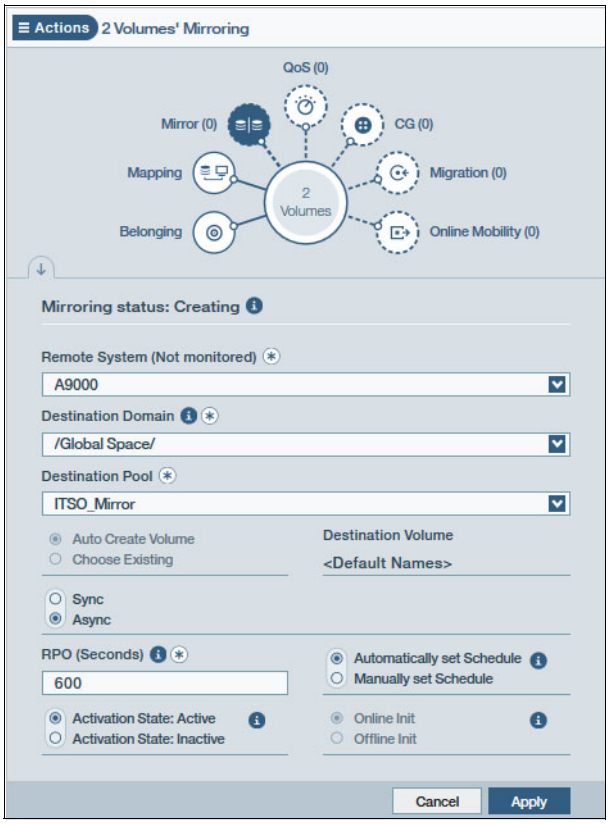

Figure 2-22 Volume asynchronous mirroring settings

You can check the mirroring initialization progress of one volume in the Volume Mirroring window. Figure 2-23 shows that the volume P8_DB_001 is initialized to 64%.

Figure 2-23 Volume mirroring initializing

You can change the listing header to show the mirror status. Figure Figure 2-24 shows the source and target volumes. The RPO and the effective RPO are shown in this example. The value **RPO OK** in the *Mirror Status* column means that the synchronisation runs within the defined RPO.

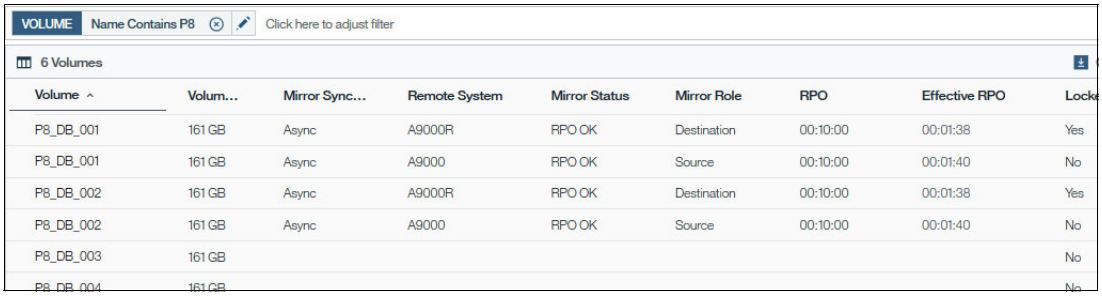

Volume	Volum...	Mirror Sync...	Remote System	Mirror Status	Mirror Role	RPO	Effective RPO	Locke
P8_DB_001	161 GB	Async	A9000R	RPO OK	Destination	00:10:00	00:01:38	Yes
P8_DB_001	161 GB	Async	A9000	RPO OK	Source	00:10:00	00:01:40	No
P8_DB_002	161 GB	Async	A9000R	RPO OK	Destination	00:10:00	00:01:38	Yes
P8_DB_002	161 GB	Async	A9000	RPO OK	Source	00:10:00	00:01:40	No
P8_DB_003	161 GB							No
P8_DB_004	161 GB							No

Figure 2-24 Asynchronous mirroring progress

You can check the initializing status of one volume by selecting the mirror information of a volume, as shown in Figure 2-25.

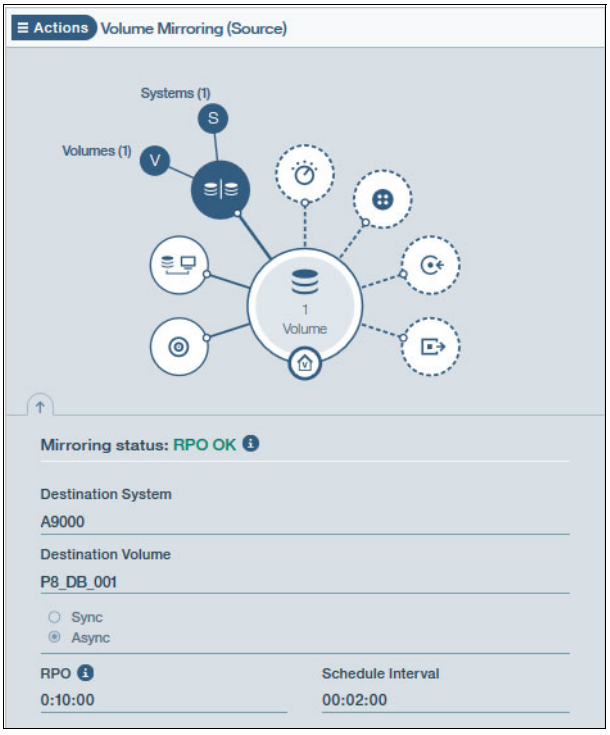

Figure 2-25 Mirroring status

In Figure 2-25 the mirroring status is RPO OK. This status means that the mirror synchronisation meets the defined RPO. You can see the snapshots associated with the mirrored volumes in the Snapshot view of the Hyper-Scale Manager, as shown in Figure 2-26 on page 49.

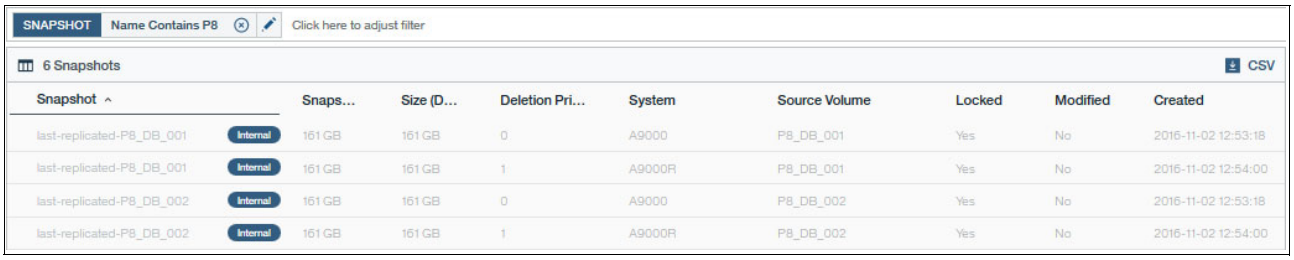

Figure 2-26 Snapshots associated with synchronically mirrored volumes

The Snapshot Properties windows shows the details of the selected snapshot which is the first line of Figure 2-26, as shown in Figure 2-27.

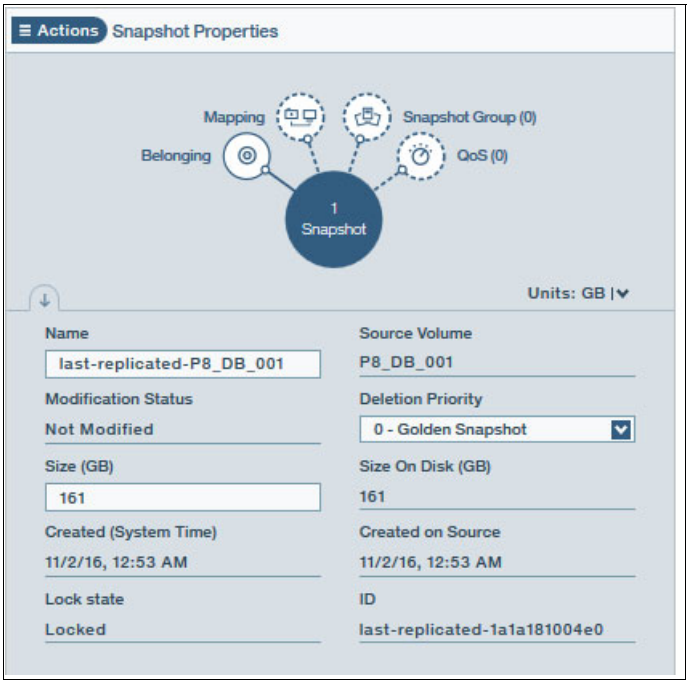

Figure 2-27 Snapshot detail

2.6.3 Using offline initialization

To avoid the impact of distance on initialization you can use offline initialization. The peers can be synchronized locally and the DR system moved to its remote site; or, if the target system is already at a remote location with limited WAN capabilities, you can apply an image backup of the source volume onto the destination. You can then activate the offline mirroring initialization. If a backup tape is physically transported to the remote site, it must be a volume image backup.

The mirror pairing is defined normally, with the addition of specifying the `offline init` option when making the definition of the pairing, as shown in Figure 2-28 on page 50. The destination volume must be entered manually. The activation state is automatically set to *inactive* when using the `offline init` option.

If a remote mirroring configuration is set up when a volume is first created (that is, before any application data is written to the volume), initialization is fast.

When a consistency group mirror coupling is created, the CG must be empty so there is no data movement and the initialization process is fast.

The mirror coupling status at the end of initialization differs for synchronous mirroring and asynchronous mirroring, but in either case, when initialization is complete, a consistent set of data exists at the remote site.

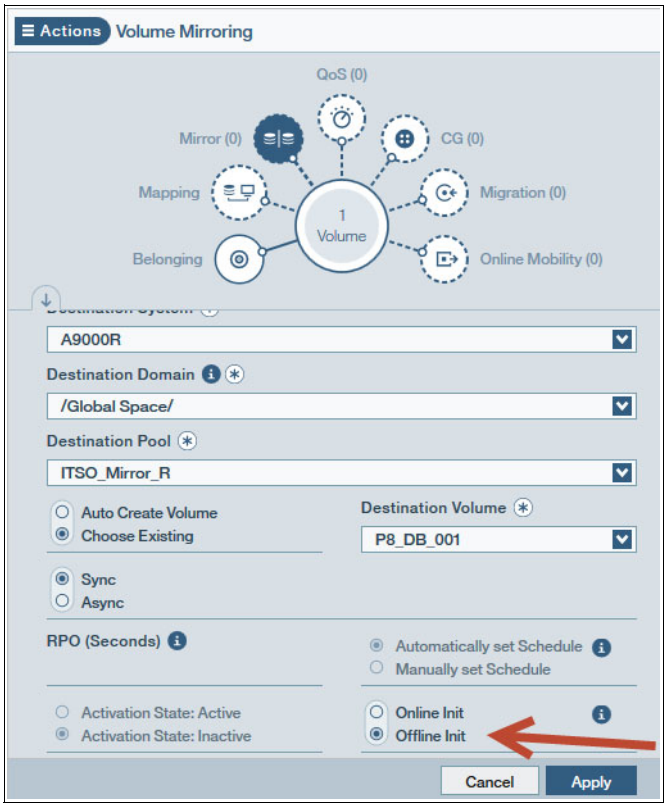

Figure 2-28 Offline init selected

2.6.4 Defining consistency group mirroring

Consistency group mirroring uses one CG on the source system and one CG on the destination (target) system. Both CGs must be empty, meaning they do not contain any volume, before defining the CG mirroring.

Select the CG in the CONSISTENCY GROUP view, and select **Mirror** in the Hub view as shown in Figure 2-29 on page 51.

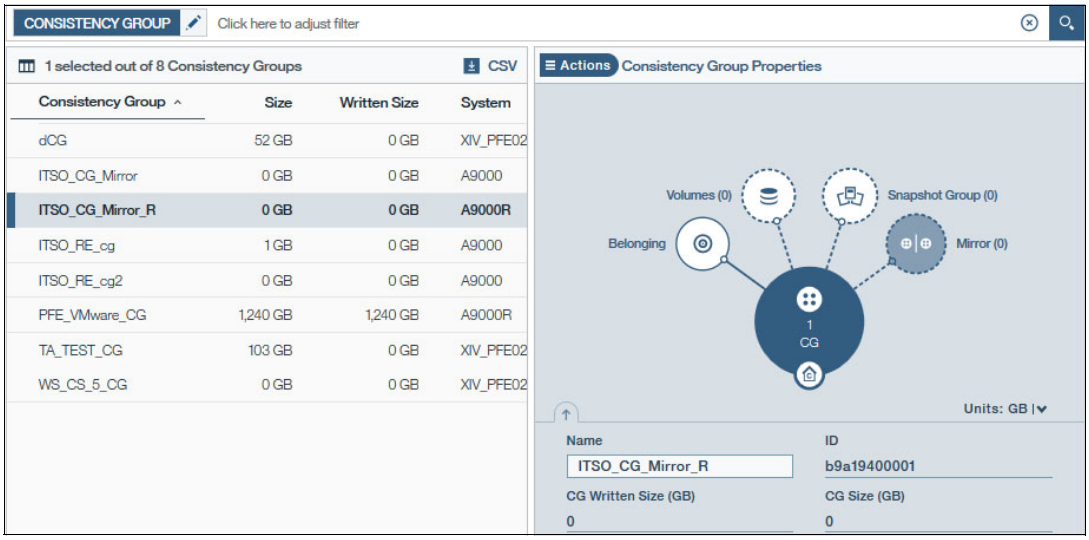

Figure 2-29 Setting the consistency group for mirroring

Then add the mirroring for the consistency group. Figure 2-30 shows the parameters for synchronous mirroring.

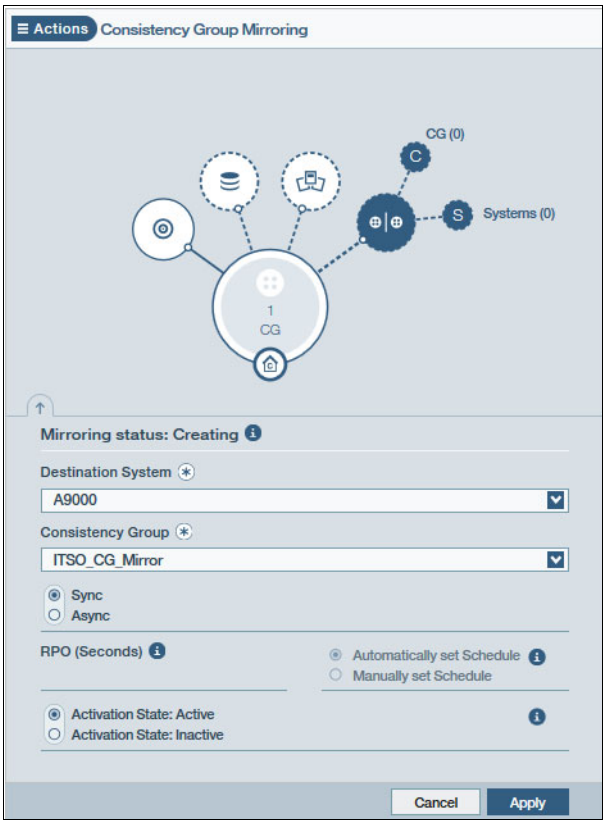

Figure 2-30 Consistency group mirror parameters

Click **Apply** to establish the volume mirrors. Because there are no volumes in the CGs, the mirroring status is synchronized, as shown in Figure 2-31 on page 52.

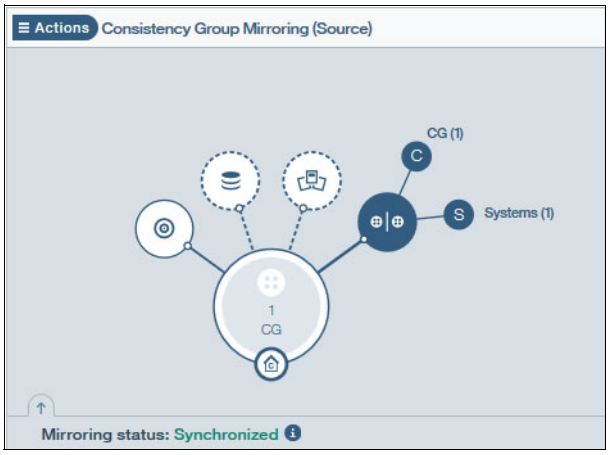

Figure 2-31 Consistency group mirroring status

2.6.5 Adding a mirrored volume to a mirrored consistency group

Before adding a volume to the CG, these prerequisites must be met:

- Volume and CG must be associated with the same source pool and target pool.
- Both volume synchronization status and mirrored CG synchronization status are either `RPO OK` for asynchronous mirroring or `Synchronized` for synchronous mirroring.
- The volume mirroring settings must be identical to those of the CG:
  - Mirroring type
  - Mirroring role
  - Mirroring status
  - Mirroring target
  - Target pool

To add a mirrored volume to a mirrored consistency group, complete the following two steps:

1. Define the volume mirror from the volume at the source system to the volume at the destination (target) system with activation state **active**.

   Monitor initialization until it is complete. Volume mirroring initialization must complete before the mirrored volume can be moved to a mirrored CG.

2. Add the source mirrored volume to the source mirrored consistency group at the source system. The destination volume is automatically added to the destination consistency group at the destination system.

After a mirrored volume has completed initialization, the source volume can be added to a pre-existing mirrored CG. Only mirrored volumes can be added to a mirrored consistency group. Select a mirrored volume in the Volume view and select **Mirror** in the Hub view and add the CG, as shown in Figure 2-32 on page 53.

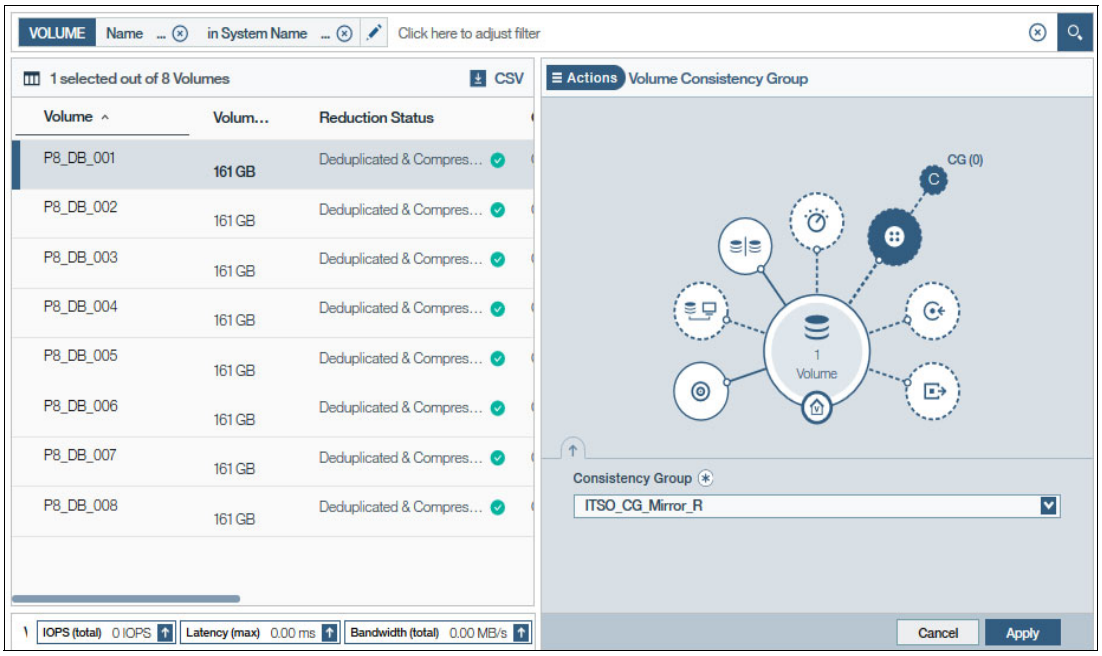

Figure 2-32 Adding a mirrored volume to a mirrored consistency group

Click **Apply** to add the volume to the chosen consistency group. All volumes have the same mirrored synchronisation state in a mirrored consistency group. In this example consistency group the volume mirror is synchronised, therefore all volumes that are added to this consistency group must be synchronised.

Select the CG in the CONSISTENCY GROUP view and select **Volumes** in the Hub view, as shown in Figure 2-33.

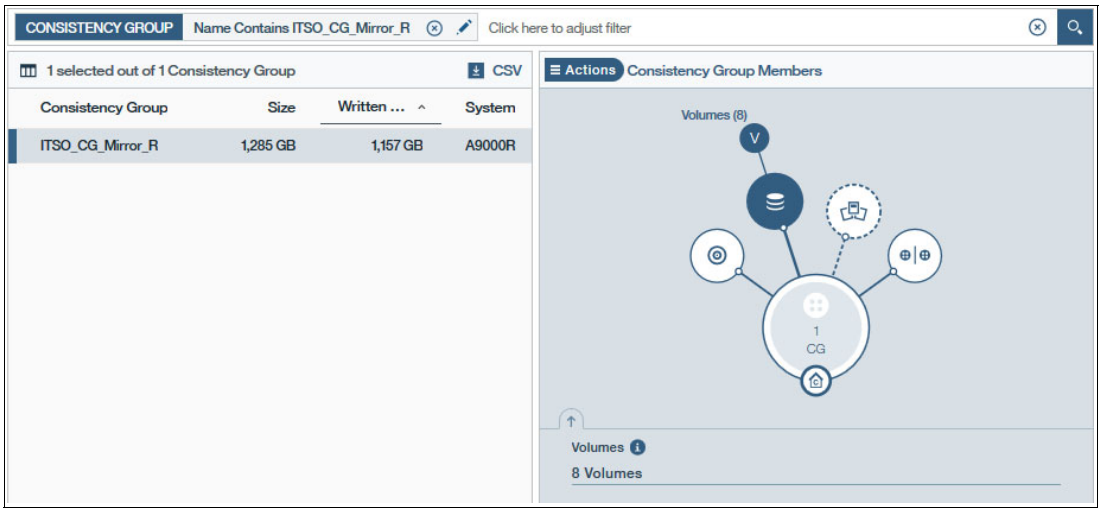

Figure 2-33 Consistency group reference to the mirrored volumes

Select the V, as shown in Figure 2-33, and the mirrored volumes of the mirrored consistency group are listed, as shown in Figure 2-34 on page 54.

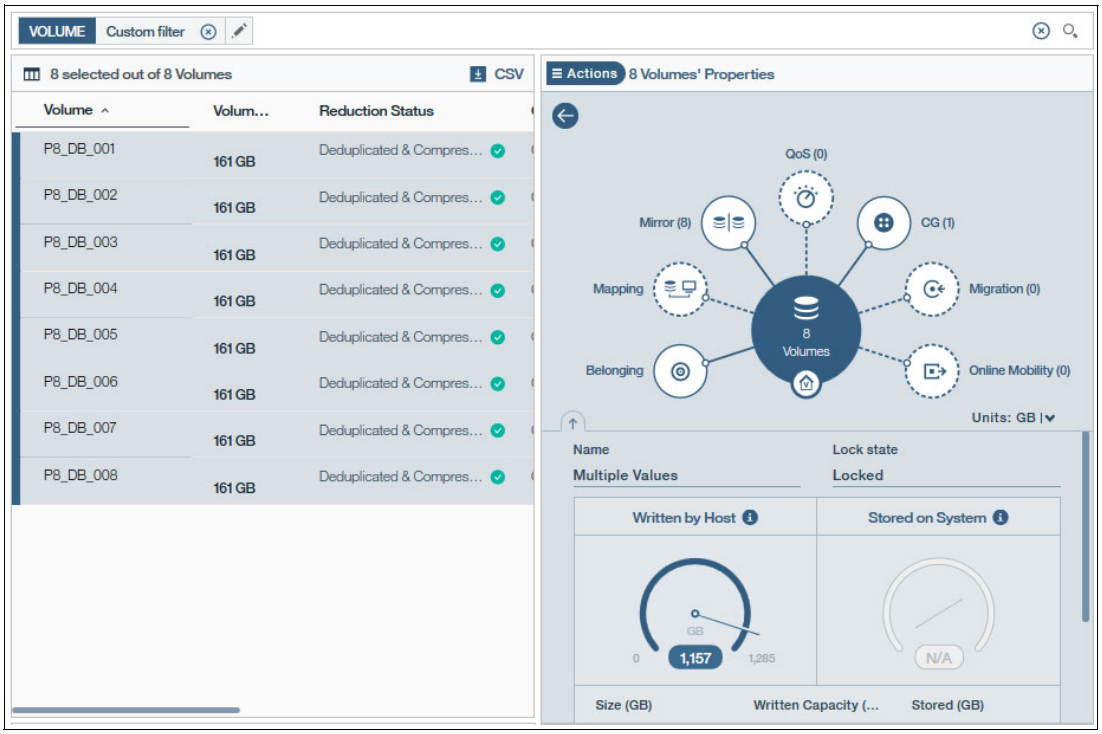

Figure 2-34 Listing of the mirrored volumes in a mirrored consistency group

A mirrored volume can only be in one mirrored CG. Adding a mirrored volume, which is a member of a mirrored CG, to another mirrored CG automatically removes the volume from the previous mirrored CG.

It is also important to realize that in a CG, all volumes have the same role, the same domain, and the same pool on each site.

All volumes in a CG have the same role. CGs are handled as a single entity, and, in asynchronous mirroring, a delay in replicating a single volume affects the status of the entire CG.

2.7 Migrating data using remote mirroring

Remote mirroring can be used to migrate data to a newly installed IBM FlashSystem A9000 and A9000R system if the source data is on an IBM FlashSystem A9000 or A9000R system. This process is basically the same as described in 2.5, “Remote mirroring scenarios” on page 42:

- Define the target system and IBM FlashSystem A9000 and A9000R connectivity
- Define mirror coupling peer volumes
- Create a mirrored consistency group
- Add volumes to the mirrored consistency group

  Adding the volumes to a consistency group is optional but doing so is convenient because the set of volumes can be treated as a group, simplifying the following processes.
- Wait for synchronization to complete
- Switch the roles of the volumes

The migrated volumes are then available for Host activity with a backup copy at the original site. When mirroring is active and synchronized (consistent), the source and destination roles of mirrored volumes or consistency groups can be switched simultaneously. Role switching is typical for returning mirroring to the normal direction after changes have been mirrored in the reverse direction after a production site switch. Role switching is also typical for any planned production site switch.

Host server write activity and remote mirroring activity must be paused briefly before and during the role switch. Additionally, in the case of asynchronous mirroring, at least one sync job must complete before the switch to ensure the expected point-in-time copy of the data exists. Figure 2-35 shows the action path for switching the role.

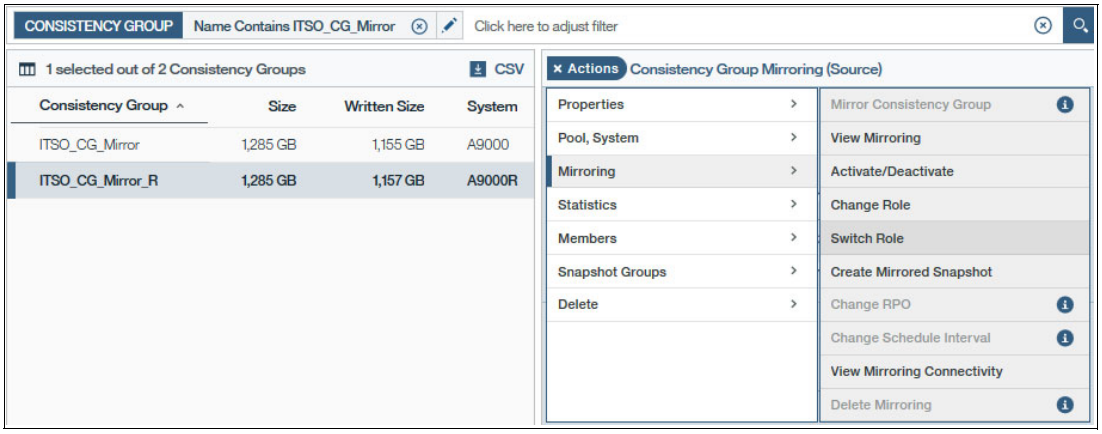

Figure 2-35 Selecting a consistency group for switching role

Select **Switch Role** for the current and new role information.

Figure 2-36 Switching the mirroring role

Click **Apply** to switch the role, as shown in Figure 2-37 on page 56.

Figure 2-37 Switched roles.

The volumes are now active on the secondary IBM FlashSystem A9000 and A9000R system and mirroring is active to the original source system. Hosts can now access the volumes at the secondary site and any changes are mirrored to the original source volumes.

2.8 Recovering from a site failure

One of the reasons for remote mirroring is to keep a copy of data in a different location in case something happens to the source site. In the case of a real disaster, the systems at the local site might need to be replaced. In such circumstances, the production activity moves to the remote site and, when the local site is restored, data can be migrated back as described in 2.7, “Migrating data using remote mirroring” on page 54.

2.8.1 System failure at the local site

Figure 2-38 shows a connectivity failure between the source system and the destination system. The mirror definitions still exist on the target system and the role of those volumes can be changed to source.

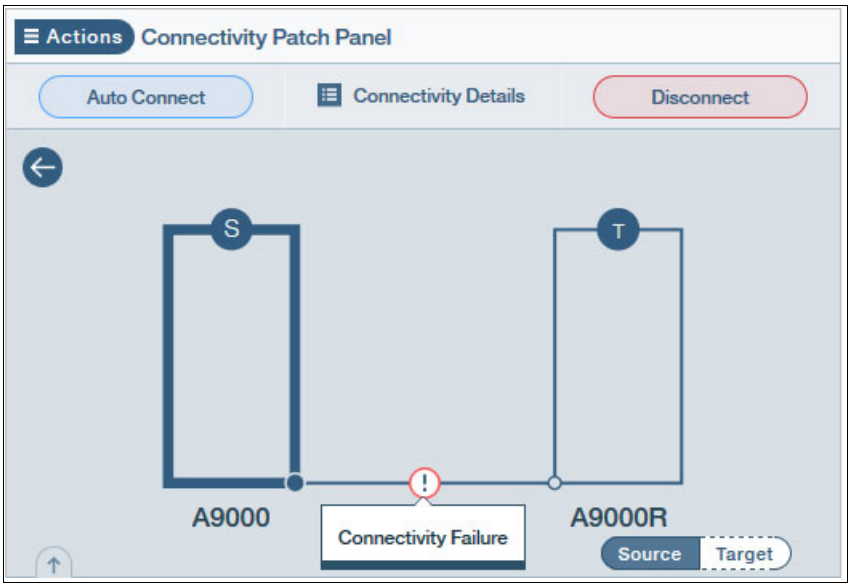

Figure 2-38 Local site system failure

Figure 2-39 shows the details of the mirrored volume status after the connectivity between the source system to the end destination system that failed.

1 selected out of 2 Volumes

Volume	Volume Size	Mirror Status	Locked	Link Status	Mirror Role
P8_DB_001	161 GB	Unsynchronized (Link ...	Yes	Disconnected	Destination
P8_DB_001	161 GB	Unsynchronized (Link ...	No	Disconnected	Source

Figure 2-39 Unsynchronized mirrored volume after link failure

Figure 2-40 on page 58 shows the details of the connectivity between the source and target system. In this example, all connections between the A9000R (the source system) and the A9000 (destination system) failed.

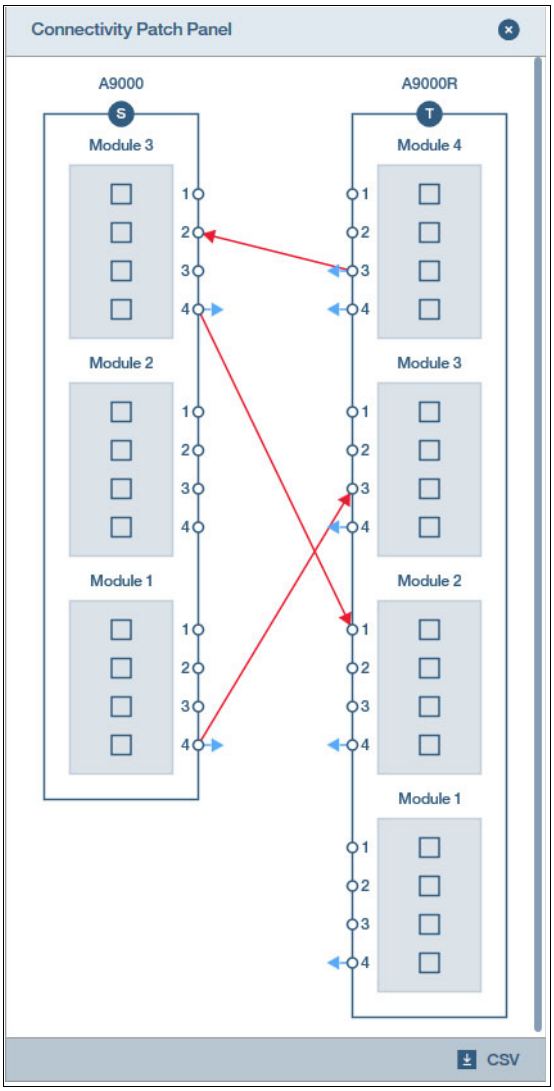

Figure 2-40 Connectivity to the destination site has failed

2.8.2 Changing role of destination volume or CG

The system at the secondary site needs to be prepared to become the production system. To do this, the role of the target volumes are changed to source volumes, and then host systems are able to access the last consistent copy of the production data.

When mirroring is active, the destination volume or CG is locked and write access is prohibited. To allow write access to a destination peer in case of failure or unavailability of the source, the destination volume role must be changed to the source role. The mirroring status is unsynchronized.

In unsynchronized state, the switch role function, as shown for a CG in Figure 2-36 on page 55, is not available. You must use the change role function as shown in Figure 2-41 on page 59. In this example the role of a mirrored volume is changed.

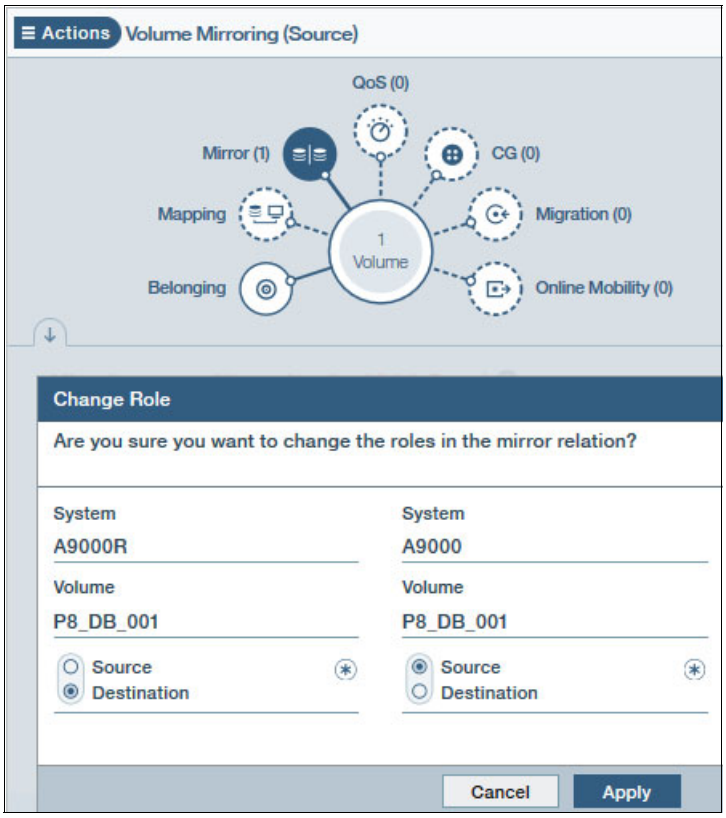

Figure 2-41 Changing role of destination volume or CG

Click **Apply** to change the role of the source volume, as shown in the error message in Figure 2-42. Only the destination role is changed.

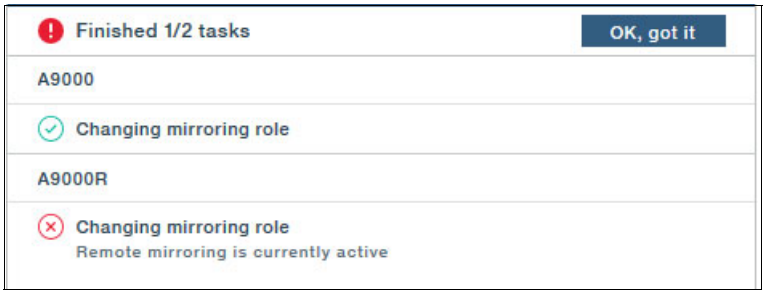

Figure 2-42 Only destination role is changed to source role

Changing the role of a volume from destination to source allows the volume to be accessed. In synchronous mirroring, changing the role also starts metadata recording for any changes that are made to the volume. This metadata can be used for resynchronization (if the new source volume remains the source when remote mirroring is reactivated). In asynchronous mirroring, changing a peer's role automatically reverts the peer to its last-replicated snapshot.

The volume on the secondary site has changed its role from destination to source. The secondary volume is now accessible from a host.

2.9 Deleting mirror coupling definitions

When a mirror coupling is deleted, all metadata and mirroring definitions are deleted, and the peers do not have any relationship at all (Figure 2-43 on page 62). However, any volume and consistency group mirroring snapshots remain on the local and remote systems. To restart mirroring, you can use offline initialization instead of a full copy of data.

If you want to delete the existing mirroring between two IBM FlashSystem A9000 or A9000R systems, you must delete the volume or CG mirroring and its mirrored volumes first. In case the link between the systems is broken, you have to use the XCLI commands on both systems, as shown in Example 2-1, for a mirrored volume.

Example 2-1 Deleting mirroring

```
# Deactivate and delete the mirror on the source system
A9000R>>mirror_deactivate vol=P8_DB_001
Warning:   Are you sure you want to deactivate mirroring? y/n: y
Command executed successfully.
A9000R>>mirror_delete vol=P8_DB_001
Warning:   Are you sure you want to delete this mirroring relationship? y/n: y
Command executed successfully.

# Deactivate and delete the mirror on the destination system
A9000>>mirror_deactivate vol=P8_DB_001
Warning:   Are you sure you want to deactivate mirroring? y/n: y
Command executed successfully.
A9000>>mirror_delete vol=P8_DB_001
Warning:   Are you sure you want to delete this mirroring relationship? y/n: y
Command executed successfully.
```

The connectivity can be deleted using XCLI commands, as shown in Example 2-2.

Example 2-2 Mirroring connectivity deletion

```
# Delete connectivity and ports on the source system
============================================================
A9000R>>target_connectivity_list
Target Name   Remote Port        FC Port         IP Interface   Active   Up
A9000         5001738056730131   1:FC_Port:4:3                  yes      no
A9000         5001738056730113   1:FC_Port:1:3                  yes      no
A9000R>>target_connectivity_delete target=A9000 fcaddress=5001738056730131
local_port=1:FC_Port:4:3
Command executed successfully.
A9000R>>target_connectivity_delete target=A9000 fcaddress=5001738056730113
local_port=1:FC_Port:1:3
Command executed successfully.

A9000R>>target_port_list
Target Name   Port Type   Active   WWPN              iSCSI Address   iSCSI Port
A9000         FC          yes      5001738056730131                  0
A9000         FC          yes      5001738056730113                  0
A9000R>>target_port_delete target=A9000 fcaddress=5001738056730131
Command executed successfully.
A9000R>>target_port_delete target=A9000 fcaddress=5001738056730113
Command executed successfully.
```

```
# Delete connectivity and ports on the destination system
A9000>>target_connectivity_list
Target Name   Remote Port        FC Port         IP Interface   Active   Up
A9000R        5001738051A60132   1:FC_Port:1:4                  yes      yes
A9000R        5001738051A60120   1:FC_Port:3:4                  yes      yes
A9000>>target_connectivity_delete target=A9000R fcaddress=5001738051A60132
local_port=1:FC_Port:1:4
Command executed successfully.
A9000>>target_connectivity_delete target=A9000R fcaddress=5001738051A60120
local_port=1:FC_Port:3:4
Command executed successfully.

A9000>>target_port_list
Target Name   Port Type   Active   WWPN               iSCSI Address   iSCSI Port
A9000R        FC          yes      5001738051A60120                   0
A9000R        FC          yes      5001738051A60132                   0
A9000>>target_port_delete target=A9000R fcaddress=5001738051A60120
Command executed successfully.
A9000>>target_port_delete target=A9000R fcaddress=5001738051A60132
Command executed successfully.
```

The last step is to delete the target definition on the source and the destination system, as illustrated in Example 2-3.

Example 2-3 Delete mirror target definition

```
# Delete targeton the source system
A9000R>>target_delete target=A9000
Command executed successfully.

# Delete targeton the destination system
A9000>>target_delete target=A9000R
Command executed successfully.
```

Typical usage of mirror deletion is also a one-time data migration using remote mirroring. This process includes deleting the mirror couplings after the migration is complete.

2.10 Recovery from a temporary site outage

A temporary site outage might require moving production host activity to the remote location. This can be accomplished using the procedure for recovering from a site failure. In this case, the local system is restored and mirroring can be resumed.

Because the most current data is at the remote site, the changes must be copied to the original site. The remote site had its role changed to the source role. Therefore, the local system must be changed to a target role. When changing the destination to a source, the former source must first be changed to the destination role (upon recovery of the primary site) *before* changing the secondary role back from source to destination.

In synchronous mirroring, changing a peer role from source to destination allows the destination to accept mirrored data from the source. It also causes deletion of metadata that was used to record any changes when the peer had the source role.

In asynchronous mirroring, changing a peer's role automatically reverts the peer to its last-replicated snapshot.

This scenario shows how to use a consistency group temporary on the secondary site. These three steps are shown:

- Changing the consistency group role of the secondary, the destination site to source.

  Hosts can then access the volumes of the consistency group for reading and writing.
- Resynchronization of the consistency group.

  The date written to the secondary site is mirrored to the primary site. The volumes of the mirrored consistency group are synchronized again.
- Changing the consistency group role of the secondary, the current source site to destination.

  The primary site is then source site again.

The sections that follow describe the Hyper-Scale Manager commands for a mirrored consistency group. Equivalent Hyper-Scale Manager commands exists for mirrored volumes, which are not members of a consistency group.

2.10.1 Changing the role of a consistency group

Because there is no connectivity to the primary site, the **`switch role`** command cannot be used to switch both roles. Therefore, the **`change role`** command is used to change the volumes to the source role. Because the volumes are in a mirrored consistency group, the **`Change Role`** command of the consistency group can be used, as shown in Figure 2-43.

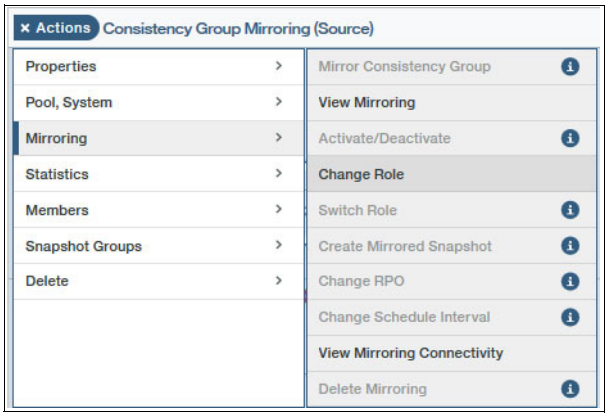

Figure 2-43 Consistency group change role command

Change the role of the primary system to source and the role of the secondary system to destination, as shown in Figure 2-44 on page 63.

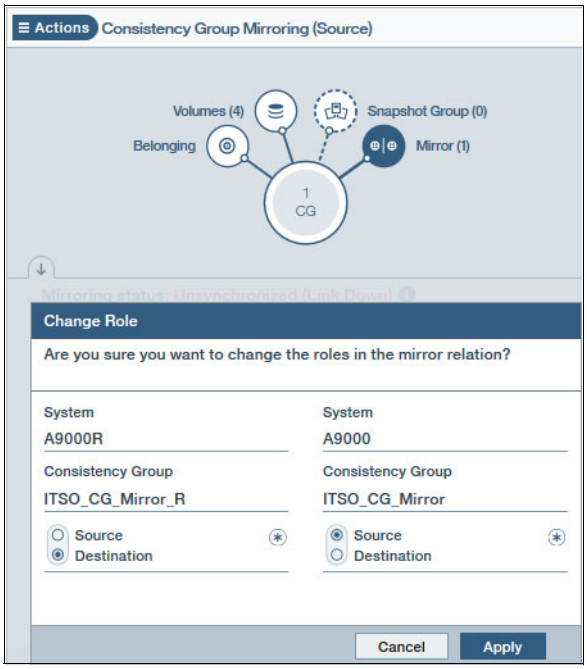

Figure 2-44 Changing the role of the consistency group

Click **Apply** to start the role changing. Because the primary site is not connected, only the secondary system's role is changed, as shown in the messages of Figure 2-45.

Figure 2-45 Changing only the secondary CG role

Both peers might temporarily have the source role when a failure at the primary site results in a disaster recovery production site switch from primary to secondary site.

Map the volumes of the secondary system to the host. The host can now use the latest mirrored data.

You can lock the volumes of the CG on the primary site to prevent any access to them in case the link between both sites is up again.

2.10.2 Reactivation and resynchronization

When the primary site becomes available again and you can switch production back to the primary site, the production changes made to the volumes at the secondary site must be resynchronized to the volumes at the primary site. After the link between the primary and secondary site is re-established, you must change the role of the primary site to destination, as described in Figure 2-44. The next step is to activate the consistency group, as shown in Figure 2-46.

Figure 2-46 Re-activating the consistency group

After re-synchronizing the volumes, the consistency group have the *synchronized* state, as shown in Figure 2-47.

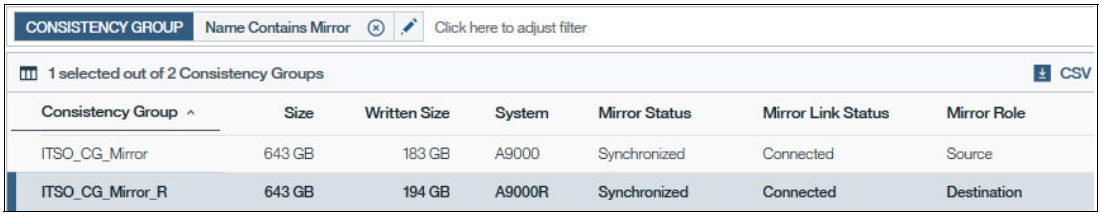

Figure 2-47 CG is synchronized and secondary site is the source site

2.10.3 Changing the role back to source of the primary site

To change the primary site role back to source, use the switch role action, as shown in Figure 2-48.

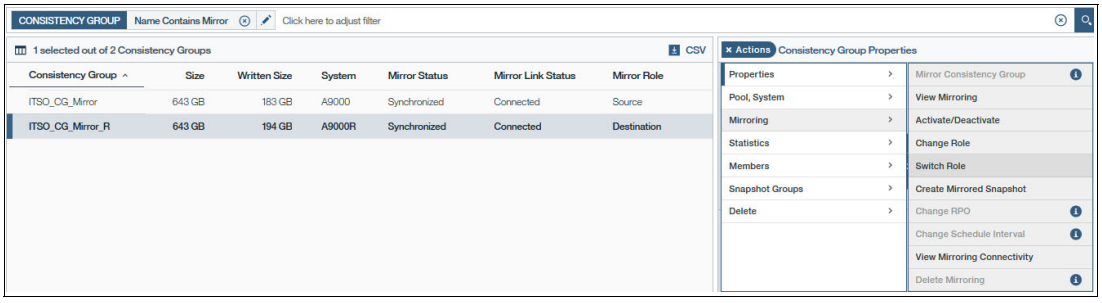

Figure 2-48 Switch role command for consistency group

Click **Apply** in Figure 2-49 to switch the roles of the primary and secondary site.

Figure 2-49 Switch role details

The primary site now has the source role again, as shown in Figure 2-50.

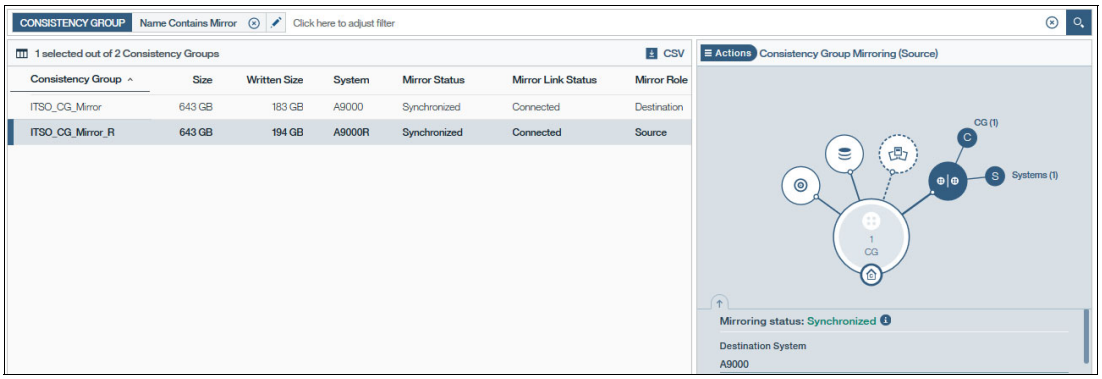

Figure 2-50 The primary system has the source role again.

The volumes of the consistency group now have the same roles as before the site failure occurred. The primary system has the source role, as shown in Figure 2-51.

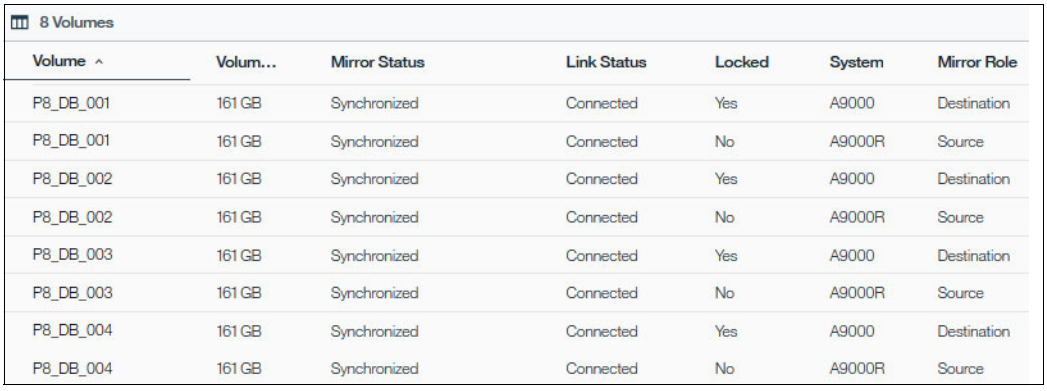

Figure 2-51 Volume mirroring roles back to the roles they had before the site failure

2.11 Temporary deactivation of mirroring

The following examples are of mirror deactivation and reactivation in the same direction:

- Remote mirroring is temporarily deactivated because of a communication failure and then automatically reactivated by the system when communication is restored.
- Remote mirroring is temporarily deactivated to create an extra copy of consistent data at the secondary site.
- Remote mirroring is temporarily deactivated by user action during peak load in an environment with constrained network bandwidth.
- Planned deactivation of remote mirroring can be done to suspend remote mirroring during a planned network outage or DR test.

Consistency group and volume state after link failure

Figure 2-52 shows the mirror status *unsychronized* of a consistency group after the link between of the source and destination system failed. In this example, ITSO_CG_Mirror_R is the source CG and ITSO_CG_Mirror the destination CG.

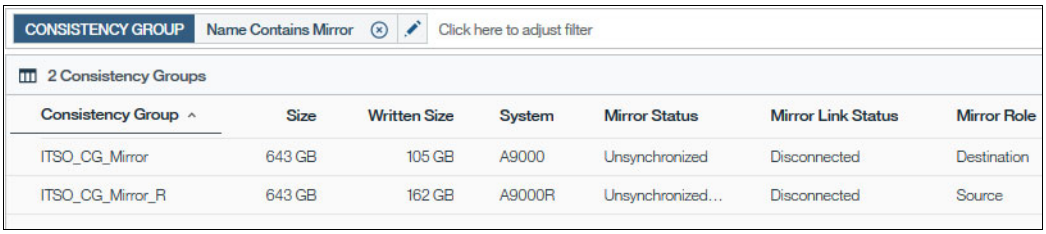

Consistency Group	Size	Written Size	System	Mirror Status	Mirror Link Status	Mirror Role
ITSO_CG_Mirror	643 GB	105 GB	A9000	Unsynchronized	Disconnected	Destination
ITSO_CG_Mirror_R	643 GB	162 GB	A9000R	Unsynchronized…	Disconnected	Source

Figure 2-52 Unsynchronized consistency group after link failure

Figure 2-53 shows the mirror status *unsychronized* of the volumes of the consistency group after the link between of the source and destination system failed.

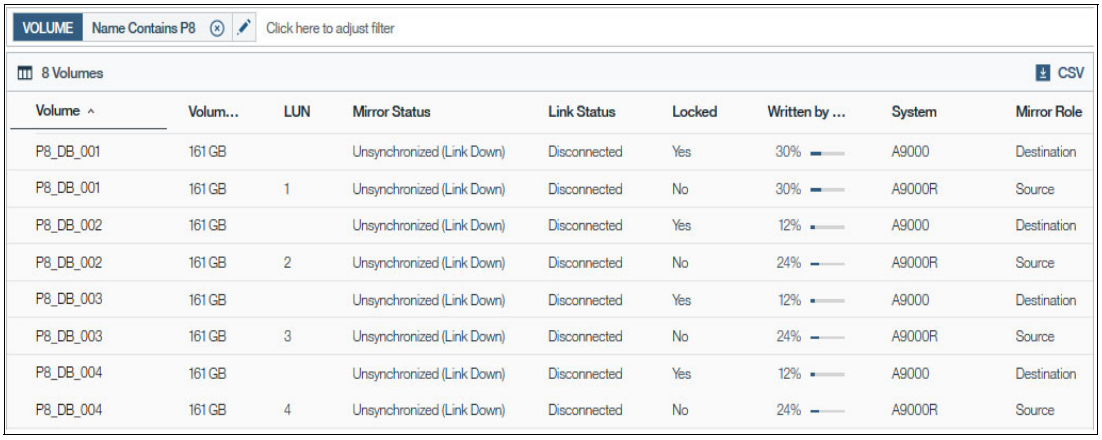

Volume	Volum…	LUN	Mirror Status	Link Status	Locked	Written by …	System	Mirror Role
P8_DB_001	161 GB		Unsynchronized (Link Down)	Disconnected	Yes	30%	A9000	Destination
P8_DB_001	161 GB	1	Unsynchronized (Link Down)	Disconnected	No	30%	A9000R	Source
P8_DB_002	161 GB		Unsynchronized (Link Down)	Disconnected	Yes	12%	A9000	Destination
P8_DB_002	161 GB	2	Unsynchronized (Link Down)	Disconnected	No	24%	A9000R	Source
P8_DB_003	161 GB		Unsynchronized (Link Down)	Disconnected	Yes	12%	A9000	Destination
P8_DB_003	161 GB	3	Unsynchronized (Link Down)	Disconnected	No	24%	A9000R	Source
P8_DB_004	161 GB		Unsynchronized (Link Down)	Disconnected	Yes	12%	A9000	Destination
P8_DB_004	161 GB	4	Unsynchronized (Link Down)	Disconnected	No	24%	A9000R	Source

Figure 2-53 Unsynchronized volumes after link failure

2.11.1 Consistency group and volume state after the link is restored

After the link is re-established again the volumes automatically re-synchronize, as shown in Figure 2-54.

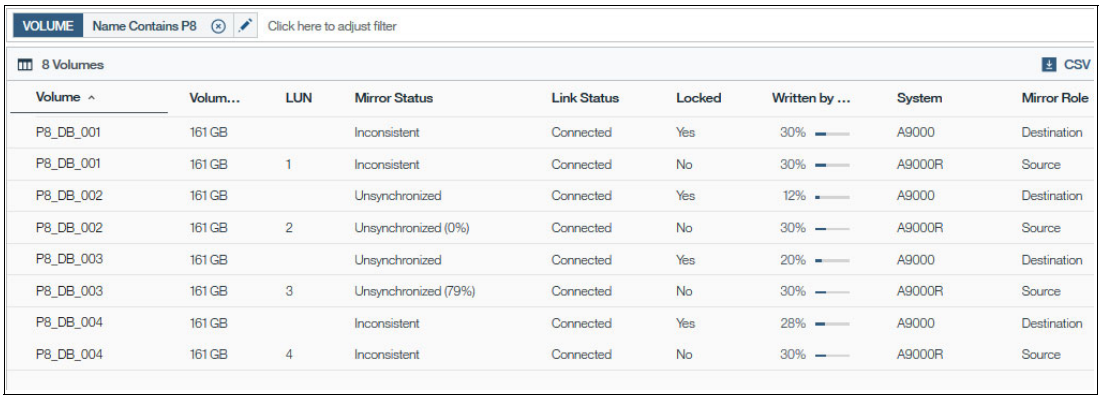

Volume	Volum...	LUN	Mirror Status	Link Status	Locked	Written by ...	System	Mirror Role
P8_DB_001	161 GB		Inconsistent	Connected	Yes	30%	A9000	Destination
P8_DB_001	161 GB	1	Inconsistent	Connected	No	30%	A9000R	Source
P8_DB_002	161 GB		Unsynchronized	Connected	Yes	12%	A9000	Destination
P8_DB_002	161 GB	2	Unsynchronized (0%)	Connected	No	30%	A9000R	Source
P8_DB_003	161 GB		Unsynchronized	Connected	Yes	20%	A9000	Destination
P8_DB_003	161 GB	3	Unsynchronized (79%)	Connected	No	30%	A9000R	Source
P8_DB_004	161 GB		Inconsistent	Connected	Yes	28%	A9000	Destination
P8_DB_004	161 GB	4	Inconsistent	Connected	No	30%	A9000R	Source

Figure 2-54 Re-synchronizing of mirrored volumes

After the volumes are resynchronized, the consistency group is in the synchronized state again, as shown in Figure 2-55.

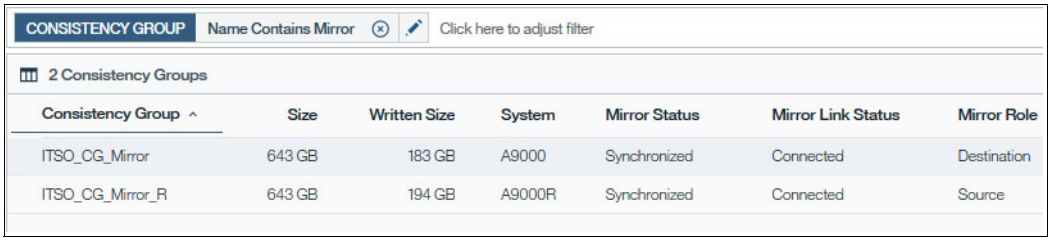

Consistency Group	Size	Written Size	System	Mirror Status	Mirror Link Status	Mirror Role
ITSO_CG_Mirror	643 GB	183 GB	A9000	Synchronized	Connected	Destination
ITSO_CG_Mirror_R	643 GB	194 GB	A9000R	Synchronized	Connected	Source

Figure 2-55 Resynchronized consistency group

2.11.2 Deactivating mirror coupling: Change recording

A mirror coupling can be deactivated by a user action. Use the **`Activate/Deactivate`** command in the Hyper-Scale Manager, as shown in Figure 2-56.

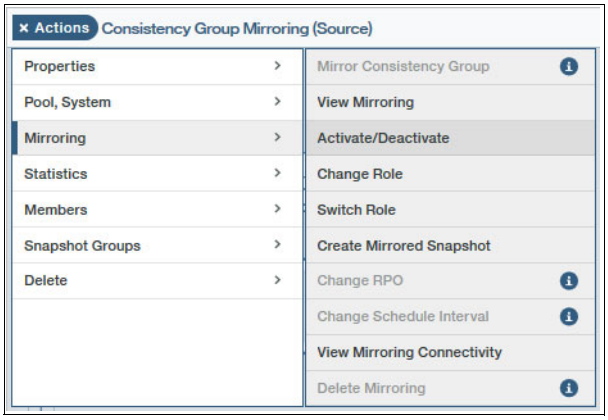

Figure 2-56 Activating or deactivating a mirrored consistency group.

Check **Activation State: Inactive**. Click **Apply** to change the state.

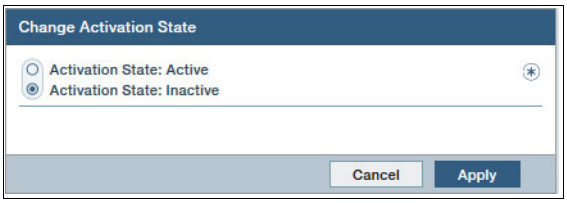

Figure 2-57 Changing the activation state of a CG

In this case, the mirror of the consistency group changes to *inactive* state, as shown in Figure 2-58. All volumes of this consistency group also have the *inactive* state.

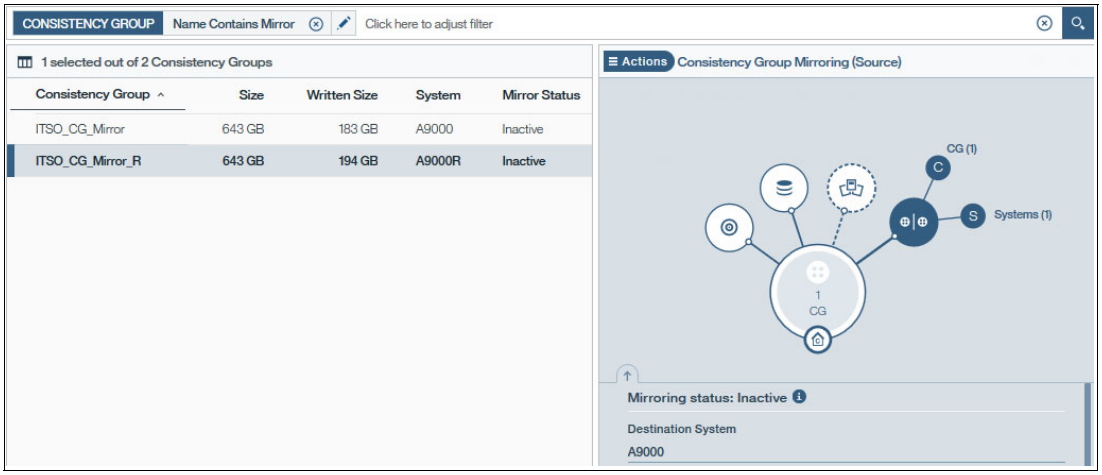

Figure 2-58 Deactivated mirroring of a consistency group

During the inactive state, a consistent set of data is available at the destination site. The currency of the consistent data ages in comparison to the source volumes, and the gap increases when mirroring is in standby mode.

In synchronous mirroring, during standby mode, the storage system metadata is used to note which parts of a source volume are changed but have not yet been replicated to the destination volume (because mirroring is not currently active).

When synchronous mirroring is reactivated by a user command or communication is restored, the metadata is used to re-synchronize changes from the source volumes to the destination volumes. Mirroring records changes for source volumes only. If you want to record changes to both peer volumes when mirroring is in standby mode, change the destination volume to a source volume.

In asynchronous mirroring, metadata is not used and the comparison between the most-recent and last-replicated snapshots indicates the data that must be replicated.

2.11.3 Mirror reactivation and resynchronization: Normal direction

In synchronous mirroring, when mirroring has been in standby mode, any changes to volumes with the source role are recorded in metadata. When mirroring is reactivated, changes recorded in metadata for the current source volumes are resynchronized to the current destination volumes.

This is the same behavior as the re-synchronization after a link failure, as shown in Figure 2-54 and Figure 2-55. The rate for this resynchronization of changes can be specified by the user in MBps using the XCLI **`target_config_sync_rates`** command or the **`Max Resync Rate`** command in the Hype-Scale Manager System Connectivity pane, as shown in Figure 2-59.

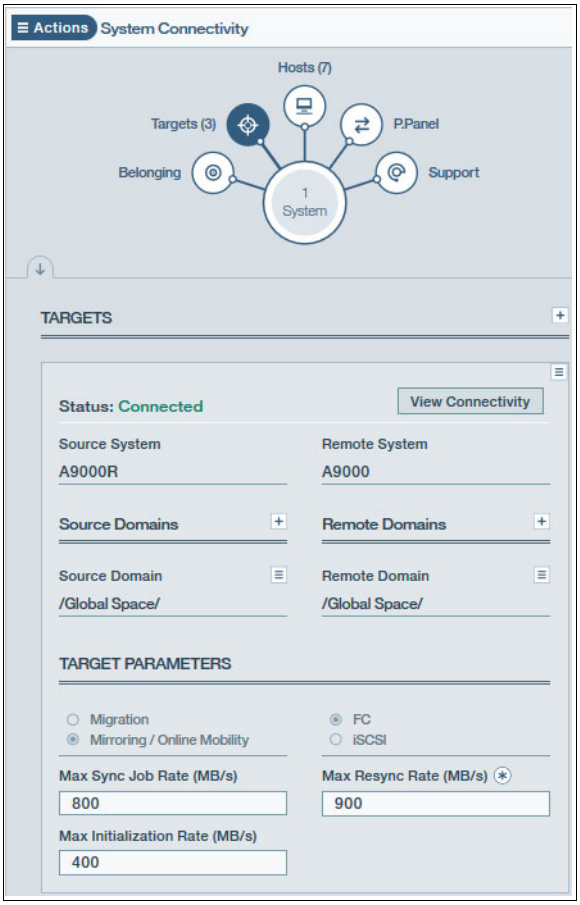

Figure 2-59 Changing the resynchronization rate

When mirroring is reactivated in the normal direction, changes recorded at the primary peers are copied to the secondary peers.

3

IBM Hyper-Scale Mobility with IBM FlashSystem A9000

IBM Hyper-Scale Mobility is a powerful function for moving volumes between IBM FlashSystem A9000 systems in a manner that is transparent to host applications.

This chapter includes the following sections:

- Introduction to IBM Hyper-Scale Mobility
- IBM Hyper-Scale Mobility design considerations
- IBM Hyper-Scale Mobility requirements
- IBM Hyper-Scale Mobility process description
- Usage prerequisites and considerations
- Management software support for IBM Hyper-Scale Mobility
- Using IBM Hyper-Scale Mobility

3.1 Introduction to IBM Hyper-Scale Mobility

Introduced with IBM FlashSystem A9000 Version 12.0.1, IBM Hyper-Scale Mobility is a feature of the IBM FlashSystem A9000 system.

IBM Hyper-Scale Mobility enables clients to move a volume from one IBM FlashSystem A9000 to another (over synchronous distances) in a manner that is transparent to host applications before, during, and after the volume migration.

Note: IBM Hyper-Scale Mobility is a no-charge feature that is included with IBM FlashSystem A9000 Version 12.0.1 and later.

IBM Hyper-Scale Mobility helps you overcome provisioning scenarios that normally challenge traditional systems. It can accommodate several critical client needs in the modern data center and cloud environment, including online data mobility, load balancing, over-provisioning, and storage system repurposing.

This volume migration capability greatly enhances IBM FlashSystem A9000 scalability and directly addresses several client storage-related issues:

- Managing storage growth
- Providing more flexibility in capacity forecasting
- Managing costs
- Scaling between IBM FlashSystem A9000 systems
- Balancing workloads across systems without service downtime for better performance
- Retiring systems gracefully and upgrading to the latest technology
- Separating storage and host maintenance cycles

IBM Hyper-Scale Mobility capabilities also provide IBM FlashSystem A9000 with improved total cost of ownership (TCO) and improved service continuity.

Consider how IBM Hyper-Scale Mobility can address the following client situations:

- Migrating all the data from a IBM FlashSystem A9000 that will be decommissioned or redeployed for a new purpose.
- IBM Hyper-Scale Mobility can be used to rebalance the workload, away from a system that is heavily used to a system that is under used. It does not matter if the capacity imbalance is performance-related or data storage-related. IBM Hyper-Scale Mobility can help resolve both situations.

3.2 IBM Hyper-Scale Mobility design considerations

IBM Hyper-Scale Mobility is designed to provide a process to move volumes between one IBM FlashSystem A9000 to another IBM FlashSystem A9000 with little host effect.

Moving volumes usually requires defining a volume on the destination (or target) IBM FlashSystem A9000, porting data, and disrupting host applications when the new volume is activated. The IBM Hyper-Scale Mobility architecture enables these basic steps to be completed with minimum host involvement and with no disruption to host applications.

The objective is to move the data to a new volume that is on another IBM FlashSystem A9000 at the same time enabling the host to view this new volume as though it is the original. This task is accomplished by redirecting input/output (I/O) activity automatically (by using a proxy method) from the source IBM FlashSystem A9000 to the destination IBM FlashSystem A9000 without needing changes to the host configuration.

The new volume must look to the host as though it is the original, which is accomplished by duplicating the volume characteristics of the original volume (for example, the worldwide name of the two volumes is identical).

The final steps of the migration require some host intervention to establish paths directly to the new volume and to remove the paths to the original volume. This task, however, does not significantly disrupt host application activity. The host I/Os never need to be interrupted during the whole migration process.

For a detailed, step-by-step illustration of the online migration process, using either the IBM Hyper-Scale Manager or the XCLI, see 3.7, “Using IBM Hyper-Scale Mobility” on page 78.

3.3 IBM Hyper-Scale Mobility requirements

To cause minimal disruption to the host, and to enable uninterrupted I/O operations, these requirements must be met:

- Multi-path driver

  The process of moving a volume concurrently with host access requires that the host operating system uses a multi-path driver to access the volume that is being moved. During the migration, new paths are added, and the old paths to the original volume are eventually removed.

- Switch zoning

  The source and the destination IBM FlashSystem A9000 must both be accessible to the host system concurrently. This might require changes to the zoning in use by the host that is involved. The two IBM FlashSystem A9000 systems also need zoning between each other to enable connections for the volume data migration.

- IBM FlashSystem A9000 systems

  There must be two IBM FlashSystem A9000 systems: One system is the source, and the other system is the target (destination).

  The new volume is automatically created on the destination system in a specified storage pool, so sufficient space must be available for this action. Consideration must also be given for any possible increase in snapshot space that might be required on the destination system.

3.4 IBM Hyper-Scale Mobility process description

From a design standpoint, the IBM Hyper-Scale Mobility process that is used to move a volume between IBM FlashSystem A9000 systems can be summarized in a sequence of stages. The process is characterized by phases and the corresponding states of the migrated volume. The phases and states are depicted in Figure 3-1 on page 74.

Volume mobility between the source system and the destination system does not interrupt host activity, and is not disrupted by rebuild, redistribution, phase-out, or failover on the destination.

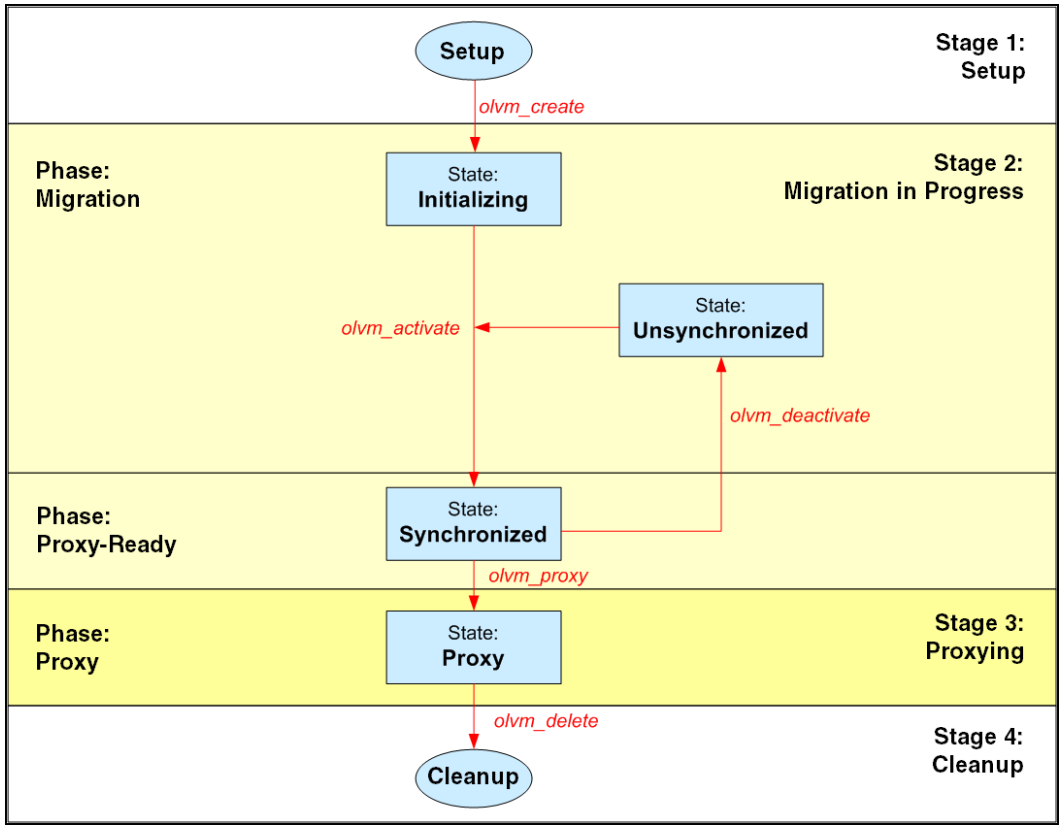

Figure 3-1 Volume migration flow of the IBM Hyper-Scale Mobility

Stage 1: Setup

The first stage of the IBM Hyper-Scale Mobility is to create the volume on the destination IBM FlashSystem A9000 and to set up the relationship between the two volumes that are needed for migration by running the **`olvm_create`** command.

Stage 2: Migration in progress

Data migration can begin by running the **`olvm_activate`** command and must complete successfully before the host can access the new volume. This is part of the Migration phase. During migration, new data is written by the host to the source and ported (copied) to the destination.

The source volume state goes from initializing to synchronized if the migration is successful (Proxy_Ready phase). The source volume can enter the unsynchronized state if there is a link disruption or deactivation. You can also set it to this state by running the **`olvm_deactivate`** command.

Stage 3: Proxying

At this point, the administrator can instruct the source storage system to redirect host I/O to the new volume on the destination storage system by running the **`olvm_proxy`** command. This stage corresponds to the Proxy phase.

In proxy mode, the source (now the proxy) no longer functions as a regular volume, and the source A9000 communicates host requests to the destination. At this point, the migration is no longer reversible.

In proxy mode, the host can remain connected to the source without a need to zone and move it to the destination. The host can be moved after the ported volume data is on the destination.

Stage 4: Cleanup

The final stage of the migration involves connecting the host directly to the new volume. This stage might require zoning changes, and the multi-path device driver must discover the new path to the ported volume.

Finally, the original paths can be removed, and the IBM Hyper-Scale Mobility relationship can be deleted by running the **`olvm_delete`** command.

3.5 Usage prerequisites and considerations

In addition to the general requirements for using IBM Hyper-Scale Mobility that are described in 3.3, “IBM Hyper-Scale Mobility requirements” on page 73, there are additional prerequisites and considerations.

3.5.1 Volume and storage pool considerations

There are specific considerations and conditions regarding volume and storage pools selection when you use IBM Hyper-Scale Mobility:

- Volumes that are already part of a mirroring relationship cannot be selected as the source volume with IBM Hyper-Scale Mobility.
- Volumes that are already part of a consistency group cannot be selected as the source volume with IBM Hyper-Scale Mobility.
- Any snapshots that exist on a volume are *deleted* when the IBM Hyper-Scale Mobility process is moved to the Proxy phase.
- The destination storage pool for an IBM Hyper-Scale Mobility relationship must contain enough free space to accommodate the destination volume. Consideration must also be given for any possible increase in snapshot space that might be required on the target system.

3.5.2 Management workstation connectivity

To perform the online migration of an A9000 volume by using the GUI or XCLI, you must verify that both the source and destination A9000 systems are configured and accessible by the GUI or XCLI.

3.5.3 Connectivity between the source and destination A9000 systems

To effectively use IBM Hyper-Scale Mobility, there must be adequate connectivity between the source and destination A9000 systems.

The IBM Hyper-Scale Mobility process uses the same mechanism as remote mirroring to synchronize source and destination volumes. Fundamentally, IBM Hyper-Scale Mobility acts similarly to synchronous remote mirroring during the migration and proxy-ready phases.

When planning IBM Hyper-Scale Mobility activities, follow the same guidelines and leading practices regarding connectivity between source and destination A9000 systems as those followed with synchronous remote mirroring. For more information, refer to Chapter 1, “Remote connectivity with IBM FlashSystem A9000 and IBM FlashSystem A9000R” on page 1.

Note, however, that as made available with A9000 Version 12.0.1, IBM Hyper-Scale Mobility is only available between two A9000 systems; it restricts that the source and target A9000 systems must be connected over Fibre Channel (FC) and be located at a short distance from each other. This is similar to Synchronous Remote Mirroring. Internet Small Computer System Interface (iSCSI) is not supported at this time.

3.5.4 Host system connectivity

Before you use IBM Hyper-Scale Mobility, you must confirm that any host that has mappings to the source volume on the source IBM FlashSystem A9000 has the appropriate access to the destination IBM FlashSystem A9000:

- If you are migrating an iSCSI logical unit number (LUN), ensure that the host can access the destination system over the Internet Protocol (IP) network.
- If you are migrating an FC LUN, ensure that the host is correctly zoned to the destination system.

3.6 Management software support for IBM Hyper-Scale Mobility

IBM Hyper-Scale Manager v5.0.1 and IBM XIV® XCLI v5.0 introduced support for the IBM Hyper-Scale Mobility feature for the IBM FlashSystem A9000. Either tool can be used to perform the tasks that are related to IBM Hyper-Scale Mobility. IBM Hyper-Scale Manager v5.1 is used for the examples and illustrations in this publication.

3.6.1 IBM Hyper-Scale Manager GUI features

In support of the IBM Hyper-Scale Mobility feature, the IBM Hyper-Scale Manager GUI was updated with new Volume Mobility options in the Remote Views menu in the left pane, as shown in Figure 3-2.

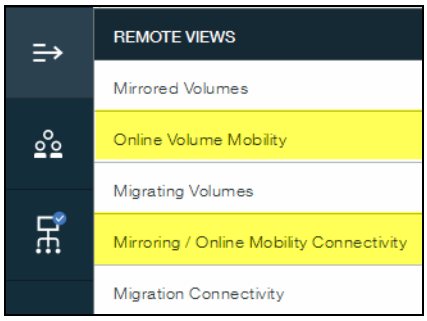

Figure 3-2 Volume Mobility menu options to support IBM Hyper-Scale Mobility

Additionally, several filters have been added to the default filters available when selecting Volumes from the left navigation pane, as shown in Figure 3-3.

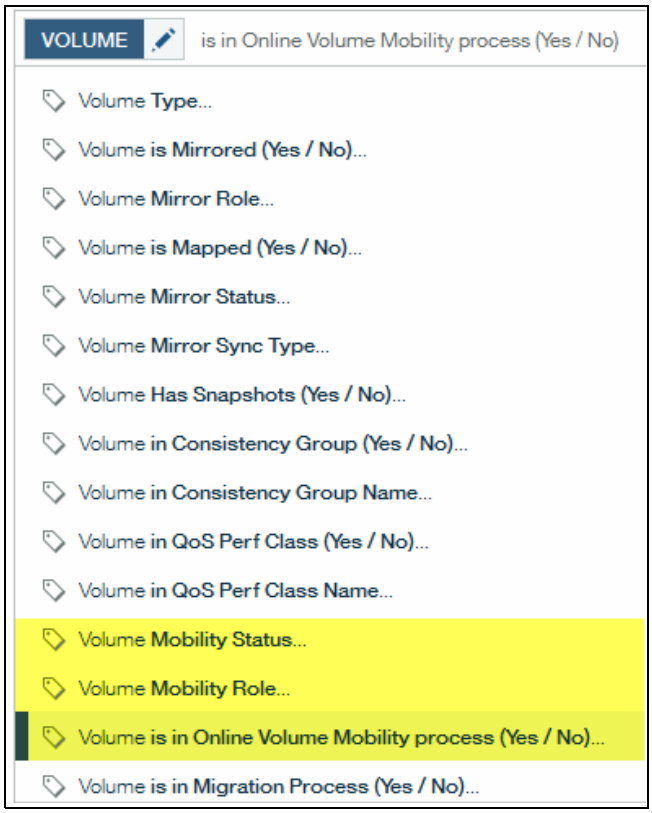

Figure 3-3 Mobility filters in the Volumes menu

Other views and GUI panels were updated to reflect support for IBM Hyper-Scale Mobility as well.

3.6.2 XCLI functions that support IBM Hyper-Scale Mobility

The XCLI includes a series of commands to support IBM Hyper-Scale Mobility:

`olvm_list`	Displays details for IBM Hyper-Scale Mobility relationships on a local system.
`olvm_create`	Defines an IBM Hyper-Scale Mobility configuration.
`olvm_activate`	Activates (starts) an IBM Hyper-Scale Mobility migration for a defined IBM Hyper-Scale Mobility relationship.
`olvm_proxy`	Moves the IBM Hyper-Scale Mobility source volume to a proxy state.
`olvm_deactivate`	Deactivates the IBM Hyper-Scale Mobility migration for a defined IBM Hyper-Scale relationship.
`olvm_abort`	Aborts a defined or activated IBM Hyper-Scale Mobility process.
`olvm_delete`	Deletes an IBM Hyper-Scale Mobility relationship and attributes.

3.7 Using IBM Hyper-Scale Mobility

The following examples depict an online volume migration of a volume by using Hyper-Scale Manager and the XCLI. In these examples, the volume that is being migrated is mapped over Fibre Channel to a Red Hat Enterprise Linux (RHEL) host, and the host is actively generating I/Os to the volume throughout the migration. The source and destination IBM FlashSystem A9000 systems are connected via Fibre Channel according to the connectivity described and outlined in Chapter 1, “Remote connectivity with IBM FlashSystem A9000 and IBM FlashSystem A9000R” on page 1.

Here are the details of the migration examples:

- The source IBM FlashSystem A9000 is named A9000_ITSO.
- The destination IBM FlashSystem A9000 is named A9000 6003308 Jazz™.
- The Hyper-Scale Manager volume to be migrated is named HS_Mob_6003310_02.
- The XCLI volume to be migrated is named HS_Mob_6003310_01.
- The RHEL host with LUN mapping is named mcnode81.
- The number of physical paths from the host to each IBM FlashSystem A9000 is two. Each HBA was zoned to four fibre ports in the source IBM FlashSystem A9000 and zoned to three fibre ports in the destination IBM FlashSystem A9000.

3.7.1 Using the Hyper-Scale Manager to use IBM Hyper-Scale Mobility

This section illustrates how to use IBM Hyper-Scale Mobility by using the IBM Hyper-Scale Manager. We have decomposed the process into a series of six tasks.

Task 1: Planning and preparation

To complete planning and preparation, follow these steps:

1. Confirm that both the source and destination IBM FlashSystem A9000 are configured and accessible in the IBM Hyper-Scale Manager and that there is FC connectivity between these systems. Figure 3-4 shows the source and destination systems with confirmed connectivity between the two systems.

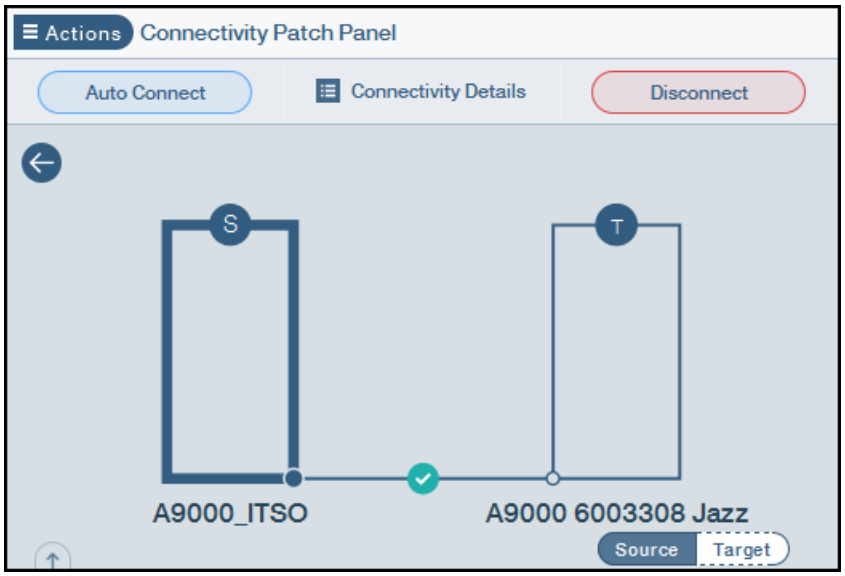

Figure 3-4 Connectivity between two A9000

2. Validate the host multipath connectivity to the source volume by running the **xiv_devlist** command of the Host Attachment Kit (HAK), as shown in Example 3-1.

Example 3-1 Check multipath connectivity

```
[root@mcnode81 ~]# xiv_devlist
XIV Devices
------------------------------------------------------------------------------
Device               Size (GB)  Paths  Vol Name           Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------------
/dev/mapper/mpathaqd  1033.6     8/8    HS_Mob_6003310_01  13443   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaqe  1033.6     8/8    HS_Mob_6003310_02  13444   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaqf  1033.6     8/8    HS_Mob_6003310_03  13445   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaqg  1033.6     8/8    HS_Mob_6003310_04  13446   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaqi  1033.6     8/8    HS_Mob_6003310_05  13447   6003310  mcnode81
```

Note the Vol ID of 13444 and XIV ID of 6003310 of volume HS_Mob_6003310_02, indicating that the volume is presented to the host from the “A9000_ITSO” IBM FlashSystem A9000.

As part of the preparation, it is also possible to define the host on the target IBM FlashSystem A9000. This action can also be completed when the migration is taking place. In this example, the host definition and port definitions are created for the host on the destination IBM FlashSystem A9000 at this time.

3. From the main Hyper-Scale Manager Dashboard, mouse over the **New** icon, click **New Item**, and select **Host** from the menu, as shown in Figure 3-5. This illustration applies to an FC-connected host. For an iSCSI-connected host, see the procedure that is explained in the IBM Redbooks publication, *IBM FlashSystem A9000, IBM FlashSystem A9000R and IBM XIV Storage System: Host Attachment and Interoperability*, SG24-8368.

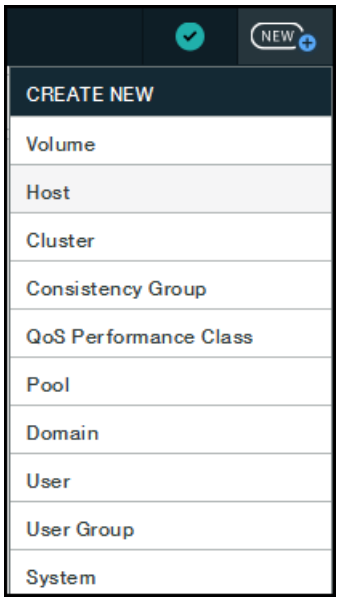

Figure 3-5 Defining a new host

4. In the Add Host panel, enter the details for the host, as shown in Figure 3-6. Ensure that the details match the corresponding host details from the source IBM FlashSystem A9000. In this example, the Name of the host is `mcnode81`, the Type is default, and the Cluster value is empty. The System value is chosen from the drop-down selection and set to the Destination IBM FlashSystem A9000 (A9000 6003308 Jazz in this example).

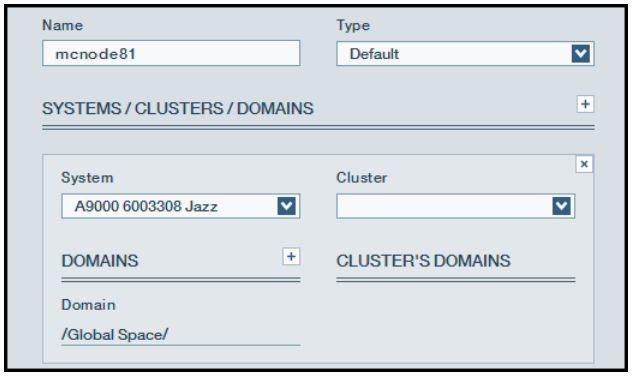

Figure 3-6 Add host details about the destination IBM FlashSystem A9000

5. Add the ports for the host. Click the plus icon (**+**)and fill in the port information, as shown in Figure 3-7. For this example, two ports are added using the WWPN addresses zoned to this host. When completed, click **Create**. Ensure that the details match the corresponding port details from the source IBM FlashSystem A9000. In this example, choose FC for the Port Type; the Port Addresses are 10000090FA731D36 and 10000090FA731D36.

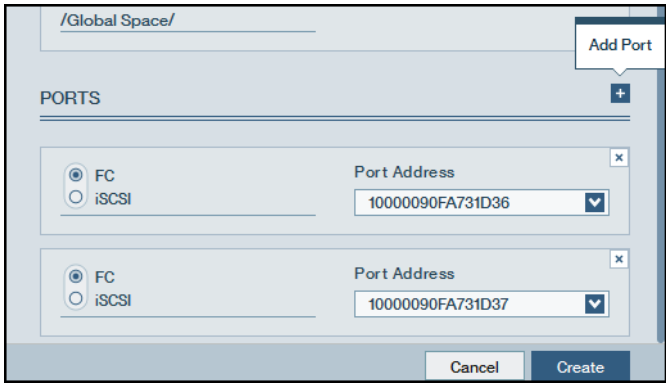

Figure 3-7 Add host ports on the destination IBM FlashSystem A9000

6. The new host and port definitions are now visible in the Hosts tab of the Hyper-Scale Manager display, as shown in Figure 3-8 (shown with the panel display added).

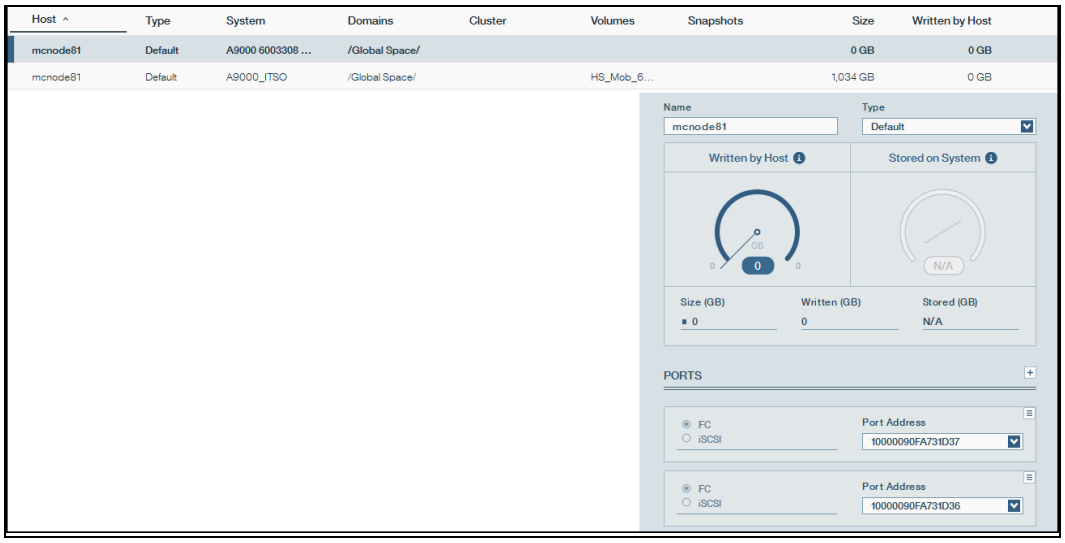

Figure 3-8 Newly defined host and port on the destination IBM FlashSystem A9000

Task 2: Setup

For Setup, complete the following steps:

1. To create the IBM Hyper-Scale Mobility relationship, select the source volume in the Volumes view and click the **Online Mobility** spoke in the Hub display on the right, as shown in Figure 3-9.

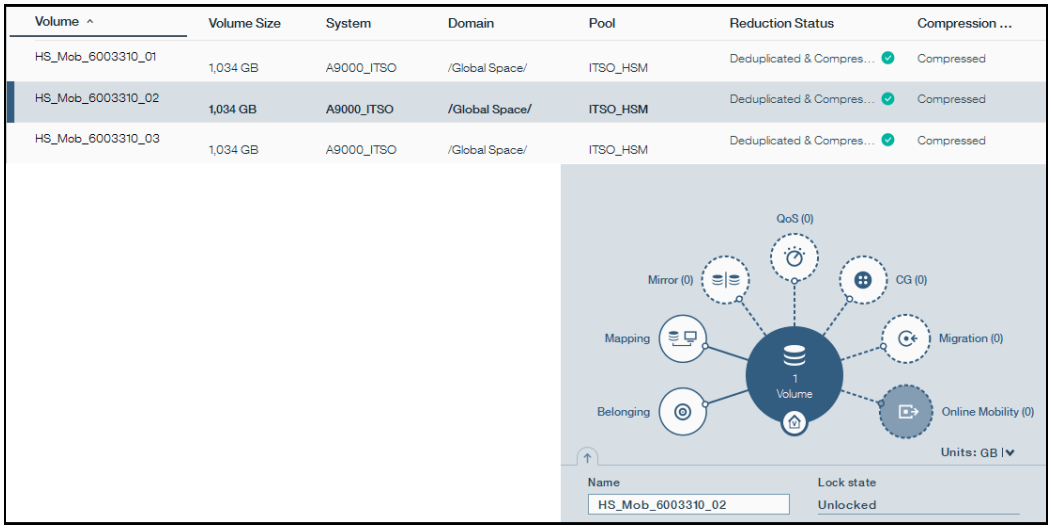

Figure 3-9 Volumes view with Hub Display in Hyper-Scale Manager

2. After selecting the Online Mobility spoke, click the plus sign (**+**) to create a new Online Mobility definition for the source volume. Then fill in the information for the Destination System, as shown in Figure 3-10.

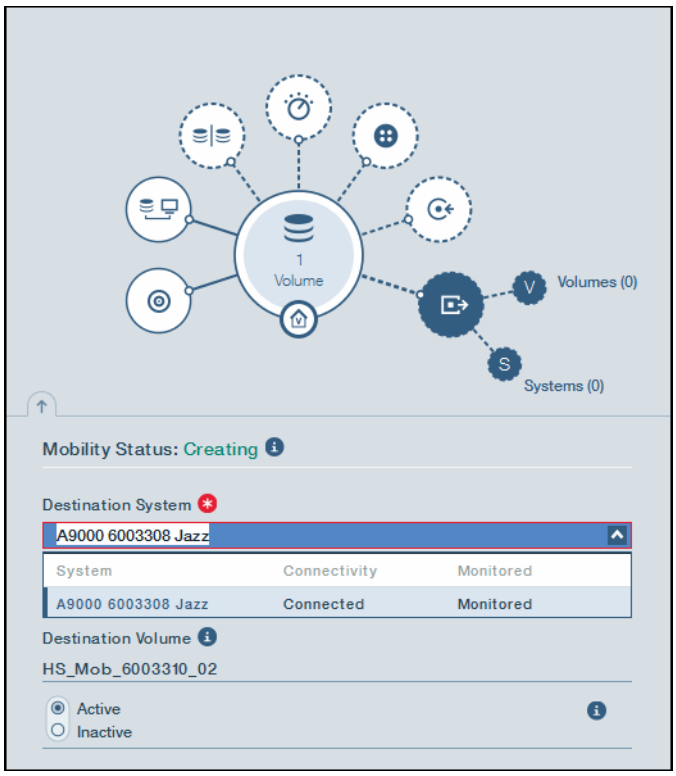

Figure 3-10 Creating a new Online Volume Migration - Destination System

Also, fill in the information for the Destination Pool, as shown in Figure 3-11.

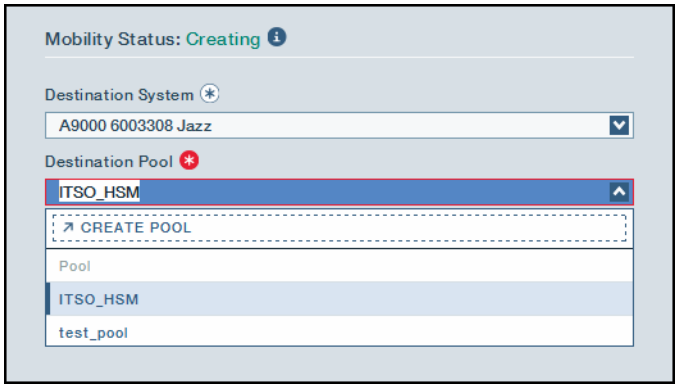

Figure 3-11 Creating a new Online Volume Migration - Destination Pool

Note: A destination Volume is created automatically on the destination system with the same name as the source Volume.

3. To immediately activate the IBM Hyper-Scale Mobility relationship upon creation, select the **Active** radio button. In this example, the Inactive radio button is used, as the next step will explicitly activate the relationship.

   Click **Create** to create the relationship.

Consideration: A volume that is part of a mirror relationship cannot be used as the source volume for IBM Hyper-Scale Mobility without first removing that mirror.

4. By opening a new tab and selecting the Online Volume Mobility menu option, as shown in Figure 3-12, the source and destination volumes can be seen for the newly created migration relationship, as shown in Figure 3-13.

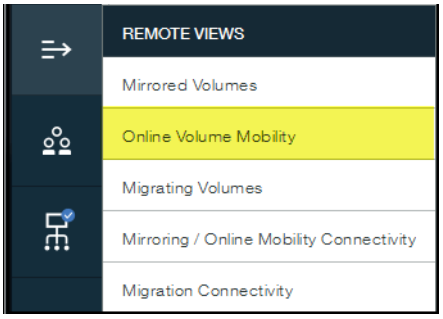

Figure 3-12 Remote Views - Online Volume Mobility

Notice that the Mobility Status for both is showing as Inactive.

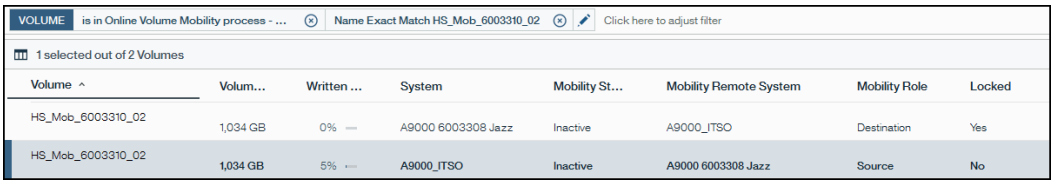

Figure 3-13 Source and Destination volumes for HS_Mob_6003310_02 relationship

Note, also, that the size of the destination volume matches the size of the source volume, but the Written (Used capacity) of the volume is 0 GB. This is expected because the Hyper-Scale Mobility relationship has not yet been activated, and the synchronization of the destination volume has not yet begun. When the Hyper-Scale Mobility relationship is in an inactive state, Host I/O continues to only go to the original (source) Volume, as described in the information pop-up shown in Figure 3-14.

Figure 3-14 Hyper-Scale Mobility - Inactive description

Note: Although the destination volume has been created, it cannot yet be mapped to the host. The destination volume in an IBM Hyper-Scale Mobility relationship cannot be mapped to a host until the Proxy phase. In addition, its Locked Status is Read Only. These are the available actions that you can perform on the volume:

- Move to Pool
- Create Snapshot
- Create Snapshot (Advanced)
- Copy this Volume
- Show Statistics
- Properties

Task 3: Migration

To complete the Migration task, perform the following steps:

1. Select the Volume in the destination IBM FlashSystem A9000 and then select the Online Mobility spoke from the Hub view on the right. From the resulting panel, select **Activate**, as shown in Figure 3-15.

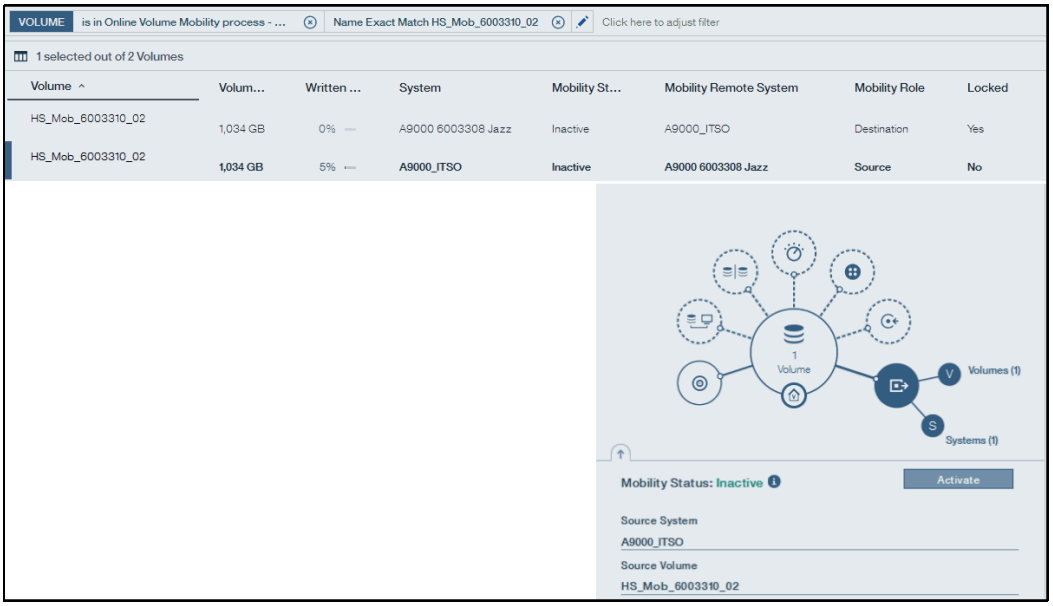

Figure 3-15 Activate the Hyper-Scale Mobility relationship

This action activates the relationship and queues the process to synchronize the destination volume on the destination IBM FlashSystem A9000.

Remember: Only one IBM Hyper-Scale Mobility process at a time is synchronized from the source IBM FlashSystem A9000.

In this example, there are no other IBM Hyper-Scale Mobility relationships on the system, so the synchronization begins immediately. The progress of the synchronization can be viewed in the Volumes tab, as shown in Figure 3-16. The status is initializing and the progress is shown as a percentage. Note that the progress is only shown for the Source volume.

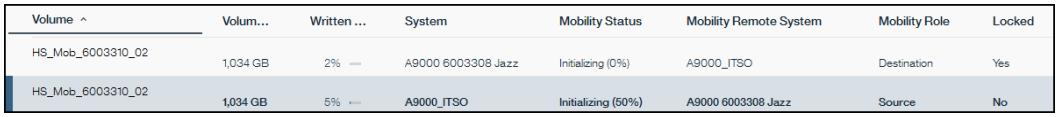

Figure 3-16 IBM Hyper-Scale Mobility status view from the Volumes tab

2. After the synchronization is complete, Mobility Status changes to Synchronized for both volumes. Note that both volumes now report the same Written percentage (Used Capacity). Volume data is mirrored and is consistent on both the source and destination volumes. Host I/O is still only being sent to the original (source) volume in this state, as shown in the information pop-up in the Hub view on the right in Figure 3-17.

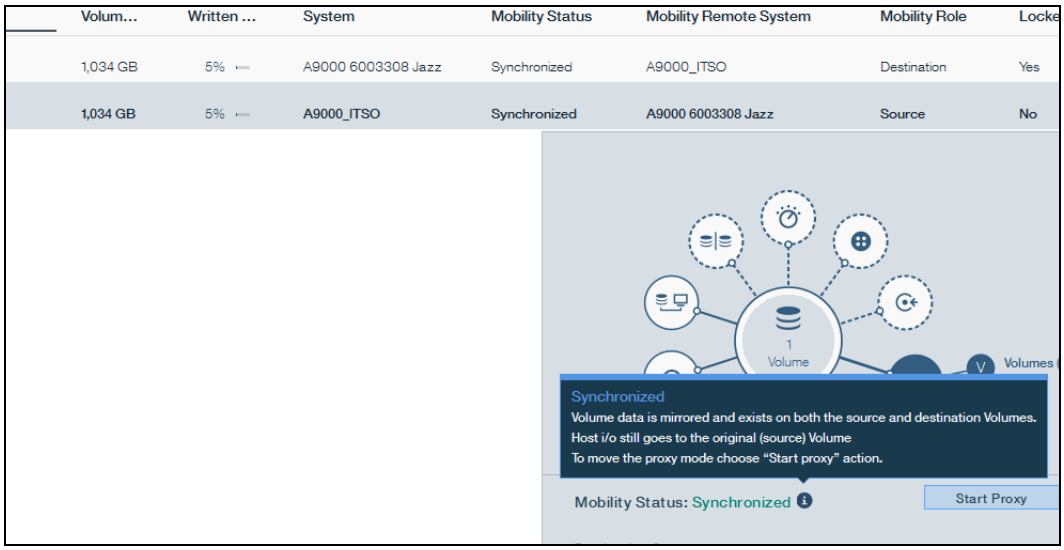

Figure 3-17 Proxy Ready and Synchronized at the source XIV Storage System

Now, the destination IBM FlashSystem A9000 contains a consistent copy of the source volume,
in addition to the necessary host and port definitions. The next task is to enter the actual migration stage.

Task 4: Proxying

Proxying means that the source IBM FlashSystem A9000 starts redirecting host I/O to the volume on the destination IBM FlashSystem A9000. To start the proxy mode, complete the following steps:

1. In the Hub view to the right of the Volumes tab, select **Start Proxy**, as shown in Figure 3-18.

Figure 3-18 Start IBM Hyper-Scale Mobility Proxy

2. As shown in Figure 3-19, a warning message is displayed, indicating that any source volume snapshots that exist are deleted if the proxy is started.

   Click **Apply** to proceed and initiate the Proxy phase.

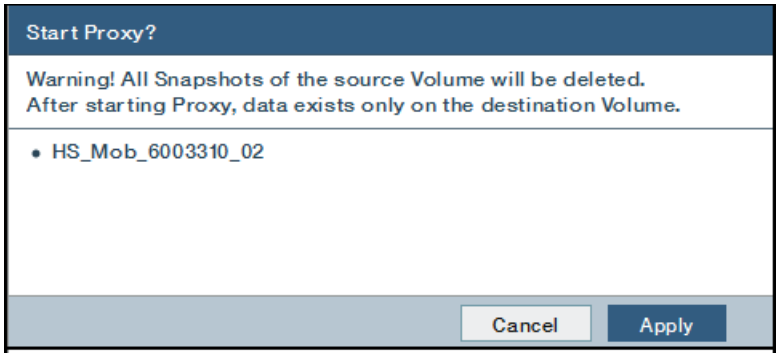

Figure 3-19 Snapshot deletion warning on proxy start

Important: Upon initiating the Proxy phase for a volume, it is no longer possible to stop the migration of this volume. Up to this point, the source volume still exists on the source IBM FlashSystem A9000, and any host I/O is still synchronized between both the source and destination volumes.

Clicking **Apply** in this window triggers the source IBM FlashSystem A9000 to proxy all subsequent host I/O to the volume on the destination IBM FlashSystem A9000 only, and the source volume becomes inconsistent and unavailable. The source volume is now just a logical placeholder for the original volume. Its size is 0 GB, and its hard (total usable) capacity is returned to the storage pool.

3. The Volumes tab now changes to show the Hyper-Scale Mobility relationship in the Proxy phase, as shown in Figure 3-20.

   Note that the source volume is disabled and labeled Proxy. Both volumes reflect the Mobility Status of Proxy Active.

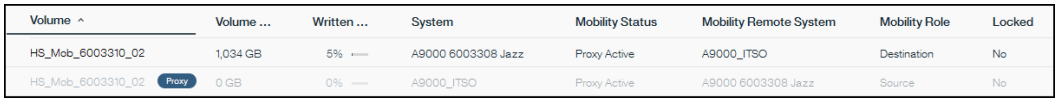

Figure 3-20 Proxy phase viewed from the Volumes Tab

4. Now that the Proxy phase is successfully initiated, the newly created destination volume is ready to map to the host on the destination IBM FlashSystem A9000. Select the destination volume in the Volumes tab and choose the Mapping spoke in the Hub display on the right, as shown in Figure 3-21.

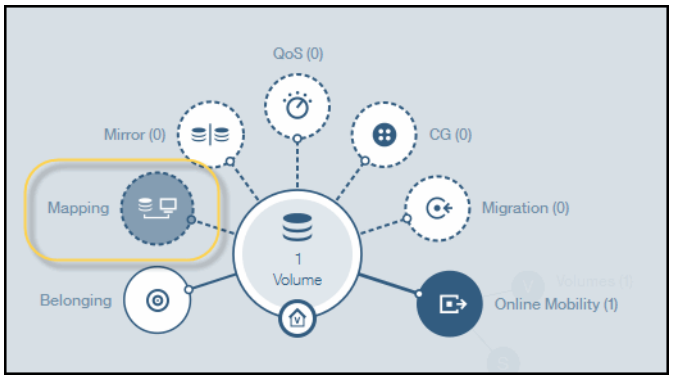

Figure 3-21 Selecting the Mapping spoke in the Hub display

5. Click the plus sign (**+**) below the Hub to add a new mapping, as shown in Figure 3-22.

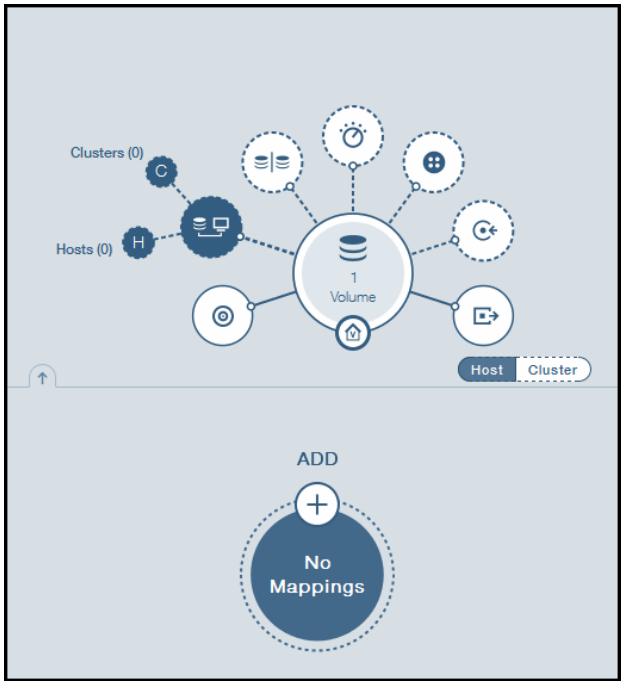

Figure 3-22 Add a new Mapping to the destination volume

6. In the window that opens, select the new host from the drop-down selection and the LUN ID to map it to. Click **Apply** to map the volume, as shown in Figure 3-23.

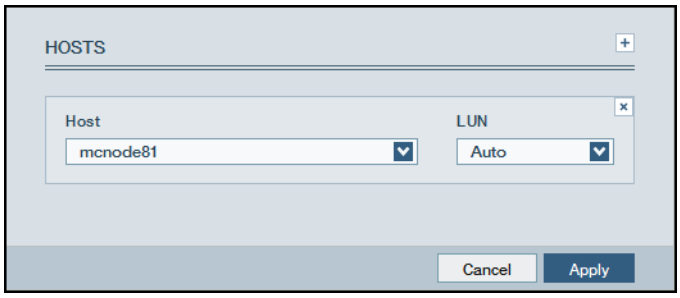

Figure 3-23 Select the host to map

7. After the destination volume is mapped to the host, a rescan of the devices from the host is needed to pick up the new paths to the volume on the destination system. Run the **xiv_fc_admin -R** command, and then run the **xiv_devlist** command from the host, as shown in Example 3-2.

   Note the 14 paths (six additional) for the destination volume, HS_Mob_6003310_02, in the bold row.

Example 3-2 Host rescan

```
[root@mcnode81 ~]# xiv_fc_admin -R
[root@mcnode81 ~]# xiv_devlist
XIV Devices
------------------------------------------------------------------------------
Device               Size (GB)  Paths  Vol Name          Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------------
/dev/mapper/mpathaqf  1033.6     8/8    HS_Mob_6003310_03  13445   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaql  1033.6     8/8    HS_Mob_6003310_01  13452   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaqo  1033.6     14/14  HS_Mob_6003310_02  13453   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxu  2067.3     6/6    HS_Mob_6003308_01  15857   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxv  2067.3     6/6    HS_Mob_6003308_03  15859   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathayb  2067.3     6/6    HS_Mob_6003308_02  15858   6003308  mcnode81
------------------------------------------------------------------------------
```

Note: The **xiv_fc_admin** and **xiv_devlist** commands are available on the host after installing the Host Attachment Kit (HAK).

Task 5: Cleanup

After the host has connectivity to the volume through the new paths to the destination IBM FlashSystem A9000, unmap the volume from the host on the source IBM FlashSystem A9000, and remove the paths to the source volume from the host.

To accomplish this task, complete the following steps:

1. Select the source volume in the Volumes tab of the Hyper-Scale Manager and select the mapping spoke in the Hub display, as shown in Figure 3-24.

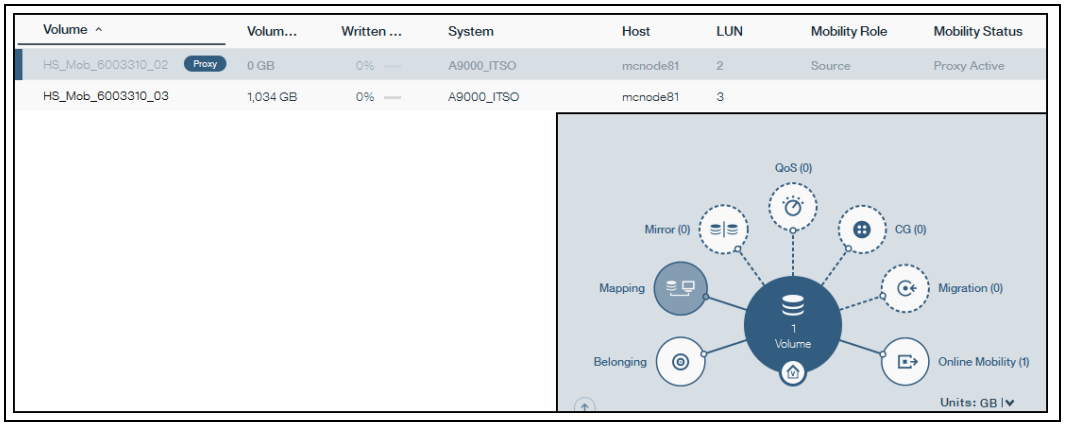

Figure 3-24 Mapping selection for the source volume

2. In the resulting Hub panel, click the **Actions** menu and select Unmap Host, as shown in Figure 3-25.

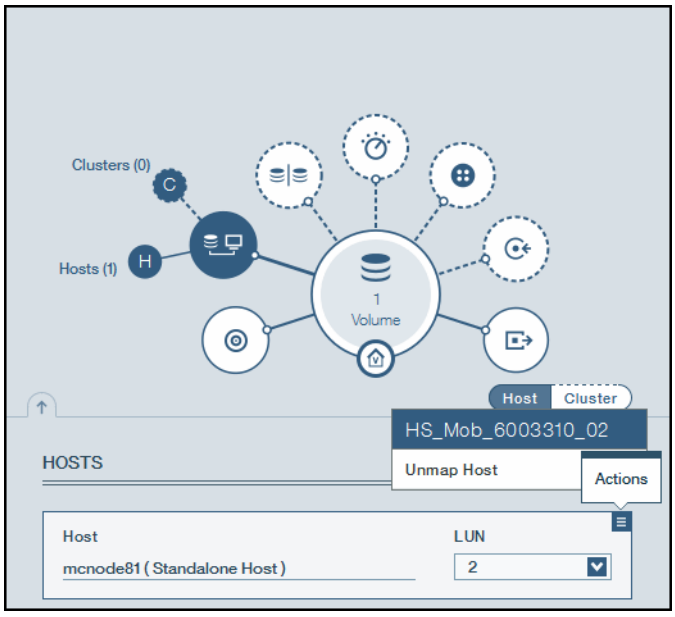

Figure 3-25 Unmap Host selection for the source volume

3. A confirmation window opens and prompts you to confirm that you want to unmap the selected volume, as shown in Figure 3-26. Click **Apply** to unmap the source volume from the host.

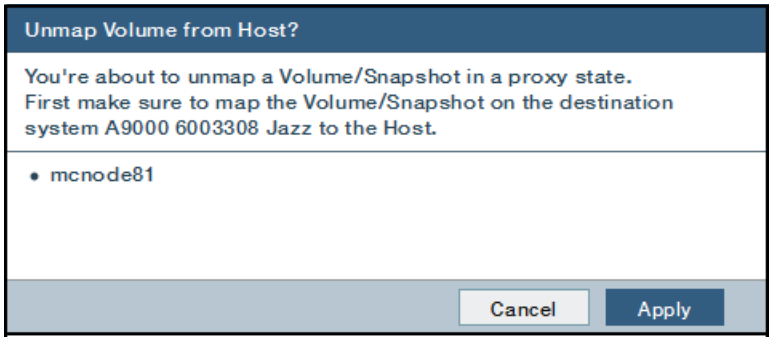

Figure 3-26 Confirm the unmapping of the source volume from the host

4. After the volume is unmapped from the host, only six enabled paths to the volume remain to the host. Run the **`xiv_fc_admin -R`** command, and then run the **`xiv_devlist`** command from the host to remove the unused paths, as shown in Example 3-3.

Example 3-3 Only six paths enabled

```
[root@mcnode81 ~]# xiv_fc_admin -R
[root@mcnode81 ~]# xiv_devlist
XIV Devices
------------------------------------------------------------------------------
Device               Size (GB)  Paths  Vol Name           Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------------
/dev/mapper/mpathaqf  1033.6     8/8    HS_Mob_6003310_03  13445   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaql  1033.6     8/8    HS_Mob_6003310_01  13452   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaqo  1033.6     6/14   HS_Mob_6003310_02  13453   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxu  2067.3     6/6    HS_Mob_6003308_01  15857   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxv  2067.3     6/6    HS_Mob_6003308_03  15859   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathayb  2067.3     6/6    HS_Mob_6003308_02  15858   6003308  mcnode81
------------------------------------------------------------------------------
```

Note the unmatching paths for the destination volume, HS_Mob_6003310_02, in the bold row. Also, note that the Vol ID is still 13453 and the XIV ID has changed to 6003308, which indicates that the host is connected *only* to the destination volume on the destination IBM FlashSystem A9000. The destination IBM FlashSystem A9000 keeps the same Vol ID, if possible.

1. Run the **xiv_fc_admin -R --clean** command, and then run the **xiv_devlist** command from the host to remove the failed paths from the RHEL host, as shown in Example 3-4.

Example 3-4 Path removal and rescan for RHEL

```
[root@mcnode81 ~]# xiv_fc_admin -R --clean
[root@mcnode81 ~]# xiv_devlist
XIV Devices
------------------------------------------------------------------------------------
Device                Size (GB)  Paths  Vol Name           Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------------------
/dev/mapper/mpathaqf  1033.6     8/8    HS_Mob_6003310_03  13445   6003310  mcnode81
------------------------------------------------------------------------------------
/dev/mapper/mpathaql  1033.6     8/8    HS_Mob_6003310_01  13452   6003310  mcnode81
------------------------------------------------------------------------------------
/dev/mapper/mpathaqo  1033.6     6/6    HS_Mob_6003310_02  13453   6003308  mcnode81
------------------------------------------------------------------------------------
/dev/mapper/mpathaxu  2067.3     6/6    HS_Mob_6003308_01  15857   6003308  mcnode81
------------------------------------------------------------------------------------
/dev/mapper/mpathaxv  2067.3     6/6    HS_Mob_6003308_03  15859   6003308  mcnode81
------------------------------------------------------------------------------------
/dev/mapper/mpathayb  2067.3     6/6    HS_Mob_6003308_02  15858   6003308  mcnode81
------------------------------------------------------------------------------------
```

2. To remove the failed paths in AIX®, use the sample script shown Example 3-5.

Example 3-5 Path removal in AIX

```
for disk in `lsdev -Cc disk | grep 2810 | awk '{ print $1 }'`
do
        for path in `lspath -l $disk -F "status connection" | grep Failed | awk '{ print
$2 }'`
        do
                echo $disk
                rmpath -l $disk -w $path -d
        done
done
```

Task 6: Post-cleanup

The previous example successfully used IBM Hyper-Scale Mobility via the Hyper-Scale Manager to move the volume HS_Mob_6003310_02 from IBM FlashSystem A9000 6003310 to IBM FlashSystem A9000 6003308 without incurring any downtime on the host. The final step is to end the proxy and delete the relationship. To accomplish this task, complete the following steps:

1. From the Volumes tab in the Hyper-Scale Manager, select the source volume and choose the Online Mobility spoke in the Hub view to the right, as shown in Figure 3-27.

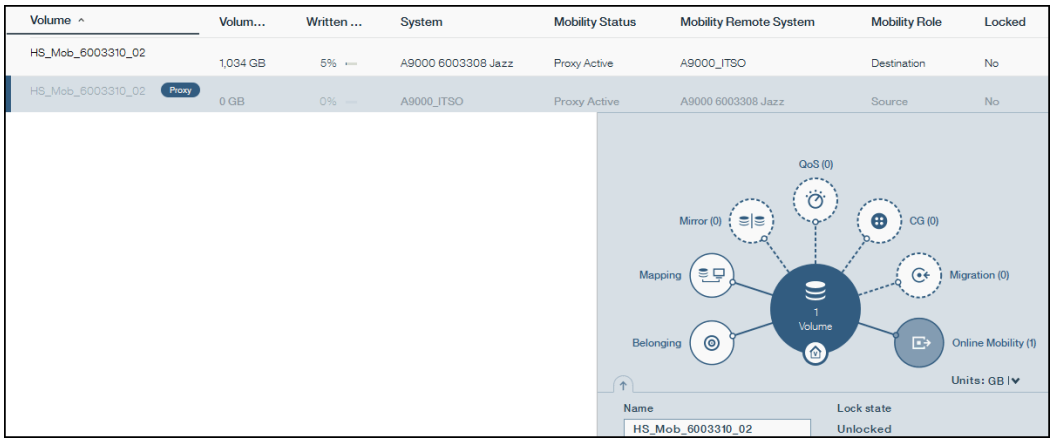

Figure 3-27 Online Mobility spoke for the source volume

2. Select End Proxy, as shown in Figure 3-28.

Figure 3-28 End the IBM Hyper-Scale Mobility proxy

3. A confirmation window opens, as shown in Figure 3-29, to ensure that you mapped the host to the destination volume and confirmed access to it. This was confirmed in the Cleanup stage.

   Click **Apply** to end the proxy.

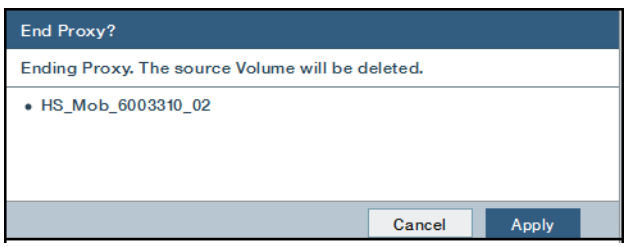

Figure 3-29 Confirm the ending of the IBM Hyper-Scale Mobility proxy

4. The proxy has been removed and the expected end state of the volumes on both the source and destination IBM FlashSystem A9000s are shown in Figure 3-30.

   On the source IBM FlashSystem A9000, volume HS_Mob_6003310_02 no longer exists; on the destination IBM FlashSystem A9000, there is a new volume named HS_Mob_6003310_02. This has the matching size and Written percentage (Used capacity) as the source volume before the Hyper-Scale Mobility process start.

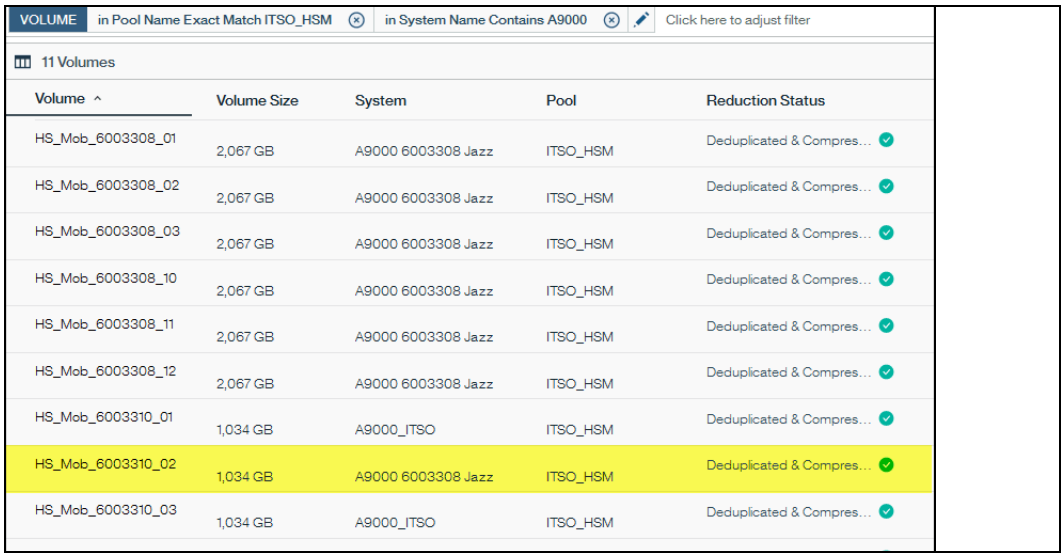

Figure 3-30 Volume list after clean-up

Confirm that the source volume is no longer on the source IBM FlashSystem A9000 and a new volume that matches the source now exists on the destination IBM FlashSystem A9000.

3.7.2 Using the XIV command-line interface with IBM Hyper-Scale Mobility

This section illustrates using the XCLI for IBM Hyper-Scale Mobility. The process consists of the same six tasks that are described and illustrated in 3.7.1, “Using the Hyper-Scale Manager to use IBM Hyper-Scale Mobility” on page 78.

Task 1: Planning and preparation

To perform an online migration of a IBM FlashSystem A9000 volume by using the XCLI, follow these steps:

1. Confirm that both the source and destination IBM FlashSystem A9000s are configured and accessible through the XCLI, and that there is Fibre Channel connectivity between these systems.

   Example 3-6 on page 94 shows the **target_list** and **`target_connectivity_list`** commands run on the source IBM FlashSystem A9000 to confirm the connectivity. The output that is shown in Example 3-6 on page 94 shows that the destination IBM FlashSystem A9000 is connected through only a single active FC connection.

Example 3-6 The target_list and target_connectivity_list commands before migration

```
#Source:
A9000_ITSO>>target_list
Name                 SCSI Type   Connected   Max Initialization Rate   Max Resync Rate   Max Syncjob Rate
A9000 6003308 Jazz   FC          yes         100                       300               300
A9000_ITSO>>target_connectivity_list
Target Name          Remote Port        FC Port         IP Interface   Active   Up
A9000 6003308 Jazz   500173800CEC0123   1:FC_Port:3:4                  yes      yes

#Destination:
A9000 6003308 Jazz>>target_connectivity_list
Target Name          Remote Port        FC Port         IP Interface   Active   Up
A9000_ITSO           500173800CEE0123   1:FC_Port:3:4                  yes      yes
```

2. In addition, confirm that any host that has mappings to the source volume on the source IBM FlashSystem A9000 has the appropriate access to the destination IBM FlashSystem A9000.

   Consider the following items:

   – For a migration of an iSCSI LUN, ensure that the host can access the destination system over the IP network.

   – For a migration of an FC LUN, ensure that the host is correctly zoned to the destination system.

3. Next, validate the host multipath connectivity to the source volume by running the **xiv_devlist** command of the HAK, as shown in Example 3-7.

   Example 3-7 The xiv_devlist command before migration

```
[root@mcnode81 ~]# xiv_devlist
XIV Devices
------------------------------------------------------------------------------
Device                Size (GB)  Paths  Vol Name           Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------------
/dev/mapper/mpathaqf  1033.6     8/8    HS_Mob_6003310_03  13445   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaql  1033.6     8/8    HS_Mob_6003310_01  13452   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaqo  1033.6     6/6    HS_Mob_6003310_02  13453   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxu  2067.3     6/6    HS_Mob_6003308_01  15857   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxv  2067.3     6/6    HS_Mob_6003308_03  15859   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathayb  2067.3     6/6    HS_Mob_6003308_02  15858   6003308  mcnode81
------------------------------------------------------------------------------
```

   Note the Vol ID of 13452 and XIV ID of 6003310, which indicates that the volume is presented to the host from the source IBM FlashSystem A9000 6003310.

   Example 3-8 on page 95 shows the output of the **vol_list** command for the source and destination IBM FlashSystem A9000s before the start of the IBM Hyper-Scale Mobility process.

Example 3-8 The vol_list command on source and destination IBM FlashSystem A9000s before migration

```
#Source:
A9000_ITSO>>vol_list pool=ITSO_HSM
Name                Size (GB)   Master Name   Consistency Group   Pool       Creator   Written (GB)
HS_Mob_6003310_03   1033                                          ITSO_HSM   admin     0
HS_Mob_6003310_11   1033                                          ITSO_HSM   admin     0
HS_Mob_6003310_12   1033                                          ITSO_HSM   admin     0
HS_Mob_6003310_01   1033                                          ITSO_HSM   admin     30

#Destination:
A9000 6003308 Jazz>>vol_list pool=ITSO_HSM
Name                Size (GB)   Master Name   Consistency Group   Pool       Creator   Written (GB)
HS_Mob_6003308_01   2067                                          ITSO_HSM   admin     53
HS_Mob_6003308_02   2067                                          ITSO_HSM   admin     0
HS_Mob_6003308_03   2067                                          ITSO_HSM   admin     0
HS_Mob_6003308_10   2067                                          ITSO_HSM   admin     3
HS_Mob_6003308_11   2067                                          ITSO_HSM   admin     0
HS_Mob_6003308_12   2067                                          ITSO_HSM   admin     0
HS_Mob_6003310_02   1033                                          ITSO_HSM             53
```

Task 2: Setup

To perform the Setup task, perform the following steps:

1. To create the IBM Hyper-Scale Mobility relationship, run the **olvm_create** command, as shown in Example 3-9. The required parameters are **vol=**, **remote_pool=**, and **target=**.

Attention: Any existing snapshot for the IBM Hyper-Scale Mobility source volume is *deleted without explicit warning* from the XCLI when the **olvm_create** command is executed.

Example 3-9 The olvm_create command

```
#Source:
A9000_ITSO>>olvm_create vol=HS_Mob_6003310_01 remote_pool=ITSO_HSM target="A9000 6003308 Jazz"
Command executed successfully.
A9000_ITSO>>olvm_list
Volume name         Role     Remote System        Active   Phase       State          Link Up
HS_Mob_6003310_01   source   A9000 6003308 Jazz   no       migration   Initializing   yes

#Destination:
A9000 6003308 Jazz>>olvm_list
Volume name         Role          Remote System   Active   Phase       State          Link Up
HS_Mob_6003310_01   destination   A9000_ITSO      no       migration   Initializing   yes
```

2. As seen in Example 3-9, running the **olvm_create** command accomplishes the following tasks:
   - Creates a volume on the destination XIV Storage System, with same name as the source volume, in the specified **remote_pool**
   - Creates a Hyper-Scale Mobility relationship between the source IBM FlashSystem A9000 volume and the destination IBM FlashSystem A9000 volume

   The **olvm_create** command only creates the relationship. The IBM Hyper-Scale Mobility relationship is not activated in this step, and the destination volume is empty.

Task 3: Migration

To complete the migration task, perform the following steps:

1. After the IBM Hyper-Scale Mobility relationship is created, activate it by running the **olvm_activate** command, as shown in Example 3-10.

Example 3-10 The olvm_activate command

```
#Source:
A9000_ITSO>>olvm_activate vol=HS_Mob_6003310_01
Command executed successfully.
A9000_ITSO>>olvm_list
Volume name         Role     Remote System        Active   Phase       State          Link Up
HS_Mob_6003310_01   source   A9000 6003308 Jazz   yes      migration   Initializing   yes

#After initialization is complete:
A9000_ITSO>>olvm_list
Volume name         Role     Remote System        Active   Phase   State          Link Up
HS_Mob_6003310_01   source   A9000 6003308 Jazz   yes      ready   Synchronized   yes

#Destination:
A9000 6003308 Jazz>>olvm_list
Volume name         Role          Remote System   Active   Phase       State          Link Up
HS_Mob_6003310_01   destination   A9000_ITSO      yes      migration   Initializing   yes

#After initialization is complete:
A9000 6003308 Jazz>>olvm_list
Volume name         Role          Remote System   Active   Phase   State        Link Up
HS_Mob_6003310_01   destination   A9000_ITSO      yes      ready   Consistent   yes
```

2. The **olvm_activate** command initializes the synchronization of the source and destination volumes. Note that the Hyper-Scale Mobility relationship becomes active (`Active=yes`). The synchronization is queued because only one IBM Hyper-Scale Mobility process is synchronized at a time.

   The time that is required to complete the synchronization depends on numerous factors, including how many migrations are active, and the amount of data that needs to be synchronized per volume. After the synchronization has completed, the Phase changes to ready on both the source and destination volumes. Additionally, the State changes to Synchronized on the source volume as the State changes to Consistent on the destination volume.

Task 4: Proxying

Proxying means that the source IBM FlashSystem A9000 starts redirecting host I/O to the volume on the destination IBM FlashSystem A9000. To start the proxy mode, complete the following steps:

1. During the migration, host and host-port definitions can be created by running the commands shown in Example 3-11. Note that these definitions can also be done before, and independently of, the migration process. Example 3-12 shows the resulting host definition on the destination IBM FlashSystem A9000.

Example 3-11 The host_define and host_add_port commands

```
#Destination:
A9000 6003308 Jazz>>host_define host=mcnode81
Command executed successfully.
A9000 6003308 Jazz>>host_add_port host=mcnode81 fcaddress=10000090FA731D36
Command executed successfully.
A9000 6003308 Jazz>>host_add_port host=mcnode81 fcaddress=10000090FA731D37
Command executed successfully.
```

Example 3-12 The host_list command on the destination IBM FlashSystem A9000

```
A9000 6003308 Jazz>>host_list host=mcnode81
Name       Type      FC Ports                          iSCSI Ports   User Group   Cluster   Performance Class
mcnode81   default   10000090FA731D36,10000090FA731D37                                      none
```

2. Run the XCLI **olvm_proxy** command to start the Proxy phase, as shown in Example 3-13.

Important: Upon initiating the Proxy phase for a volume, it is no longer possible to stop the IBM Hyper-Scale Mobility for this volume. Up to this point, the source volume still exists on the source IBM FlashSystem A9000, and any host I/O is still synchronized between both source and destination volumes.

Answering y to confirm the command execution shown in Example 3-13 triggers the source IBM FlashSystem A9000 to proxy all subsequent host I/O to the volume on the destination IBM FlashSystem A9000 only, and the source volume becomes inconsistent and unavailable. The source volume is now just a logical placeholder for the original volume. Its size is 0 GB, and its hard (total usable) capacity is returned to the storage pool.

Example 3-13 The olvm_proxy command

```
#Source:
A9000_ITSO>>olvm_proxy vol=HS_Mob_6003310_01

Warning:   Are you sure you want to move the volume HS_Mob_6003310_01 to a
Proxy state? Source volume and all volume snapshots will be deleted. y/n: y
Command executed successfully.
A9000_ITSO>>olvm_list
Volume name         Role     Remote System        Active   Phase   State   Link Up
HS_Mob_6003310_01   source   A9000 6003308 Jazz   yes      proxy   Proxy   yes

#Destination:
A9000 6003308 Jazz>>olvm_list
Volume name         Role          Remote System   Active   Phase   State     Link Up
HS_Mob_6003310_01   destination   A9000_ITSO      yes      proxy   Proxied   yes
```

3. To verify the start of the Proxy phase execute the **olvm_list** command, as shown in Example 3-13 on page 97. The State of the source volume changes to Proxy, and the State of the destination volume changes to Proxied. All host I/O to the source volume is redirected (proxied) to the destination volume. In Example 3-14, the output of the **vol_list** command on each IBM FlashSystem A9000 now shows that the volume named HS_Mob_6003310 is no longer showing on the source IBM FlashSystem A9000 but has been moved to the destination IBM FlashSystem A9000.

Example 3-14 The vol_list command after the Hyper-Scale Process is in State Proxy/Proxied

```
#Source:
A9000_ITSO>>vol_list pool=ITSO_HSM
Name               Size (GB)   Master Name   Consistency Group   Pool       Creator   Written (GB)
HS_Mob_6003310_03  1033                                          ITSO_HSM   admin     0
HS_Mob_6003310_11  1033                                          ITSO_HSM   admin     0
HS_Mob_6003310_12  1033                                          ITSO_HSM   admin     0

#Destination:
A9000 6003308 Jazz>>vol_list pool=ITSO_HSM
Name               Size (GB)   Master Name   Consistency Group   Pool       Creator   Written (GB)
HS_Mob_6003308_01  2067                                          ITSO_HSM   admin     53
HS_Mob_6003308_02  2067                                          ITSO_HSM   admin     0
HS_Mob_6003308_03  2067                                          ITSO_HSM   admin     0
HS_Mob_6003308_10  2067                                          ITSO_HSM   admin     3
HS_Mob_6003308_11  2067                                          ITSO_HSM   admin     0
HS_Mob_6003308_12  2067                                          ITSO_HSM   admin     0
HS_Mob_6003310_02  1033                                          ITSO_HSM             53
HS_Mob_6003310_01  1033                                          ITSO_HSM             30
```

4. After the Proxy phase is successfully initiated, the destination volume can be mapped to the host by running the **map_vol** command, as shown in Example 3-15.

Example 3-15 The map_vol command

```
#Destination:

A9000 6003308 Jazz>>map_vol vol=HS_Mob_6003310_01 host=mcnode81 lun=3
Command executed successfully.
```

5. After mapping, rescan the devices from the host to pick up the new paths to the volume on the destination IBM FlashSystem A9000. Run the **xiv_fc_admin -R** command, and then run the **xiv_devlist** command from the host, as shown in Example 3-16.

Example 3-16 The xiv_fc_admin -R and xiv_devlist commands

```
[root@mcnode81 ~]# xiv_fc_admin -R
[root@mcnode81 ~]# xiv_devlist
XIV Devices
------------------------------------------------------------------------------
Device               Size (GB)  Paths  Vol Name           Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------------
/dev/mapper/mpathaqf  1033.6     8/8    HS_Mob_6003310_03  13445   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaql  1033.6     14/14  HS_Mob_6003310_01  13452   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaqo  1033.6     6/6    HS_Mob_6003310_02  13453   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxu  2067.3     6/6    HS_Mob_6003308_01  15857   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxv  2067.3     6/6    HS_Mob_6003308_03  15859   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathayb  2067.3     6/6    HS_Mob_6003308_02  15858   6003308  mcnode81
------------------------------------------------------------------------------
```

Note the 14 paths (six additional) for the destination volume, HS_Mob_6003310_01, in the bold row.

Task 5: Cleanup

To complete the Cleanup task, perform the following steps:

1. After validating that the host has connectivity to the destination volume through the new paths to the destination IBM FlashSystem A9000, unmap the source volume on the source IBM FlashSystem A9000 from the host, as shown in Example 3-17.

Example 3-17 The unmap_vol command

```
#Source:
A9000_ITSO>>unmap_vol vol=HS_Mob_6003310_01 host=mcnode81
Command executed successfully.
A9000_ITSO>>olvm_list
Volume name         Role     Remote System        Active   Phase   State   Link Up
HS_Mob_6003310_01   source   A9000 6003308 Jazz   yes      proxy   Proxy   yes

#Destination:
A9000 6003308 Jazz>>olvm_list
Volume name         Role          Remote System   Active   Phase   State     Link Up
HS_Mob_6003310_01   destination   A9000_ITSO      yes      proxy   Proxied   yes
```

2. After unmapping the volume from the host, only six enabled paths to the volume remain to the host. From the host run the **xiv_fc_admin -R** command followed by the **xiv_devlist** command, as shown in Example 3-18.

Example 3-18 6 out of the 14 paths enabled

```
[root@mcnode81 ~]# xiv_fc_admin -R
[root@mcnode81 ~]# xiv_devlist
XIV Devices
------------------------------------------------------------------------------
Device               Size (GB)  Paths  Vol Name           Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------------
/dev/mapper/mpathaqf  1033.6     8/8    HS_Mob_6003310_03  13445   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaql  1033.6     6/14   HS_Mob_6003310_01  13452   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaqo  1033.6     6/6    HS_Mob_6003310_02  13453   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxu  2067.3     6/6    HS_Mob_6003308_01  15857   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxv  2067.3     6/6    HS_Mob_6003308_03  15859   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathayb  2067.3     6/6    HS_Mob_6003308_02  15858   6003308  mcnode81
------------------------------------------------------------------------------
```

Note the unmatching paths for the destination volume, HS_Mob_6003310_01, in the bold row. Also, note the Vol ID is still 13452 and the XIV ID has changed to 6003308, which indicates that the host is connected *only* to the destination volume on the destination IBM FlashSystem A9000. The destination IBM FlashSystem A9000 keeps the same Vol ID, if possible.

3. Run the **xiv_fc_admin -R --clean** command, and then run the **xiv_devlist** command from the host to remove the failed paths from the RHEL host, as shown in Example 3-19. To remove the failed paths from an AIX host, a sample script is shown in Example 3-5 on page 91.

Example 3-19 Path removal and rescan for RHEL

```
[root@mcnode81 ~]# xiv_fc_admin -R --clean
[root@mcnode81 ~]# xiv_devlist
XIV Devices
------------------------------------------------------------------------------
Device               Size (GB)  Paths  Vol Name           Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------------
/dev/mapper/mpathaqf  1033.6     8/8    HS_Mob_6003310_03  13445   6003310  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaql  1033.6     6/6    HS_Mob_6003310_01  13452   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaqo  1033.6     6/6    HS_Mob_6003310_02  13453   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxu  2067.3     6/6    HS_Mob_6003308_01  15857   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathaxv  2067.3     6/6    HS_Mob_6003308_03  15859   6003308  mcnode81
------------------------------------------------------------------------------
/dev/mapper/mpathayb  2067.3     6/6    HS_Mob_6003308_02  15858   6003308  mcnode81
------------------------------------------------------------------------------
```

Task 6: Post-cleanup

The previous example successfully used IBM Hyper-Scale Mobility via the XCLI to move the volume HS_Mob_6003310_01 from IBM FlashSystem A9000 6003310 to IBM FlashSystem A9000 6003308 without incurring any downtime on the host. The final step is to end the proxy and delete the relationship: To accomplish this task, complete the following steps:

1. Run the **olvm_delete** command, as shown in Example 3-20.

Example 3-20 The olvm_delete command

```
#Source:
A9000_ITSO>>olvm_delete vol=HS_Mob_6003310_01

Warning:   Are you sure you want to delete IBM Hyper-Scale Mobility relationship? y/n: y
Command executed successfully.
A9000_ITSO>>olvm_list
No olvms match the given criteria

#Destination:
A9000 6003308 Jazz>>olvm_list
No olvms match the given criteria
```

2. The proxy has been removed and the expected end state of the volumes on both the source and destination IBM FlashSystem A9000s are shown in Example 3-21. On the source IBM FlashSystem A9000, volume HS_Mob_6003310_02 no longer exists; on the destination IBM FlashSystem A9000, there is a new volume named HS_Mob_6003310_02. This has the matching size and Written percentage (Used capacity) as the source volume before the Hyper-Scale Mobility process start.

Example 3-21 The vol_list by pool command

```
#Source:
A9000_ITSO>>vol_list pool=ITSO_HSM
Name                Size (GB)   Master Name   Consistency Group   Pool       Creator   Written (GB)
HS_Mob_6003310_03   1033                                          ITSO_HSM   admin     0
HS_Mob_6003310_11   1033                                          ITSO_HSM   admin     0
HS_Mob_6003310_12   1033                                          ITSO_HSM   admin     0

#Destination:
A9000 6003308 Jazz>>vol_list pool=ITSO_HSM
Name                Size (GB)   Master Name   Consistency Group   Pool       Creator   Written (GB)
HS_Mob_6003308_01   2067                                          ITSO_HSM   admin     53
HS_Mob_6003308_02   2067                                          ITSO_HSM   admin     0
HS_Mob_6003308_03   2067                                          ITSO_HSM   admin     0
HS_Mob_6003308_10   2067                                          ITSO_HSM   admin     3
HS_Mob_6003308_11   2067                                          ITSO_HSM   admin     0
HS_Mob_6003308_12   2067                                          ITSO_HSM   admin     0
HS_Mob_6003310_02   1033                                          ITSO_HSM             53
HS_Mob_6003310_01   1033                                          ITSO_HSM             30
```

4

Data migration with IBM FlashSystem A9000 and IBM FlashSystem A9000R

This chapter provides a description of the migration process, which is a built-in feature of IBM FlashSystem A9000 and IBM FlashSystem A9000R. The Data Migration Utility can migrate data from almost any storage system to IBM FlashSystem A9000 and A9000R. The chapter also includes the pre-migration and post-migration steps.

This chapter includes the following sections:

- Migration overview
- Detailed migration steps for an iSCSI connection
- Detailed migration steps for a Fibre Channel connection

4.1 Migration overview

IBM FlashSystem A9000 and A9000R includes, at no extra charge, a powerful data migration capability. The Data Migration Utility can migrate data from almost any storage system to IBM FlashSystem A9000 and A9000R.

At the start of the migration, application hosts are offline for only a short time as they are connected to IBM FlashSystem A9000 or A9000R. The logical volumes, also known as logical unit numbers (LUNs), that are migrated are reallocated from the former storage server to IBM FlashSystem A9000 or A9000R and are then natively presented again to the host. Meanwhile, the data is transparently migrated in the background, in a controlled fashion.

The migration overview, which is depicted in Figure 4-1, assumes a migration from another storage system to IBM FlashSystem A9000 or A9000R. The storage system can be an IBM storage system or another storage system, as long as that storage system supports Linux hosts. The migration process is the same for all storage systems.

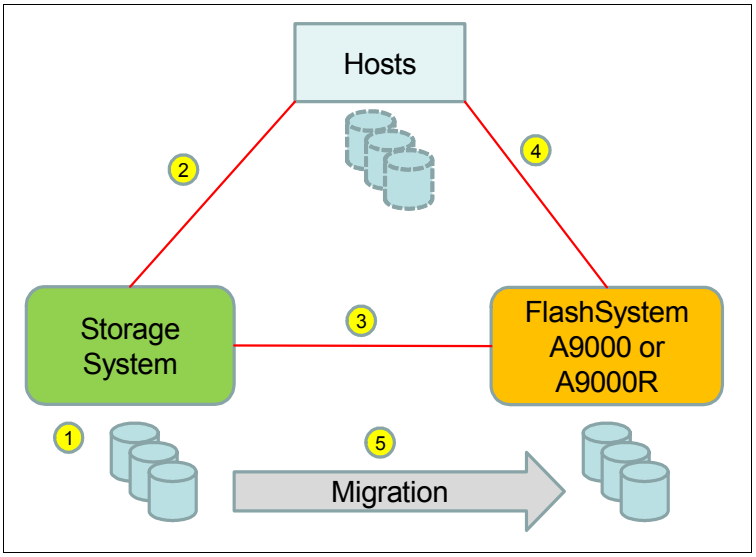

Figure 4-1 Migration overview

Figure 4-1 shows the following series of processes (the following numbers correspond to the numbers in the figure):

1. Volumes and pools that are defined in the Storage System are initially used by host systems.
2. In the illustration, hosts are connected over Fibre Channel (FC) or internet Small Computer System Interface (iSCSI) targets and the LUNs that are defined in the Storage System are mapped to those hosts.
3. An FC or iSCSI target is defined in IBM FlashSystem A9000 or A9000R for the Storage System, as though it were a host. Target migration volumes in IBM FlashSystem A9000 or A9000R are defined and mapped to the remote LUNs in the Storage System.
4. A new FC or iSCSI host object for the host has to be defined in IBM FlashSystem A9000 or A9000R. At this stage the storage system has two hosts that are defined in the system (the actual host and IBM FlashSystem A9000 or A9000R).
5. The migration feature is now enabled in IBM FlashSystem A9000 or A9000R. Before this enablement starts, the host connectivity to the Storage System has to be terminated

(apparent downtime for the host) and the volumes being migrated are manually presented by IBM FlashSystem A9000 or A9000R to the host.

4.2 Detailed migration steps for an iSCSI connection

Note: We chose to migrate from IBM FlashSystem A9000 to IBM FlashSystem A9000R. These systems were used for illustration because they were available for our experimentation. Migration can be performed from any storage system that can attach to a Linux host. The setup details differ based on the storage system.

At a high level, we used the following steps to migrate volumes from an IBM FlashSystem A9000 (or any storage system) to IBM FlashSystem A9000R:

1. Initial connection and preparation:
   - Establish connection between IBM FlashSystem A9000 and IBM FlashSystem A9000R.
   - Prepare an iSCSI qualified name (IQN) table for all of the systems that are involved.
   - Assume that the volumes and the host are already defined in the IBM FlashSystem A9000.
   - Define IBM FlashSystem A9000R to the IBM FlashSystem A9000 with type **Default**.
   - Define IBM FlashSystem A9000 in IBM FlashSystem A9000R (as a migration target).
2. Perform pre-migration tasks for each host that is attached to the IBM FlashSystem A9000 that will be migrated:
   - Back up host data.
   - Stop all I/O from the host to the LUNs on IBM FlashSystem A9000.
   - Update the Host Attachment Kit (HAK) to latest version.

Important: If your migration source is not an IBM XIV, IBM FlashSystem A9000, or IBM FlashSystem A9000R, it is important that you uninstall any specific multipath software (for example IBM SDD) and install the IBM Host Attachment Kit (HAK).

   - A reboot of the host might be needed depending on the multipath software changes requirements.
3. Define and test the data migration volume:
   - On the IBM FlashSystem A9000, remap volumes from the old host to IBM FlashSystem A9000R.
   - On IBM FlashSystem A9000R, create the data migration tasks and test them.
4. Activate the data migration tasks on IBM FlashSystem A9000R.
5. Define the host in IBM FlashSystem A9000R and bring the host and applications online:
   - Define the host and iSCSI name on IBM FlashSystem A9000R.
   - Map volumes to the host on IBM FlashSystem A9000R.
   - Verify that IBM FlashSystem A9000R storage is visible.
   - Depending on the host a rescan on the host for the volumes on IBM FlashSystem A9000R might be necessary. Enable and start the host applications.
6. Complete the data migration on IBM FlashSystem A9000R:
   - Monitor IBM FlashSystem A9000R migration tasks.
   - On completion, delete the migration tasks.

Tip: Print these overview steps and refer to them as you perform a migration.

The following steps are specific to the example that is mentioned here and focus only on the iSCSI-attached volumes in an IBM FlashSystem A9000 with a Windows host.

4.2.1 Initial connection and preparation

For the initial connection setup, start by connecting IBM FlashSystem A9000R to IBM FlashSystem A9000 that you are migrating from.

Because the migration and the host attachment are through iSCSI, you must ensure that appropriate firewall ports are opened to allow iSCSI communications.

Important: If the Internet Protocol (IP) network includes firewalls, TCP port 3260 must be open for iSCSI host attachment and migration to work.

It is also a good opportunity for you to update the host operating system (OS) patches and drivers to the latest supported levels for the storage system that you are migrating (IBM FlashSystem A9000 in our scenario).

Prepare an iSCSI qualified name (IQN) table

A preferred practice is to document all of the IQNs of all of the systems that are involved, as shown in Table 4-1.

Table 4-1 Example of an iSCSI qualified name (IQN) table

Machine type	iSCSI name
Windows host	iqn.1991-05.com.microsoft:windows-ilplid1
IBM FlashSystem A9000	iqn.2005-10.com.xivstorage:01322131
IBM FlashSystem A9000R	iqn.2005-10.com.xivstorage:01320902

On a Windows host, click **Control Panel** → **iSCSI initiator**. You can determine the IQNs from the iSCSI initiator System Properties display.

Using the Hyper-Scale Manager you can see the iSCSI port address in the Host Properties window, as shown in Figure 4-2 on page 107.

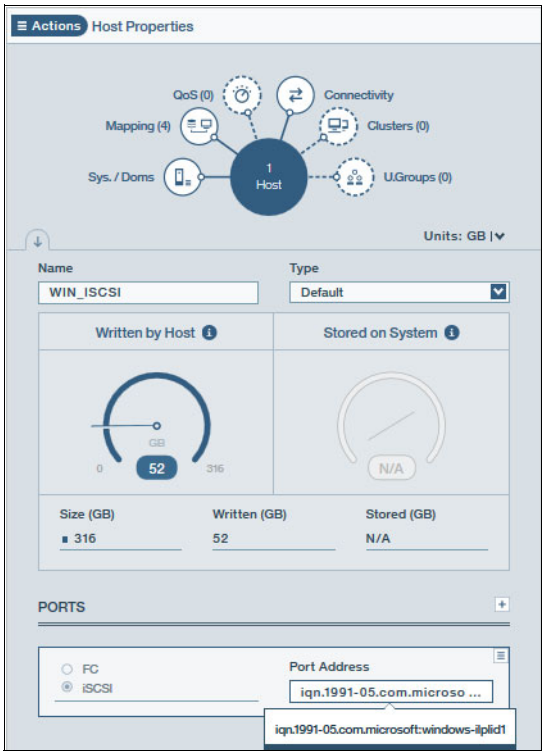

Figure 4-2 iSCSI qualified name in the Host Properties window

Volumes and host that are defined in IBM FlashSystem A9000

The example assumes that the pools and the volumes that you want to migrate existed in IBM FlashSystem A9000. This example also assumes that the host that uses the volumes existed in IBM FlashSystem A9000.

Define IBM FlashSystem A9000R on IBM FlashSystem A9000 as host

IBM FlashSystem A9000R must be defined as a **Default** host on IBM FlashSystem A9000. Use the Hyper-Scale Manger **Create New** → **Host** to define IBM FlashSystem A9000R as a host. In this example it is defined as iSCSI host, as shown in Figure 4-3.

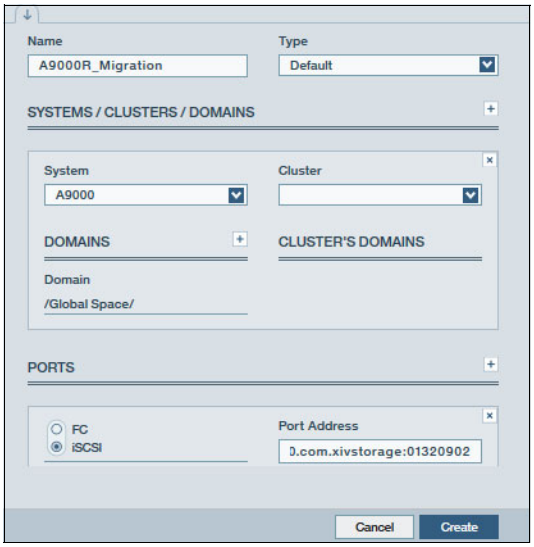

Figure 4-3 Hosts in the Hyper-Scale Manager

Select **iSCSI** and enter the iSCSI qualified name, as shown in Figure 4-3 on page 107.

Define IBM FlashSystem A9000 on IBM FlashSystem A9000R as a target

After the physical connectivity is established and IBM FlashSystem A9000R is defined to IBM FlashSystem A9000 (or any iSCSI-attached storage), you must also define the IBM FlashSystem A9000 on IBM FlashSystem A9000R.

You must define the storage system target. Also, you must define the connectivity between IBM FlashSystem A9000R and IBM FlashSystem A9000 through the iSCSI target initiator. Complete these steps:

1. By using the Hyper-Scale Manager GUI, log in to IBM FlashSystem A9000R system with the storage administrator credentials.
2. Click **Actions** → **Targets** → **View/Modify Targets** and click the icon to add a target, as shown in Figure 4-4.

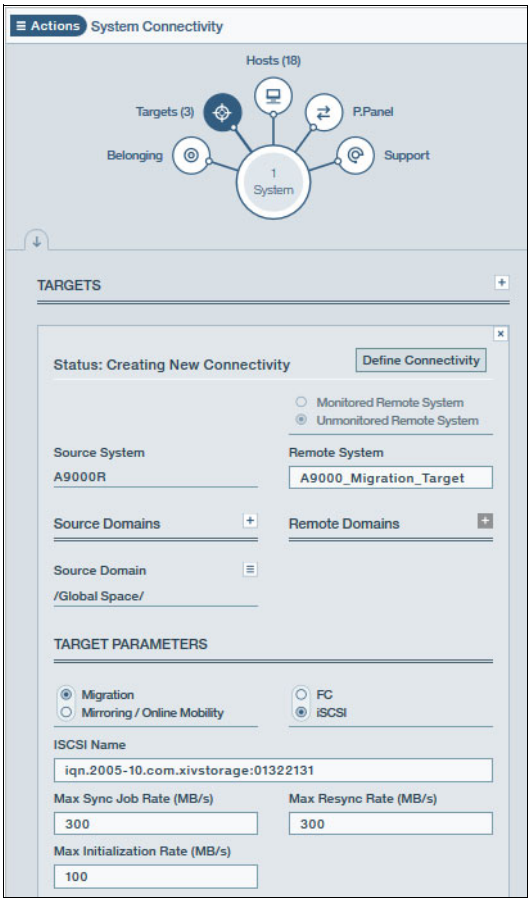

Figure 4-4 Defining a target for migration in IBM FlashSystem A9000R

3. Enter the following information and select the following options, as shown in Figure 4-5:

   Remote System This field is for the current system, for example, `A9000_Migration_Target`. You have to select **Unmonitored Remote System** to be able to enter the name.

   Source Domains The default is **/Global Space/**, which means *no-domain* unless specified

   Remote Domain Depending on the previous selection of domain, other options might be available.

   Target Type Two options exist: Migration and Mirroring. Select **Migration**.

   Target Protocol Fibre Channel or iSCSI. For this example, **iSCSI** is selected.

   iSCSI (Initiator) Name The iSCSI target that IBM FlashSystem A9000R attaches its volumes from. In this example, IBM FlashSystem A9000 is the iSCSI target, so select the corresponding iSCSI name from Table 4-1 on page 106.

   Max. Sync Job Rate Depending on the network bandwidth, the default value is set to 300 MBps. --- Not applicable for migration

   Max. Resync Rate Depending on the network bandwidth, the default value is set to 300 MBps. --- Not applicable for migration

   Max. Init Rate Depending on the network bandwidth, the default value is set to 100 MBps.

4. After the target is defined, the target appears in the System Connectivity display as shown in Figure 4-5.

   Select **Define Connectivity** to define the iSCSI connections.

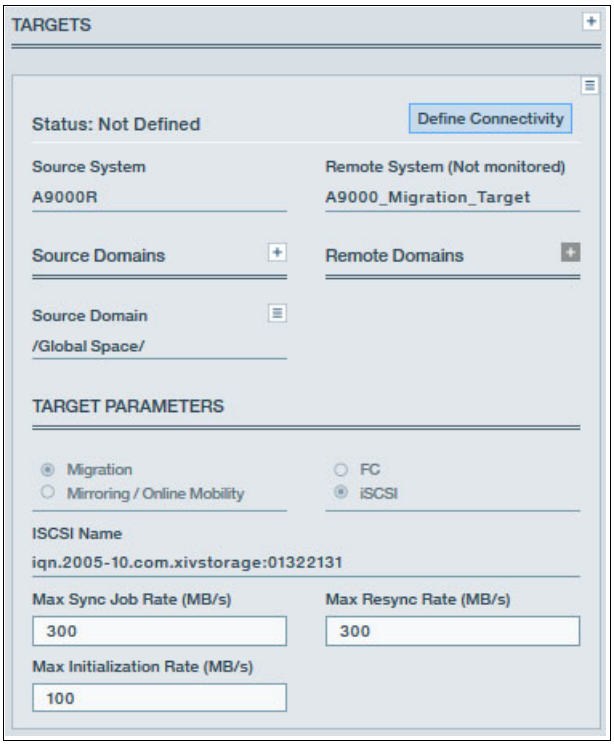

Figure 4-5 Defining target ports in IBM FlashSystem A9000R

5. In the window that is displayed, complete the following steps:

   a. Enter the IP addresses for the iSCSI connectivity of IBM FlashSystem A9000 to define the connection between the two systems, as shown in Figure 4-6.

Figure 4-6 Define the IP ports for the target system

   b. Select **Auto Connect** to establish the connections between IBM FlashSystem A9000R and the target system.

   c. If the connection between IBM FlashSystem A9000R and the target system (IBM FlashSystem A9000) is successful, the connection is displayed as connected (with a check mark) and you can select **Connectivity Details** to view the connectivity, as shown in Figure 4-7.

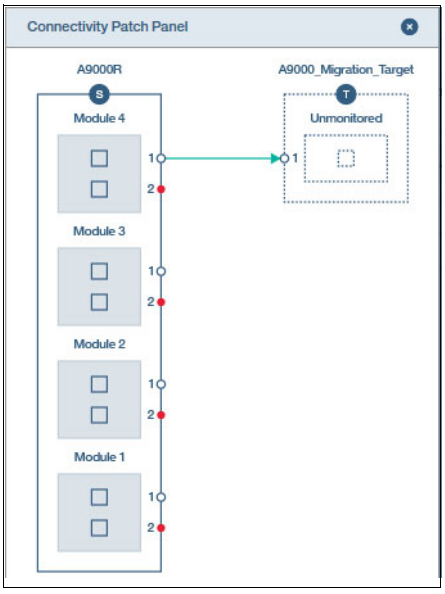

Figure 4-7 Connectivity details

4.2.2 Perform pre-migration tasks for each host that is migrated

Perform the following pre-migration tasks for each host that you are migrating:

- Back up the volumes that are migrated.
- Stop all I/O from the host to the LUNs on the storage that is migrated.

Back up the volumes that are migrated

It is a good practice to take a full restorable backup before any data migration activity. Verify that all of the data is restorable from the backup and that no backup media errors exist.

In addition to a regular backup, a point-in-time copy of the LUNs that are being migrated (if available) is an extra level of protection so that you can perform a rapid rollback, if necessary.

Stop all I/O from the host to the LUNs on the storage that is migrated

Before the actual migration can begin, the application must be quiesced and the file system must be synchronized to ensure that the application data is in a consistent state. Because the host might need to be rebooted several times before the application data becomes available again, the following steps might be required:

- Set applications to *not* automatically start when the host operating system restarts.
- Stop file systems from being automatically remounted upon system boot.

Note: In clustered environments, such as Windows or Linux, you might choose to work with only one node until the migration is complete. If so, consider shutting down all other nodes in the cluster.

4.2.3 Define and test data migration volumes

To define and test the data migration volumes, perform these tasks:

- Remapping old volumes to the new host in IBM FlashSystem A9000
- Define migration volumes in IBM FlashSystem A9000R
- Test data migration volumes in IBM FlashSystem A9000R

Remapping old volumes to the new host in IBM FlashSystem A9000

Complete the following steps to remap the old volumes to the new host:

Important: Make sure the LUNs that are migrated do not have any mappings to another host, as this can cause data corruption.

Use the Volume Mapping windows to map the IBM FlashSystem A9000 volumes to the host created for the IBM FlashSystem A9000R connection as shown in Figure 4-8.

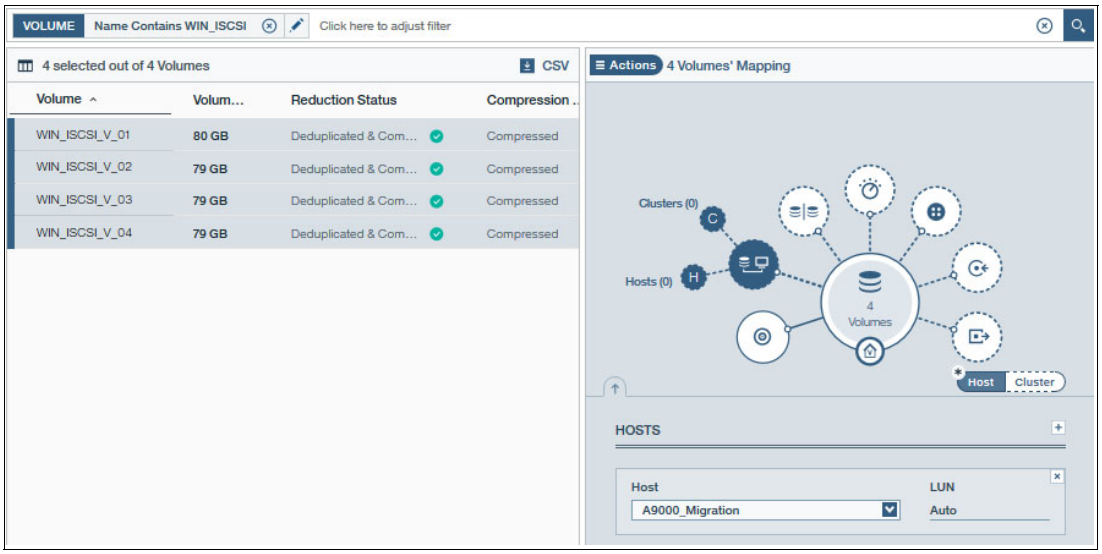

Figure 4-8 Mapping the volumes to IBM FlashSystem A9000R

Define migration volumes in IBM FlashSystem A9000R

IBM FlashSystem A9000R can determine the size of the IBM FlashSystem A9000 (or any iSCSI-attached storage) volumes and create the corresponding IBM FlashSystem A9000R volumes quickly when the data migration object is defined.

Use this method to help avoid potential problems when you manually calculate the real block size of a volume.

Complete these steps:

1. In the Hyper-Scale Manager GUI, select the pool to use for the migrated volumes and then select **Actions** → **Migration** → **Migrate remote volume to Pool**, as shown in Figure 4-9.

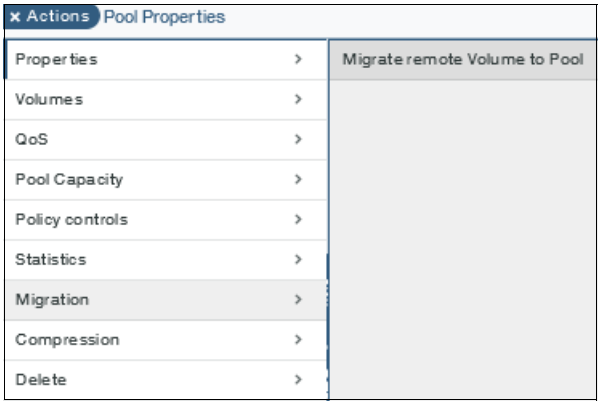

Figure 4-9 Selecting Pool for Migration in IBM FlashSystem A9000R

2. Select the settings for the following entries, as shown in Figure 4-27:

Target	Select the target system as defined in “Define IBM FlashSystem A9000 on IBM FlashSystem A9000R as a target” on page 108.
Target LUN	This field needs to match the remote logical unit number (LUN ID), for example, 1, 2 or 3.
Destination Volume	Provide a name for the destination volume, for example, `WIN_ISCSI_M_1`.
Create Volume	Select the check box if you want to create destination volumes.
Update Source	Select this option if you want to update the source volumes when the hosts make changes.
Active	Select this box if you want to activate the migration.

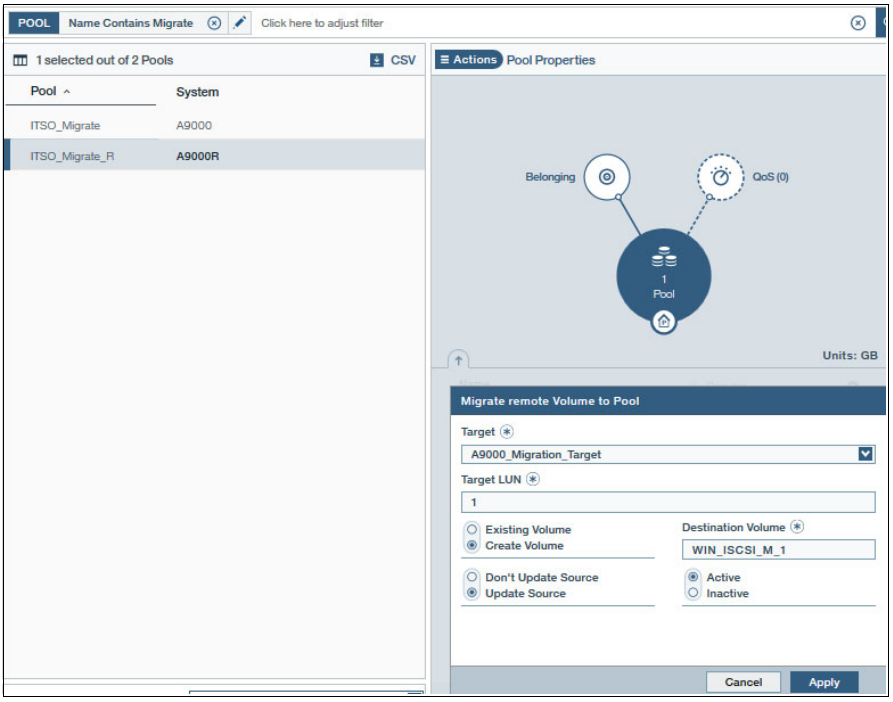

Figure 4-10 Creating migration volumes in IBM FlashSystem A9000R

3. Click **Apply**.
4. Repeat steps 2 and 3 for all of the migrated volumes with their respective remote LUNs.

After all of the migrated volumes are created, you can list the migrated volumes by using a migration volume filter, as shown in Figure 4-11.

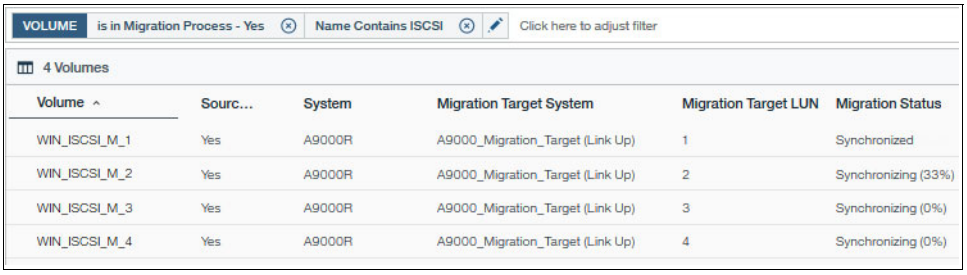

Figure 4-11 Migrated volumes view

Test data migration volumes in IBM FlashSystem A9000R

Select the migration volumes and then click **Actions** → **Test**.

The test action reads the first block of the volume and can therefore make sure that the system is able to read from the mapped volume. In some active/passive storage systems it is possible to read the configuration over the passive controller, but not read the data. The test is making sure there is connectivity to the controller that owns the LUN. If any issues exist with the data migration object, the test fails and the issues that are encountered are reported.

4.2.4 Activating data migration on the migrated volumes

If you selected **Inactive** in the Pool Properties window when creating the migration volumes on IBM FlashSystem A9000R, you have to start the migration process manually. When the data migration is initiated, the data is copied sequentially in the background from the IBM FlashSystem A9000 volume to IBM FlashSystem A9000R. The host reads and writes data to the new storage system without being aware of the background I/O that is being performed. If a write to a block comes from host that has not been migrated yet, the source storage system also is updated with this write.

Important: After the data migration is activated, the data migration can be inactivated. However, we do *not* advise that you inactivate it. When the data migration is inactivated, the host is no longer able to read or write to the source migration volume and all host I/O stops. *Do not inactivate the migration with host I/O running.*

When migration is finished choose **Disconnect** and not inactivate.

To activate the data migration, complete these steps:

1. Select the Migration volumes by using a volume migration filter.
2. Select **Volumes** → **Actions** → **Migration** → **Change Activation State** → **Activate**.
3. Select **Active,** as shown in Figure 4-12.

Figure 4-12 Activating migration for the migration volumes

4.2.5 Define the host in IBM FlashSystem A9000R and bring the host online

Before you complete the data migration and allocate the volumes to the host, the host must be defined in IBM FlashSystem A9000R. Volumes are then mapped to the hosts or clusters. Upon completion of these steps, the volumes are available to the host for read and write operations. If the host cannot view the volumes, ensure that the Host Attachment Kit (HAK) is installed and configured and that no other multipath software is installed on the host.

To attach a host to the volumes that are being migrated, complete these steps:

1. Define the host in IBM FlashSystem A9000R.

   Use the **New Item** → **Host** window to create a new host and add ports to the host.

2. Map the LUNs that are being migrated to the host.

   Use the **Volume** → **Actions** → **Mapping** → **View/Modify** window to map volumes to the host.

4.2.6 Complete the data migration on IBM FlashSystem A9000R

After the volumes and data access are confirmed, you can start the host application and verify the operations. The migration tasks run in the background and allow normal host access to the newly mapped volumes.

Make sure, that all of the devices come back online and that the applications can start.

Occasionally, a host might not need to be online during the migration (such as after hours, not in production, or if the migration is completed within the customer change window). It can be offline and be brought back online after the migration is complete.

Note: In clustered environments, the usual recommendation is for only one node of the cluster to be initially brought online after the migration is started, and to have all other nodes stay offline until the migration is complete. After completion, update all other nodes (drivers and Host Attachment Kit) in the same way that the primary node was updated during the initial outage.

Complete the data migration on IBM FlashSystem A9000R

To complete the data migration, complete the steps that are described in this section.

Data migration progress

Figure 4-11 shows one volume migrated, one volume migrating. The process is sequential so when one volume completes its migration, the next volume initializes. Figure 4-13 shows that all volumes have completed the migration and have the migration status **Synchronized**.

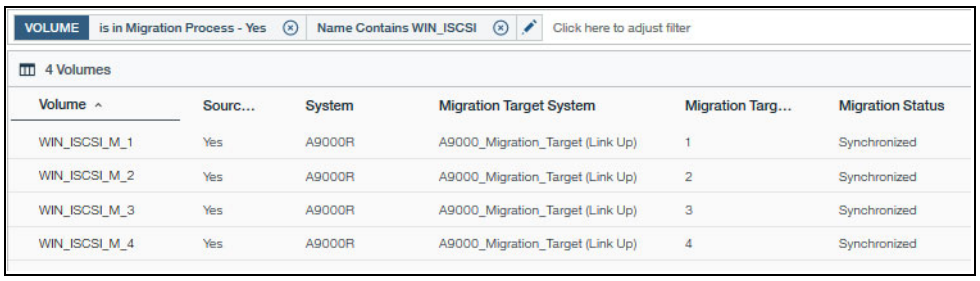

Volume	Sourc...	System	Migration Target System	Migration Targ...	Migration Status
WIN_ISCSI_M_1	Yes	A9000R	A9000_Migration_Target (Link Up)	1	Synchronized
WIN_ISCSI_M_2	Yes	A9000R	A9000_Migration_Target (Link Up)	2	Synchronized
WIN_ISCSI_M_3	Yes	A9000R	A9000_Migration_Target (Link Up)	3	Synchronized
WIN_ISCSI_M_4	Yes	A9000R	A9000_Migration_Target (Link Up)	4	Synchronized

Figure 4-13 Data migration process completed

After all of a volume's data is copied, the data migration achieves synchronization status. After synchronization is achieved, all read requests are served by the new IBM FlashSystem A9000R. If source updating was selected, IBM FlashSystem A9000R continues to write data to itself and the storage system that is being migrated until the data migration is terminated.

Disconnect data migration

After synchronization is achieved, the data migration object can be safely disconnected without host interruption. Verify that the migration is complete before you proceed. Select the volume and verify the migration status as displayed in Figure 4-13.

Important: If this migration is an online migration, do *not* inactivate the data migration before deletion because inactivation causes host I/O to stop, which can cause data corruption.

You can perform this task without host or server interruption if the volume completed synchronization. Select the volumes you want to disconnect and chose the disconnect action, as shown in Figure 4-14. If the data migration is incomplete, a warning message appears.

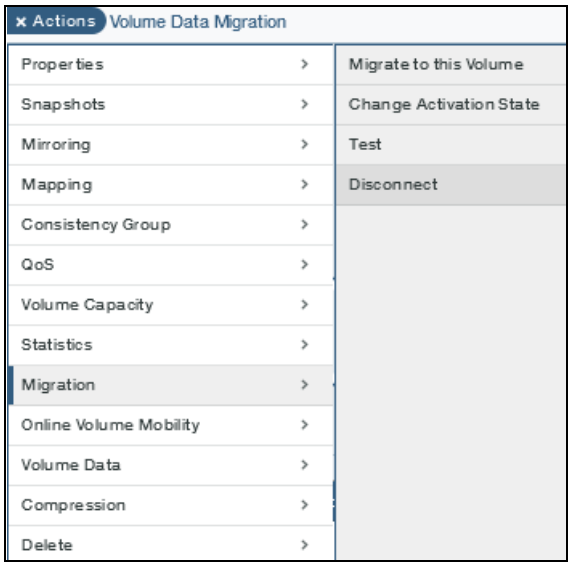

Figure 4-14 Disconnect data migration

Data migration is now completed.

4.2.7 Removing Migration Connectivity

The steps to remove the migration connectivity on A9000R for an iSCSI connection are the same steps as for an Fibre Channel connection as described in 4.3.7, "Removing Migration Connectivity".

4.3 Detailed migration steps for a Fibre Channel connection

Note: We chose to migrate from an IBM XIV Storage System to IBM FlashSystem A9000. These systems were used for illustration because they were available for our experimentation. Migration can be performed from any storage system that can attach to a Linux host. The setup details differ based on the storage system.

At a high level, we used the following steps to migrate volumes from an XIV Storage System (or any FC attached storage system) to IBM FlashSystem A9000:

1. Initial connection and preparation:
   a. Establish connection between the XIV Storage System (or any storage system) and IBM FlashSystem A9000.
   b. Prepare a World Wide Port Name (WWPN) table for all of the systems that are involved.
   c. Assume that the volumes and the host are already defined in the XIV Storage System.
   d. Define IBM FlashSystem A9000R as host object on the XIV Storage System with type **Default**.
   e. Define the XIV Storage System in IBM FlashSystem A9000R (as a migration target).
2. Perform pre-migration tasks for each host that is attached to the XIV that will be migrated:
   a. Back up host data.
   b. Stop all I/O from the host to the LUNs on the XIV storage.
   c. Update the Host Attachment Kit (HAK) to latest version

      Important: If your migration source is not an XIV, IBM FlashSystem A9000, or IBM FlashSystem A9000R, it is important that you uninstall any specific multipath software (for example IBM SDD) and install the IBM Host Attachment Kit (HAK).

   d. A reboot of the host might be needed depending on the multipath software changes requirements.
3. Define and test the data migration volume:
   a. On the XIV storage, remap volumes from the old host to IBM FlashSystem A9000.
   b. On IBM FlashSystem A9000, create the data migration tasks and test them.
4. Activate the data migration tasks on IBM FlashSystem A9000.
5. Define the host in IBM FlashSystem A9000 and bring the host and applications online:
   a. Define the host on IBM FlashSystem A9000.
   b. Map volumes to the host on IBM FlashSystem A9000.
   c. Verify that IBM FlashSystem A9000 storage is visible.
   d. Depending on the host a rescan on the host for the volumes on IBM FlashSystem A9000 might be necessary. Enable and start the host applications.
6. Complete the data migration on IBM FlashSystem A9000:
   a. Monitor IBM FlashSystem A9000 migration tasks.
   b. On completion, delete the migration tasks.

Tip: Print these overview steps and refer to them as you perform a migration.

The following steps are specific to the example that is mentioned here and focus only on the FC connection between an XIV Storage System (or any storage system) and IBM FlashSystem A9000 (or A9000R).

4.3.1 Initial connection and preparation

For the initial connection setup, start by connecting IBM FlashSystem A9000 to the XIV system that you are migrating from.

Because the migration is through FC, you must ensure a correct Storage Area Network (SAN) configuration between the XIV Storage System and IBM FlashSystem A9000. Also data connection between IBM FlashSystem A9000 and the host must be correct.

It is also a good opportunity for you to update the host operating system (OS) patches and drivers to the latest supported levels for the storage system that you are migrating (the XIV Storage System in our scenario).

Prepare a Fibre Channel WWPN table

A preferred practice is to document all WWPNs of the systems that are involved, as shown in Table 4-2.

Table 4-2 Example of FC ports

System	WWPN	TYPE
IBM FlashSystem A9000	50:01:73:80:56:73.01:14	Initiator
IBM FlashSystem A9000	50:01:73:80:56:73.01:34	Initiator
XIV Storage System	50:01:73:80:9c:4a:01:60	Target
XIV Storage System	50:01:73:80:9c:4a:01:90	Target

Volumes and host that are defined in the XIV Storage System

The example assumes that the pools and the volumes that you want to migrate existed in the XIV Storage System. This example also assumes that the host that uses the volumes existed in the XIV Storage System.

Define IBM FlashSystem A9000 on the XIV Storage System as host

IBM FlashSystem A9000 must be defined as a **Default** host on the XIV Storage System. In the XIV Storage System GUI, click **Hosts and Clusters** → **Hosts and Clusters** (Figure 4-15).

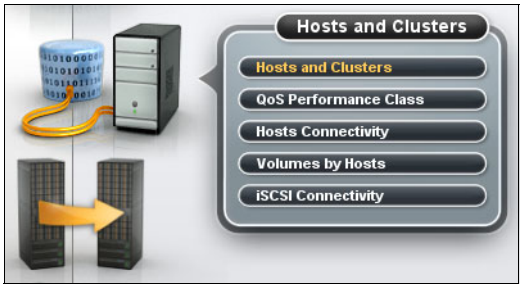

Figure 4-15 Hosts and Clusters in the XIV Storage System GUI

Choose **Add Host** from the top menu bar. Add a default host (`Data_Migration_Host`) with the name of the target IBM FlashSystem A9000R, as shown in Figure 4-16 on page 119.

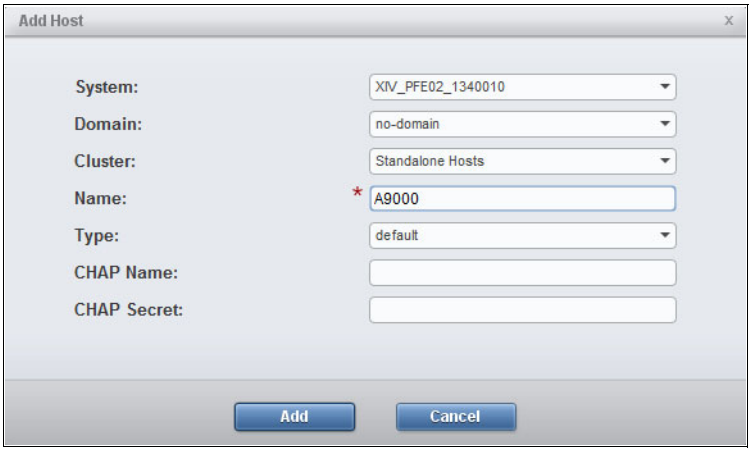

Figure 4-16 Add Host window for the data migration host

On the Add Port window, add the host ports with the WWPNs of the IBM FlashSystem A9000, as shown in Figure 4-17.

Figure 4-17 Adding the FC port connections for the data migration host

Repeat steps in Figure 4-17 for all IBM FlashSystem A9000 ports listed in Table 4-2. Figure 4-18 shows the two ports of the table added to the host.

Figure 4-18 Host and associated FC ports

Define the XIV Storage System on IBM FlashSystem A9000 as a target

After the physical connectivity is established and IBM FlashSystem A9000 is defined on the XIV Storage System (or any FC-attached storage), you must also define the XIV Storage System on IBM FlashSystem A9000 as a target.

Define the connectivity between IBM FlashSystem A9000 and the XIV Storage System via FC. Complete these steps:

1. By using the Hyper-Scale Manager GUI, log in to IBM FlashSystem A9000 system with the storage administrator credentials.
2. Click **Actions** → **Targets** → **View/Modify Targets** and click the icon to add a target, as shown in Figure 4-19.

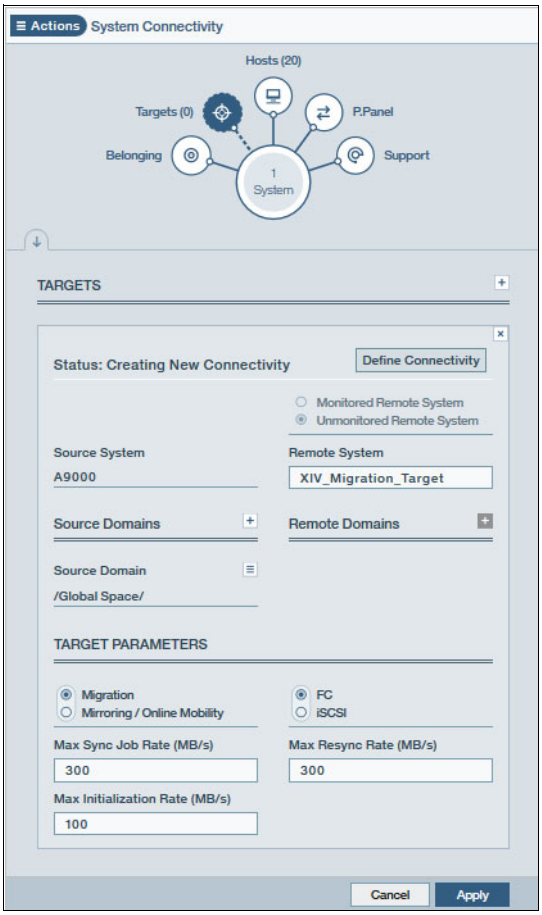

Figure 4-19 Defining a Target in IBM FlashSystem A9000

3. Enter the following information and select the following options, as shown in Figure 4-19:

Remote System	This field is for the current system, for example, `XIV_Migration_Target`. You have to select **Unmonitored Remote System** to be able to enter the name.
Source Domains	The default is **/Global Space/**, which means *no-domain* unless specified.
Remote Domain	Depending on the previous selection of domain, other options might be available.
Target Type	Two options exist: Migration and Mirroring. Select **Migration**.
Target Protocol	Fibre Channel (FC) or iSCSI. For this example, **FC** is selected.
Max. Sync Job Rate	Depending on the network bandwidth, the default value is set to 300 MBps. Not applicable for migration.
Max. Resync Rate	Depending on the network bandwidth, the default value is set to 300 MBps. Not applicable for migration.
Max. Init Rate	Depending on the network bandwidth, the default value is set to 100 MBps.

4. After the target is defined, the target appears in the System Connectivity display, as shown in Figure 4-20.

   Select **Define Connectivity** to define the FC connections.

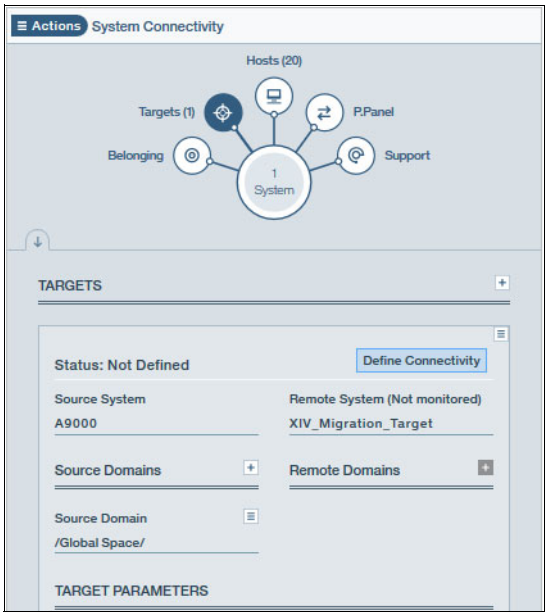

Figure 4-20 Defining connectivity in IBM FlashSystem A9000

5. In the window that is displayed, complete the following steps:
   a. Enter the WWPN as listed in Table Table 4-2 of the XIV Storage System to define the connection between the two systems, as shown in Figure 4-21.

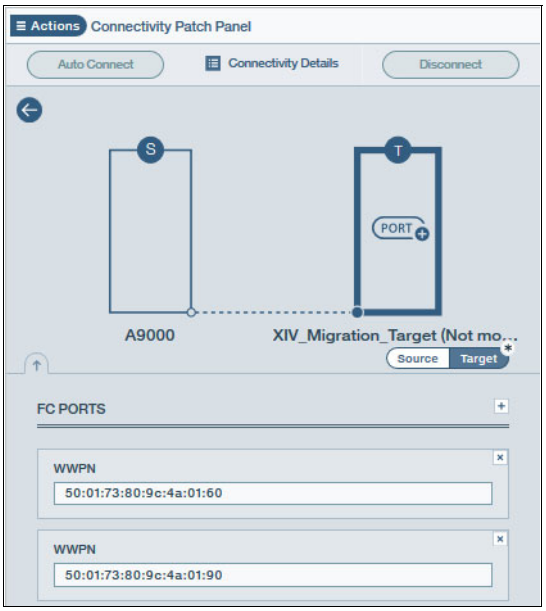

Figure 4-21 Define the FC ports for the target system

   b. Select **Auto Connect** to establish the connections between IBM FlashSystem A9000 and the target system.
   c. If the connection between IBM FlashSystem A9000 and the target system (the XIV Storage System) is successful, the connection is displayed as connected (with a check mark) and you can select **Connectivity Details** to view the connectivity, as shown in Figure 4-22.

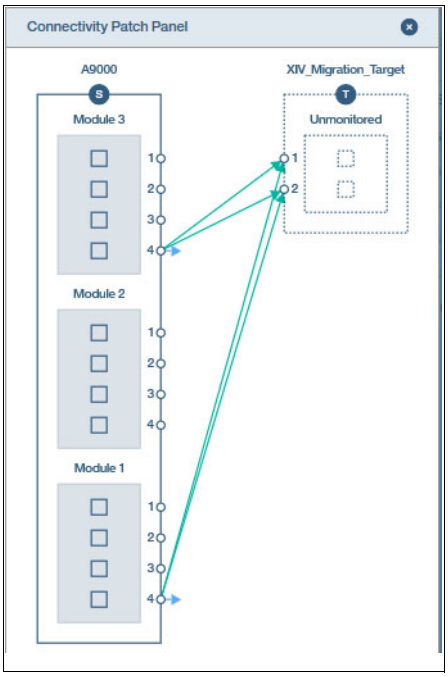

Figure 4-22 Migration connectivity details

4.3.2 Perform pre-migration tasks for each host that is migrated

The following pre-migration tasks are preferred practices to be performed for each host that you are migrating:

- Back up the volumes that are migrated.
- Stop all I/O from the host to the LUNs on the storage that is migrated.

Back up the volumes that are migrated

Its good practice to take a full restorable backup before any data migration activity. Verify that all of the data is restorable from the backup and that no backup media errors exist.

In addition to a regular backup, a point-in-time copy of the LUNs that are being migrated (if available) is an extra level of protection so that you can perform a rapid rollback, if necessary.

Stop all I/O from the host to the LUNs on the storage that is migrated

Before the actual migration can begin, the application must be quiesced and the file system must be synchronized to ensure that the application data is in a consistent state. Because the host might need to be rebooted several times before the application data becomes available again, the following steps might be required:

- Set applications to *not* automatically start when the host operating system restarts.
- Stop file systems from being automatically remounted on boot.

Note: In clustered environments, such as Windows or Linux, you might choose to work with only one node until the migration is complete. If so, consider shutting down all other nodes in the cluster.

4.3.3 Define and test data migration volumes

To define and test the data migration volumes, you perform these tasks:

- Remapping old volumes to the new host in the XIV Storage System
- Define migration volumes in IBM FlashSystem A9000
- Test data migration volumes in IBM FlashSystem A9000R

Remapping old volumes to the new host in the XIV Storage System

Complete the following steps to remap the old volumes to the new host:

Important: Make sure the LUNs that is going to be migrated do not have any mappings to an old host, this can cause data corruption.

1. In the XIV Storage System GUI, log in with the storage administrator credentials.
2. Click **Hosts and Clusters** → **Hosts and Clusters**.
3. Select the data migration host in the XIV Storage System that was defined in “Define the XIV Storage System on IBM FlashSystem A9000 as a target” on page 120.
4. Right-click the selected host, and click **Modify LUN Mapping**.
5. Select the volumes to be mapped (the volumes that were earlier mapped to the actual host), as shown in Figure 4-23, and click **Map**.

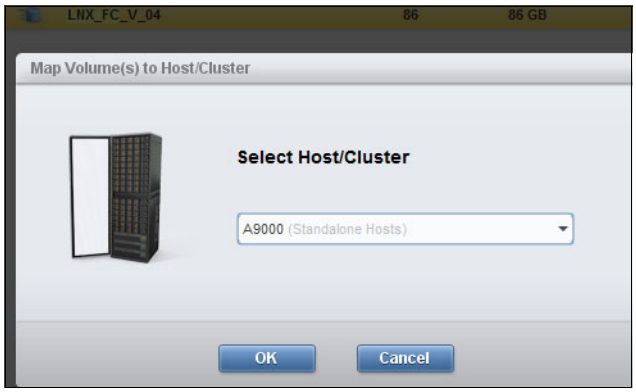

Figure 4-23 LUN mapping to the data migration host

Figure 4-24 shows the selected LUNs mapped to IBM FlashSystem A9000.

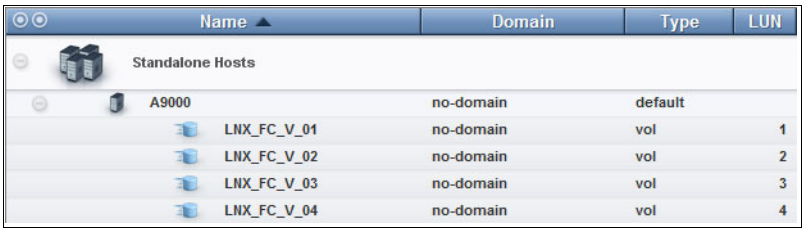

Figure 4-24 Luns mapped to IBM FlashSystem A9000

Define migration volumes in IBM FlashSystem A9000

IBM FlashSystem A9000 can determine the size of the XIV Storage System (or any iSCSI-attached storage) volumes and create the corresponding IBM FlashSystem A9000 volumes quickly when the data migration object is defined.

Use this method to help avoid potential problems when you manually calculate the real block size of a volume.

Complete these steps:

1. In the Hyper-Scale Manager GUI, select the pool to use for the migrated volumes and then select **Actions** → **Migration** → **Migrate remote Volume to Pool,** as shown in Figure 4-25.

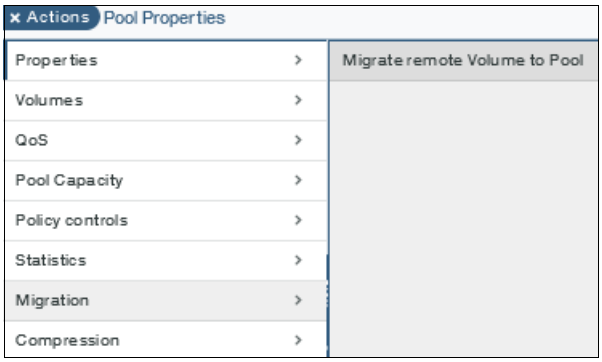

Figure 4-25 Selecting Migration in IBM FlashSystem A9000

2. Get the LUN IDs of the volumes mapped to host A9000 on the XIV Storage System as shown in Figure 4-26.

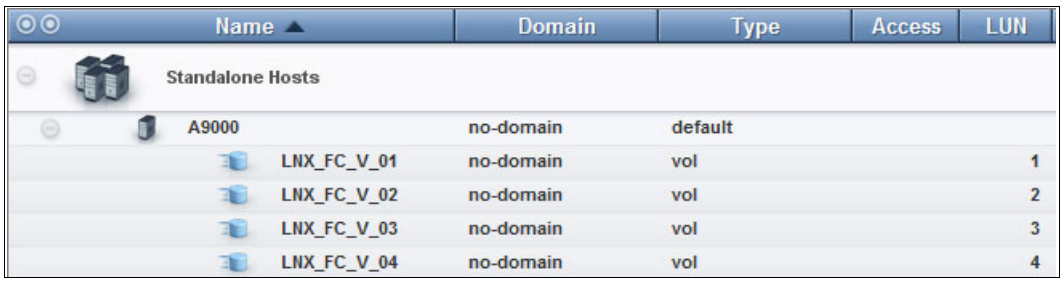

Figure 4-26 XIV Storage System LUN IDs for host A9000

Select the settings for the following entries, as shown in Figure 4-27 on page 126:

Target Select the target system that was defined earlier in the XIV Storage System.

Target LUN This field needs to match the remote logical unit number (LUN ID), for example, 1, 2, 3, and 4.

Destination Volume Provide a name for the destination volume, for example, `LNX_FC_M01`.

Create Volume Select the check box if you want to create destination volumes.

Update Source Select this option if you want to update the source volumes when the hosts make changes.

Active Select this box if you want to activate the migration.

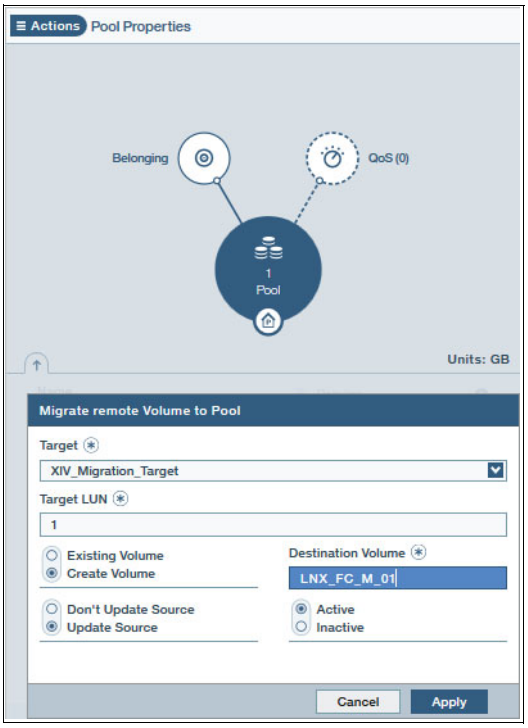

Figure 4-27 Creating migration volumes in IBM FlashSystem A9000

3. Click **Apply**.
4. Repeat steps 2 and 3 for all of the migrated volumes with their respective remote LUNs.

After all of the migrated volumes are created, you can list the migrated volumes by using a migration volume filter, as shown in Figure 4-28.

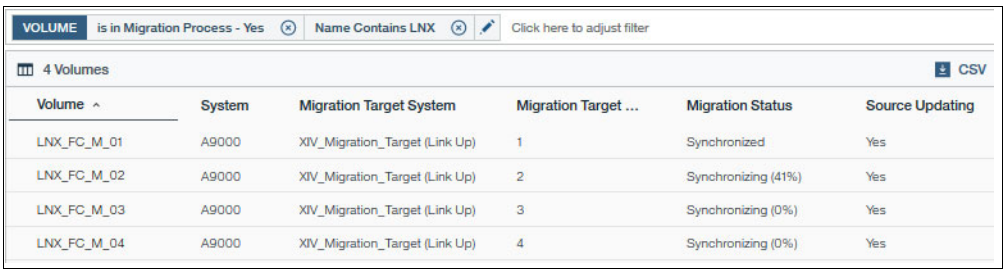

Volume	System	Migration Target System	Migration Target ...	Migration Status	Source Updating
LNX_FC_M_01	A9000	XIV_Migration_Target (Link Up)	1	Synchronized	Yes
LNX_FC_M_02	A9000	XIV_Migration_Target (Link Up)	2	Synchronizing (41%)	Yes
LNX_FC_M_03	A9000	XIV_Migration_Target (Link Up)	3	Synchronizing (0%)	Yes
LNX_FC_M_04	A9000	XIV_Migration_Target (Link Up)	4	Synchronizing (0%)	Yes

Figure 4-28 Migrated volumes view

Test data migration volumes in IBM FlashSystem A9000R

Select the migration volumes and then select **Actions** → **Test**.

The test action reads the first block of the volume and can therefore make sure that the system is able to read from the mapped volume. In some active and passive storage systems it is possible to read the configuration over the passive controller, but not read the data. The test is making sure that there is connectivity to the controller that owns the LUN. If any issues exist with the data migration object, the test fails and the issues that are encountered are reported.

4.3.4 Activating data migration on the migrated volumes

If you checked Inactive in the Pool Properties window when creating the migration volumes on IBM FlashSystem A9000R, you have to start the migration process manually. When the data migration is initiated, the data is copied sequentially in the background from the XIV Storage System volume to IBM FlashSystem A9000. The host reads and writes data to the new storage system without being aware of the background I/O that is being performed. If a write to a block comes from host that has not been migrated yet, the source storage system is also updated with this write.

Important: After the data migration is activated, the data migration can be inactivated. However, we do *not* advise that you inactivate it. When the data migration is inactivated, the host is no longer able to read or write to the source migration volume and all host I/O stops. *Do not inactivate the migration with host I/O running.*

When migration is finished, choose **Disconnect** and not inactivate.

To activate the data migration, complete these steps:

1. Select the Migration volumes by using a volume migration filter.
2. Select **Volumes** → **Actions** → **Migration** → **Change Activation State** → **Activate**.
3. Select **Active**, as shown in Figure 4-29.

Figure 4-29 Activating migration for the migration volumes

4.3.5 Define the host in IBM FlashSystem A9000 and bring the host online

Before you complete the data migration and map the volumes to the host, the host must be defined in IBM FlashSystem A9000. Volumes are then mapped to the hosts or clusters. Upon the completion of these steps, the volumes are available to the host for read and write operations. If the host cannot view the volumes, ensure that the Host Attachment Kit (HAK) is installed and configured and no other multipath software is installed on the host.

To attach a host to the volumes that are being migrated, complete these steps:

1. Define the host to IBM FlashSystem A9000R.

   Use the **New Item** → **Host** window to create a new host and add the ports to the host.

2. Map the LUNs that are being migrated to the host.

   Use the **Volume** → **Actions** → **Mapping** → **View/Modify** window to map volumes to the host.

4.3.6 Complete the data migration on IBM FlashSystem A9000

After the volumes and data access are confirmed, you can start the host application and verify the operations. The migration tasks run in the background and allow normal host access to the newly mapped volumes.

Make sure that all of the devices come back online and that the applications can start.

Occasionally, a host might not need to be online during the migration (such as after hours, not in production, or if the migration is completed within the customer change window). It can be offline and be brought back online after the migration is complete.

Note: In clustered environments, the usual recommendation is for only one node of the cluster to be initially brought online after the migration is started, and to have all other nodes stay offline until the migration is complete. After completion, update all other nodes (driver, host attachment package, and so on) in the same way that the primary node was updated during the initial outage.

Complete the data migration on IBM FlashSystem A9000

To complete the data migration, complete the steps that are described in this section.

Data migration progress

Figure 4-28 on page 126 shows one volume migrated, one volume migrating. The process is sequential so when one volume completes its migration, the next volume initializes.

Figure 4-30 shows that all volumes have completed the migration and have the migration status **Synchronized**.

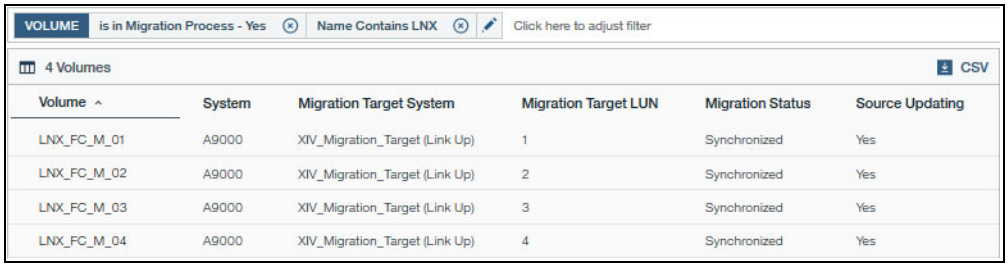

Volume	System	Migration Target System	Migration Target LUN	Migration Status	Source Updating
LNX_FC_M_01	A9000	XIV_Migration_Target (Link Up)	1	Synchronized	Yes
LNX_FC_M_02	A9000	XIV_Migration_Target (Link Up)	2	Synchronized	Yes
LNX_FC_M_03	A9000	XIV_Migration_Target (Link Up)	3	Synchronized	Yes
LNX_FC_M_04	A9000	XIV_Migration_Target (Link Up)	4	Synchronized	Yes

Figure 4-30 Data migration process completed

After all of a volume's data is copied, the data migration achieves synchronization status. After synchronization is achieved, all read requests are served by the new IBM FlashSystem A9000. If source updating was selected, IBM FlashSystem A9000R continues to write data to itself and the storage system that is being migrated until the data migration is terminated.

Disconnect data migration

After synchronization is achieved, the data migration object can be safely disconnected without host interruption. Verify that the migration is complete before you proceed. Select the volume and verify the migration status as displayed in Figure 4-30.

Important: If this migration is an online migration, do *not* inactivate the data migration before deletion because inactivation causes host I/O to stop, which can cause data corruption.

You can perform this task without host/server interruption if the volume completed synchronization. Select the volumes you want to disconnect and choose the disconnect action, as shown in Figure 4-31. If the data migration is incomplete, a warning message appears.

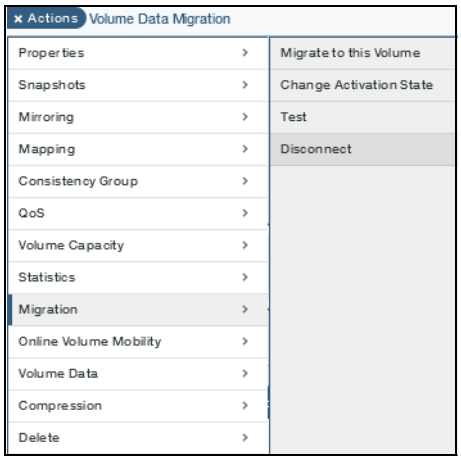

Figure 4-31 Disconnect data migration

Data migration is now completed.

4.3.7 Removing Migration Connectivity

To remove the FC migration connectivity, complete these steps:

- Remove migration Target on A9000
- Remove the host A9000 used for migration on the XIV Storage System.

Remove migration Target on A9000

In Hyper-Scale manager select A9000 in the System view and list the targets in the System Connectivity windows:

1. Select **Actions** → **Targets** → **View/Modify Targets** → **Actions Menu** of the migration target `XIV_Migration_Target`, as shown in Figure 4-32.

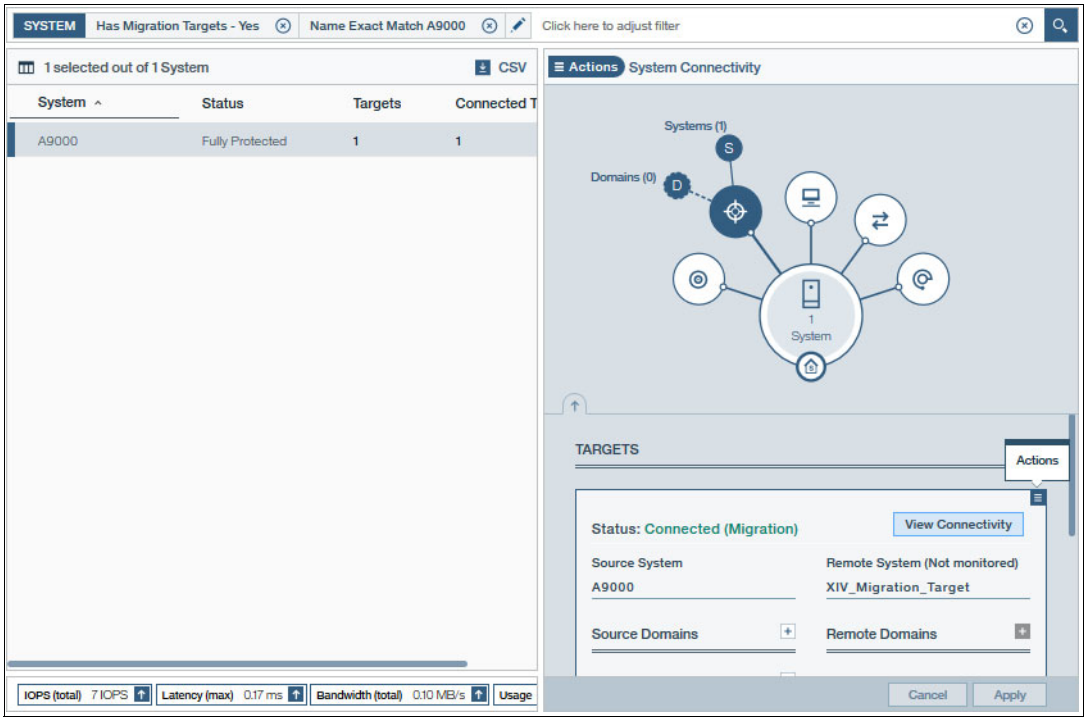

Figure 4-32 Actions Menu in the System Connectivity window

2. Select Disconnect from the actions menu, as shown in Figure 4-33.

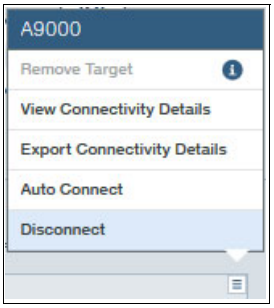

Figure 4-33 Disconnection migration target

3. Using the same action menu as shown in Figure 4-32 on page 130, select Remove Target, as shown in Figure 4-34.

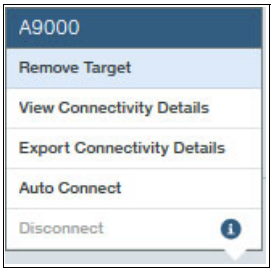

Figure 4-34 Removing migration Target

After successful target removal the target is not listed anymore in the System Connectivity window.

Remove the host A9000 on the XIV Storage System

To remove the host, complete the following steps:

1. In the XIV GUI select the volumes mapped to **A9000** and use the Unmap action to unmap the volumes from the host **A9000,** as shown in Figure 4-35.

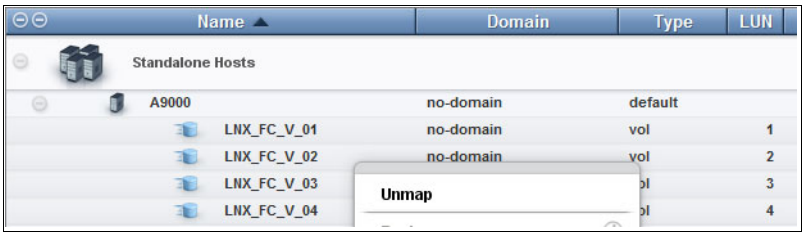

Figure 4-35 Unmap volumes

2. In the XIV GUI select the host **A9000** and use the Delete action to delete the host used for migration as shown in Figure 4-36.

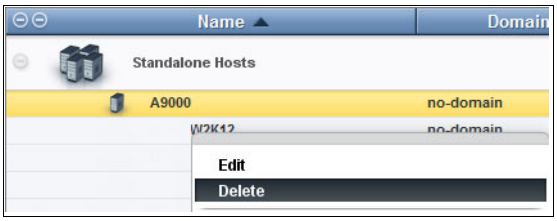

Figure 4-36 Deleting the host used for migration.

The migration connectivity is now deleted.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *IBM FlashSystem A9000 and IBM FlashSystem A9000R Architecture and Implementation,* SG24-8345
- *IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM XIV Storage System: Host Attachment and Interoperability*, SG24-8368
- *IBM XIV Storage System Business Continuity Functions*, SG24-7759
- *IBM Hyper-Scale Manager for IBM Spectrum Accelerate Family IBM XIV, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate*, SG24-8376

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

`http://www.ibm.com/redbooks`

Other publications and Online resources

- IBM FlashSystem A9000 on IBM Knowledge Center:

  `http://www.ibm.com/support/knowledgecenter/STJKN5`

  The following publications are at this website:

  - *IBM FlashSystem A9000 Command-Line Interface (CLI) Reference Guide*, SC27-8559
  - *IBM FlashSystem A9000 Product Overview,* GC27-8583-00
  - *Hyper-Scale Manager 5.0 REST API Specification*, SC27-6440-01
  - *Hyper-Scale Manager 5.0 User Guide*, SC27-8560
  - *IBM FlashSystem A9000 Models 9836-415 and 9838-415 Deployment Guide*, GC27-8564
- IBM FlashSystem A9000R on IBM Knowledge Center:

  `http://www.ibm.com/support/knowledgecenter/STJKMM`

  The following publications are at this website:

  - *IBM FlashSystem A9000R Command-Line Interface (CLI) Reference Guide*, SC27-8711
  - *IBM FlashSystem A9000R Product Overview,* GC27-8558-00
  - *IBM FlashSystem A9000R Model 415 Deployment Guide*, GC27-8565

- IBM FlashSystem A9000 product page:

  http://www.ibm.com/systems/storage/flash/a9000
- IBM FlashSystem A9000R product page:

  http://www.ibm.com/systems/storage/flash/a9000R
- IBM Fix Central:

  http://www.ibm.com/support/fixcentral/
- For the latest hardware and software requirements, visit the IBM System Storage Interoperation Center (SSIC) website:

  http://www.ibm.com/systems/support/storage/ssic/interoperability.wss
- IBM Offering Information page (announcement letters and sales manuals):

  http://www.ibm.com/common/ssi/index.wss?request_locale=en
- On this page, enter A9000, select the information type, and click **Search**. On the next page, narrow your search results by geography and language.

Help from IBM

IBM Support and downloads

http://www.ibm.com/support

IBM Global Services

http://www.ibm.com/services

REDP-5401-00

ISBN DocISBN

Printed in U.S.A.