IBM

IBM Hyper-Scale in XIV Storage

- Powerful new approach to storage management scalability
- Volume mobility and consistency across XIV systems
- Capacity planning and reporting

Bertrand Dufrasne
Christian Burns
Wenzel Kalabza
Sandor Lengyel
Patrick Schill
Christian Schoessler

Redpaper

ibm.com/redbooks

International Technical Support Organization

IBM Hyper-Scale in XIV Storage

November 2014

REDP-5053-01

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

Second Edition (November 2014)

This edition applies to the IBM XIV Storage System Gen 3 Release 3.4 with XIV Software Version 11.4.

Contents

Preface

Organizations are struggling with an unprecedented surge in data, both user-created and machine-created. With an increasing reliance on this rapidly growing data, effective scaling of data storage represents a key challenge for all enterprises.

IBM® Hyper-Scale helps you easily overcome provisioning scenarios that normally challenge traditional systems. IBM Hyper-Scale can accommodate several critical client scenarios for centralized administration, data mobility, load balancing, over-provisioning, and storage system repurposing.

IBM Hyper-Scale introduces the following major concepts and technologies, with the IBM XIV® Storage System (Gen3 models):

- IBM Hyper-Scale Manager is a flexible, consolidated multi-system management application that builds upon the originally released *Multi-System Manager* with XIV Storage Software V11.2. IBM Hyper-Scale Manager is based on, and seamlessly integrated with, the XIV graphical user interface (GUI) and spans multiple XIV systems. It virtually transforms multiple systems into a single system, allowing customers to centrally manage up to 144 XIV systems.
- IBM Hyper-Scale Mobility was released with XIV Storage System V11.3. It is a powerful function for moving volumes between storage containers transparently, with no disruption to host applications.
- IBM Hyper-Scale Consistency was released with the IBM XIV Storage Software Version 11.4. It offers cross-system consistency, enabling coordinated snapshots across independent XIV systems. It helps to ensure data protection across multiple XIV systems.

This IBM Redpaper™ publication provides a broad understanding of the IBM Hyper-Scale feature. This publication is intended for XIV clients and users who want a practical understanding of IBM Hyper-Scale concepts and usage. For information about the IBM Hyper-Scale concepts, refer to the IBM white paper, *IBM Hyper-Scale and Its Implementation in XIV Storage: A powerful new approach to storage management scalability*:

http://public.dhe.ibm.com/common/ssi/ecm/en/tsl03121usen/TSL03121USEN.PDF

Authors

This paper was produced by a team of specialists from around the world working for the IBM International Technical Support Organization (ITSO), at the IBM European Storage Competence Center in Mainz, Germany.

Bertrand Dufrasne is an IBM Certified Consulting IT Specialist and Project Leader for IBM System Storage® disk products at the ITSO, San Jose Center. He has worked at IBM in various information technology (IT) areas. He has authored many IBM Redbooks® publications, and has also developed and taught technical workshops. Before joining the ITSO, he worked for IBM Global Services as an Application Architect. He holds a Master's degree in electrical engineering.

Christian Burns is an IBM Storage Solution Architect based in New Jersey. As a member of the Storage Solutions Engineering team in Littleton, MA, he works with clients, IBM Business Partners, and IBMers worldwide, designing and implementing storage solutions that include various IBM products and technologies. Christian's areas of expertise include IBM Real-time Compression™, SVC, IBM XIV, and IBM FlashSystem™. Before joining IBM, Christian was the Director of Sales Engineering at IBM Storwize®, before its becoming IBM Storwize. He brings over a decade of industry experience in the areas of sales engineering, solution design, and software development. Christian holds a BA degree in Physics and Computer Science from Rutgers College.

Wenzel Kalabza is a certified XIV Product Field Engineer (PFE) based in the storage competence center in Mainz, Germany. Wenzel joined IBM in 1998 as customer quality engineer for IBM disk drive failure and performance analysis. He joined the Back Office for the high end storage system (ESS) in June 2002. In 2005, Wenzel started a PFE role for the IBM disk storage DS6000™. In June 2008, he became a PFE for the XIV storage product. Wenzel holds a degree in Electrical Engineering and Power Economy, and several storage-related certifications.

Sandor Lengyel is an IBM Technical Advisor in the EMEA region for the XIV Storage System, based in Vac, Hungary. He joined IBM in 1998 as a Test Engineer and worked at IBM Vac Manufacturing for IBM System Storage disk products. In 2010, Sandor joined the Storage Solution Competence Center as a Technical Advisor and certified for XIV Administration and Specialist. Until 2011, he also supported the ProtecTier, Scale Out Network Attached Storage, and FlashSystem storage products. He holds a degree in Information Technology Engineering.

Patrick Schill is a Senior Certified Storage IT specialist within IBM GBS Federal Sector. He has 12 years of experience designing high performance, turnkey solutions supporting various complex host platforms, and application workloads. Patrick's expertise and focus surrounds x86 Virtualization: Hypervisors, System Storage, Microsoft messaging solutions, and a variety of high transactional I/O database designs that support cloud, and analytical solutions. Patrick holds a Bachelor of Science degree from DePaul University.

Christian Schoessler is an IT Specialist in the ATS Disk Solution Europe Team in Mainz, Germany. He has been working as a consultant for IBM for 13 years on different storage products. His Technical Sales Support within the ATS Team has focused for three years on XIV. Other areas of expertise include Flash, Performance and Copy Services calculations. He holds a degree in Physics from the Technical University of Darmstadt.

Thanks to the following people for their contributions to this project:

Diane Benjuya, Ramy Buechler, Barak Davidov, Rami Elron, Sagy Erlichman, Jim Sedgwick, Yossi Siles, George Thomas, and Moshe Weiss
IBM

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.

Obtain more information about the residency program, browse the residency index, and apply online at the following website:

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:

  ibm.com/redbooks

- Send your comments in an email:

  redbooks@us.ibm.com

- Mail your comments:

  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:

  https://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

Summary of changes

This section describes the technical changes made in this edition of the paper and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for IBM Hyper-Scale in XIV Storage
as created or updated on November 17, 2014.

November 2014, Second Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

New information

- Section 1.5, “XIV Capacity planning with Hyper-Scale Manager” on page 32
- Section 3.3, “Creating a cross-system consistency group by using GUI” on page 79

Notices

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

1

Managing XIV with the IBM Hyper-Scale Manager

This chapter covers the capabilities of the IBM Hyper-Scale Manager. Hyper-Scale Manager provides operational efficiency for large and multi-site deployments. It enables you to monitor and configure many systems simultaneously, providing the advantages of the XIV graphical user interface (GUI) features across multiple systems. The following topics are included:

- Hyper-Scale Manager server software
- Hyper-Scale Manager server software installation
- Hyper-Scale Manager overview
- Bulk configuration of multiple systems
- Clone existing system configuration
- Copy and paste operation
- Progress indicator

Important: Illustrations in this chapter mostly apply to an IBM XIV Storage System Gen3 fully configured with 2 terabyte (TB) drives.

For additional information about topics presented in this chapter, refer to the IBM Knowledge Center for XIV at:

http://www-01.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv_gen3kcwelcomepage.html?cp=STJTAG%2F1

1.1 IBM Hyper-Scale Manager software

The XIV Management GUI offers the following connection options:

- The XIV Management GUI direct connection
- The server connection, through the Hyper-Scale Manager server

The XIV Storage Management GUI direct connection is covered in the IBM Redbooks publication, *IBM XIV Storage System Gen3: Architecture, Implementation, and Usage*, SG24-7659.

The Hyper-Scale Manager enables the XIV GUI to access and operate on multiple XIV systems *concurrently*. This capability should not be confused with the ability to have multiple systems registered in the XIV GUI (in direct mode), which is only able to configure one system at a time.

The Hyper-Scale Manager server is used to communicate with the XIV Storage System Software, which in turn interacts with the XIV Storage System hardware. It reduces operational complexity and enhances capacity planning through integrated management of multiple XIV systems. It is ideal for large and multi-site XIV deployments.

The Hyper-Scale Manager provides the following functions:

- Allows integrated management of all XIV Storage Systems across the enterprise, including central configuration of user access rights, hosts, and event rules.
- Provides powerful health monitoring by integrating events and alerts across the managed XIV Systems.
- Supports extensive capacity reporting. The system collects capacity-related data over time. Such data can be processed later by analytics applications to reveal insights that could be used to improve capacity planning.
- Hyper-Scale Manager's GUI automatically opens the relevant view and highlights new objects (such as pools, volumes, hosts) after their creation.
- Extends management scalability to smartphones and tablets.

Hyper-Scale Manager also provides support for the RESTful application programming interface (API). It provides XIV administrators with an open management API that allows retrieval of information from all systems at once. The RESTful API support in XIV is described in the IBM Redpaper, *RESTful API Support in XIV*, REPD-5064.

Figure 1-1 depicts how the IBM Hyper-Scale Manager interacts with the XIV GUI and XIV Systems, through configuration options.

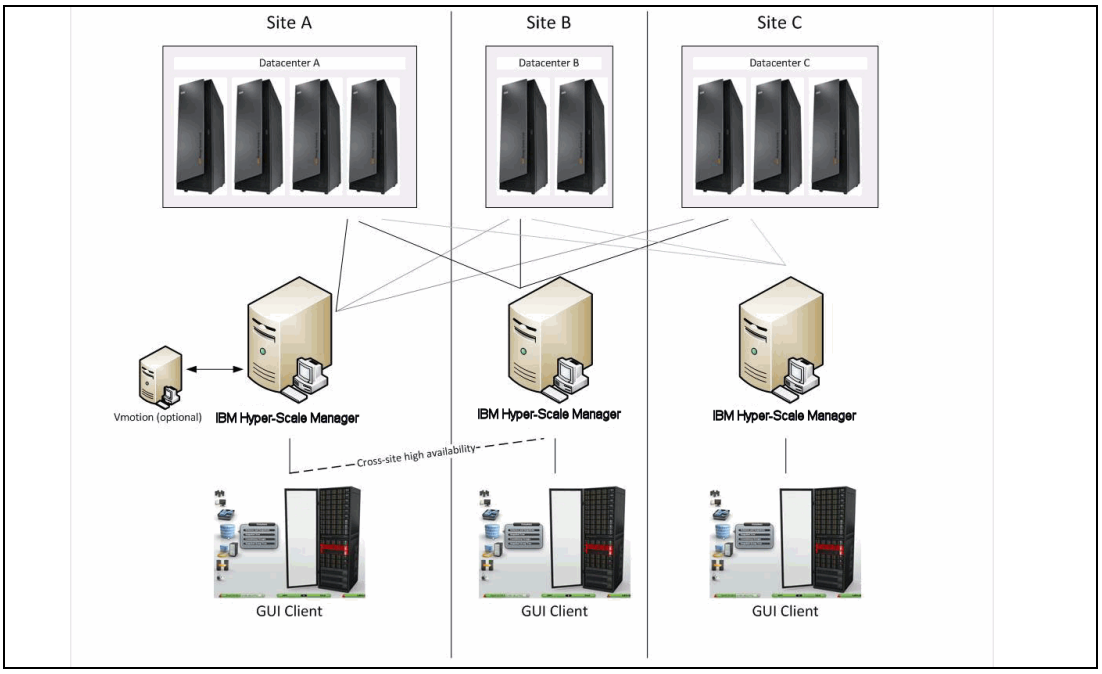

Figure 1-1 Hyper-Scale Manager options

The Hyper-Scale Manager server can exist as a single instance on a virtual machine server, or on several servers. The XIV GUI provides the interface to the Hyper-Scale Manager server. This results in the same functions available in direct mode to be expanded to monitor and configure multiple XIV Storage Systems collectively.

1.2 IBM Hyper-Scale Manager server installation

This section illustrates the step-by-step installation of the Hyper-Scale Manager server software. The Hyper-Scale Manager server is implemented as a VMware virtual appliance over a VMware ESX server (VMware Hypervisor only), or as a stand-alone installation on Red Hat Enterprise Linux (RHEL). The Hyper-Scale Manager server software can be downloaded from the following website:

`http://www.ibm.com/support/fixcentral`

From the successive pull-down menus, starting with the Product Group, make the appropriate selection to retrieve the XIV Storage System and appropriate version of the Management Tool (4.4.x.x).

1.2.1 Installing Hyper-Scale Manager server virtual appliance

Important: The minimum requirements for installing the Hyper-Scale Manager server software on a VMware ESX server (versions 4, i4, 4.1, i4.1, i5) are as follows:

Processor	Dual-core processor (Six-core processor is better)
Memory	4 gigabytes (GB) of RAM (8 GB of RAM or more is better)
VM Disk capacity	76 GB (thin-provisioned possible)
Screen resolution	1024 x 768 (1024 x 768 - 1920 x 1200 is better)

Table 1-1 and Table 1-2 describe additional specifications that are required for the installation and support of the Hyper-Scale Manager server.

Table 1-1 Hyper-Scale Manager server specifications

Specification Table
An ESX (VMware) Server that will host the Hyper-Scale Manager server
Distance between the Hyper-Scale Manager server and the XIV systems it monitors: Allowed latency maximum of 200 milliseconds (use ping timing to verify)
Maximum number of GUI connections to one server: 15
Maximum XIV Systems administered by one server: 144

Table 1-2 Hyper-Scale Manager server ports

Ports that need to be open	Port numbers
Outbound/Inbound: Remote Method Invocation (RMI)	1199 - 1209
Inbound only Secure Shell/Secure Copy Protocol (SSH/SCP)	22
Outbound only: Simple Mail Transfer Protocol (SMTP)	25
Communication with XIV systems	7778
RESTful API	8443

Graphical user interface specifications

When used with the Hyper-Scale Manager server, the XIV Storage Management GUI can be installed on a Windows workstation. The requirements remain the same as when used in direct mode.

IBM Hyper-Scale Manager software installation steps

Complete the following tasks to install the XIV Storage Management software for server mode:

1. Deploy the Hyper-Scale Manager server virtual appliance on a VMware vSphere ESXi server.
2. Install the Hyper-Scale Manager.

After those tasks, you must restart the system, and you can then start to initialize the Hyper-Scale Manager inventory, which is described in 1.3.3, “Using the IBM Hyper-Scale Manager features” on page 24.

Deploying the Hyper-Scale Manager server virtual appliance

A virtual machine (VM) must be deployed with the Hyper-Scale Manager server as a virtual appliance. The VM installation deploys the Open Virtualization Format (OVF*)* image of the Hyper-Scale Manager server on a VMware ESXi platform.

Follow this procedure to deploy the virtual appliance:

1. Deploy `*.ovf` on a VMWare ESX server.
2. Start the VM.
3. At the VMware console, log in with a user of `root` and a password of `xivmsRoot`.
4. The installation wizard opens. The wizard guides you through the Hyper-Scale Manager configuration.

Installing the Hyper-Scale Manager

Follow the installation wizard to install the Hyper-Scale Manager on the preinstalled VM. The wizard welcome dialog lists the installation steps, as shown in Figure 1-2.

```
--------------- IBM XIV Hyper-Scale Manager v1.x.x.x --------------
-------------------------------------------------------------------
Welcome to the IBM XIV Multi-System Manager installation procedure.
This setup will guide you through the installation process.

Installation Steps
--------------
Step 1: License Agreement
Step 2: Password Change
Step 3: Networking
Step 4: NTP & Regional Settings
Step 5: SMTP
Proceed? [Y/N] >
```

Figure 1-2 Welcome menu

Follow the listed steps to perform the installation:

1. The first step requires approval of the license agreement. Press Enter to view the Hyper-Scale Manager license agreement. After reading the license agreement, enter `1` to approve or `2` to decline it. You are not able to proceed with the installation or use the Hyper-Scale Manager if you do not approve the license agreement.

2. The second step of the installation is to change the root password. Proceed as shown in Figure 1-3.

```
Step 2 >> Password Change
------------------------
Please change root password.
Password must be at least 6 characters long
New password for root >
Type the password again for verification >
```

Figure 1-3 Password change

3. The next step is to configure the network. Select between Static Internet Protocol (`Static IP`) and Dynamic Host Configuration Protocol (`DHCP`), as illustrated in Figure 1-4 on page 6.

```
 Step 3: Network Configuration
#-----------------------------------
# Active network configuration:
#-----------------------------------
# IP mode
# Hostname
# Domain search list
# Domain name servers

1) Configure Static IP
2) Configure DHCP
3) Return (not supported, installation mode)
Setup/Network/IP>
```

Figure 1-4 Network settings panel example

Follow these substeps to complete the network configuration:

a. If Static IP was chosen, as shown in Figure 1-5, enter your IP address to continue.

```
---> Step 3: Networking
------------------------------

#-----------------------------------
# Active Network configuration:
#-----------------------------------
# IP mode
# Hostname
# Domain search list
# Domain name servers

Configure static IP address
Enter IP address []:
```

Figure 1-5 Network Input panel example

b. After the IP address is entered, enter the additional information regarding the static IP. Follow the prompts to configure the DNS addresses, as illustrated in Figure 1-6 on page 7.

```
Step 3 -Network Configuration
-----------------------------
#-----------------------------------
# Active Network configuration:
#-----------------------------------
# IP mode
# Hostname
#-----------------------------------
# New Network configuration:
#-----------------------------------
# IP mode STATIC
# IP address 1.1.1.1
# Netmask 1.1.1.1
# Gateway 1.1.1.1
# Domain search list ibm.com xiv.ibm.com
# Domain name servers 1.1.1.1 2.2.2.2
# Hostname
Configure DNS addressess
1) Add domain to search list
2) Remove domain from search list
3) Add name server
4) Remove name server
5) Clear name servers and domain search list
6) Continue to hostname configuration
Setup/Network/DNS>
```

Figure 1-6 Network summary panel example

c. Select `6` when done. Enter the host name and then approve the configuration. The Domain Name System (DNS) parameters are displayed on the panel.

d. Select `Y` to approve and proceed to the next step (Step 4). Selecting `N` goes back to the beginning of Step 3 of the installation.

4. The next step, which is shown in Figure 1-7, is to set up Network Time Protocol (NTP) and regional settings.

```
Step 4 -NTP & Regional Configuration
------------------------------------
#-----------------------------------
# Active NTP/Regional configuration:
# ----------------------------------
# NTP status DISABLED
# Date 12/31/2012 23:59:59
# Timezone UTC
Configure NTP servers
1) Add NTP server
2) Remove NTP server
3) Return to previous menu
Setup/Regional Settings/NTP>
```

Figure 1-7 Hyper-Scale Manager server enable NTP server

Select whether to enable NTP and then the time zone, date, and time:

a. Select `1` to enable NTP and then add NTP servers. Then, select `2` to set the time zone. First, select a global region, as shown in Figure 1-8.

```
Please identify a location so that time zone rules can be set correctly.
Please select a continent or ocean.
1) Africa
2) Americas
3) Antarctica
4) Arctic Ocean
5) Asia
6) Atlantic Ocean
7) Australia
8) Europe
9) Indian Ocean
10) Pacific Ocean
11) none - I want to specify the time zone using the Posix TZ format.
```

Figure 1-8 Hyper-Scale Manager server Continent Select panel

b. Identify a location so that time zone rules can be set correctly. After a selection is made and a country is selected, the NTP server is enabled

5. Proceed with the SMTP configuration, as shown in Figure 1-9 (this step is optional).

```
#-----------------------------------
# New SMTP configuration:
#-----------------------------------
# SMTP status ACTIVE
# SMTP server 1.1.1.1
# SMTP port 25
# Sender sender@myDomain.com
# Destinations destination@myDomain.com
Configure SMTP destinations
1) Add destination
2) Remove destination
3) Remove all
4) Send confirmation mail
5) Continue
6) Return to previous menu
Setup/SMTP/Destinations
```

Figure 1-9 Hyper-Scale Manager SMTP enable panel

SMTP configuration enables sending emails to the server administrator in a service failure:

a. Type in the SMTP configuration and follow the possible choices that are shown in Figure 1-9.

b. When all of the necessary configuration information is provided, select `5` and continue to the final portion of the installation.

6. At this stage, the installation is complete and confirmed as successful (Figure 1-10).

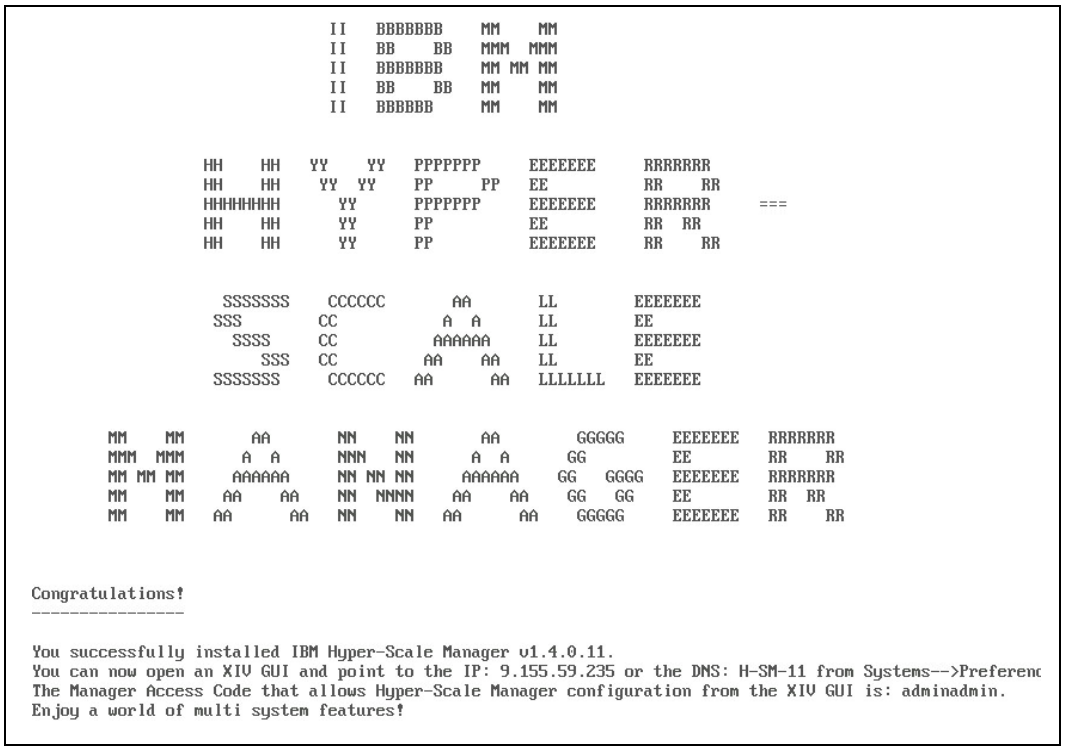

Figure 1-10 Hyper-Scale Manager installation successful

7. Reboot the server to activate the configured Hyper-Scale Manager installation.
8. Next, initialize the Hyper-Scale Manager inventory, which is described in 1.3.3, “Using the IBM Hyper-Scale Manager features” on page 24.

1.2.2 Installing Hyper-Scale Manager server for Red Hat Enterprise Linux

The Hyper-Scale Manager server is available in release 1.4.x and later. It is also available as a stand-alone installation on Red Hat Enterprise Linux (RHEL). It is downloadable as a bin file from IBM Fix Central:

http://www.ibm.com/support/fixcentral

Requirements: To install Hyper-Scale Manager server on RHEL, you need the following system elements:

- 64-bit RHEL 6 or later.
- At least 1 GB of disk space in the `/tmp` directory.
- 76 GB free space under the given installation path (the path is configurable).
- Installation must be done with the root user.

To install the server, follow these steps:

1. Transfer the bin file to the RHEL server, for example with the **`scp`** command from Linux, or **`winscp`** from windows.
2. Start the installation with the **`./IBM-Hyper-Scale-Manager_1.5.0-56.bin`** command. It takes you into the guided installation wizard. You need to make the bin file executable, for example with the **`chmod 777 IBM-Hyper-Scale-Manager_1.5.0-56.bin`** command.

3. After license agreement acceptance, the installation should look similar to the display shown in Figure 1-11.

```
Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, CLICKING ON
AN "ACCEPT" BUTTON, OR OTHERWISE USING THE PROGRAM,
LICENSEE AGREES TO THE TERMS OF THIS AGREEMENT. IF YOU ARE
ACCEPTING THESE TERMS ON BEHALF OF LICENSEE, YOU REPRESENT
AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND LICENSEE
TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

* DO NOT DOWNLOAD, INSTALL, COPY, ACCESS, CLICK ON AN
"ACCEPT" BUTTON, OR USE THE PROGRAM; AND


Press Enter to continue viewing the license agreement, or
enter "1" to accept the agreement, "2" to decline it, "3"
to print it, "4" to read non-IBM terms, or "99" to go back
to the previous screen.
1
Installing IBM Hyper-Scale Manager RPM package
########################################## [100%]
########################################## [100%]
Please wait ...
```

Figure 1-11 Hyper-Scale Manager RHEL installation

4. Next, initialize the Hyper-Scale Manager inventory, which is described in 1.3.3, “Using the IBM Hyper-Scale Manager features” on page 24.

1.2.3 Installing another IBM Hyper-Scale Manager server

The number of Hyper-Scale Manager servers that you can install is limitless. In practice, the reason for installing multiple servers is to ensure more reliability. If one server hosting the Hyper-Scale Manager stops running for any reason, you can rely on another server to still be able to manage your XIV Systems.

Each of these Hyper-Scale Manager servers can monitor up to 144 XIV systems in your organization. There is no communication between the Hyper-Scale Manager servers themselves. However, two Hyper-Scale Manager servers can have the same inventory, so that you can easily switch between them, therefore providing redundancy.

1.2.4 Upgrading the IBM Hyper-Scale Manager

The existing XIV Storage Systems profile file, which is a list of defined XIV Storage Systems and groups, is on the Hyper-Scale Manager server. This information is backed up as part of the Hyper-Scale Manager upgrade process. You can also create another profile using an export file to save this information as a backup, or to import it into another Hyper-Scale Manager server.

The local XIV Storage Systems profile file can be saved by exporting it from the XIV Storage Management GUI, as seen in Figure 1-12. It is saved as an Extensible Markup Language (XML) file and is then available for importing into the newly installed Hyper-Scale Manager GUI. This action overwrites the existing systems profile. The GUI is running on a workstation, and exports data from the Hyper-Scale Manager server that it is connected to.

Tip: The saved XIV Storage Systems profile can be imported on additional systems on which the XIV Storage Management GUI is installed, to ensure identical system views across all installations.

Save the local XIV Storage Systems profile file by exporting it from the XIV Storage Management GUI, as shown in Figure 1-12.

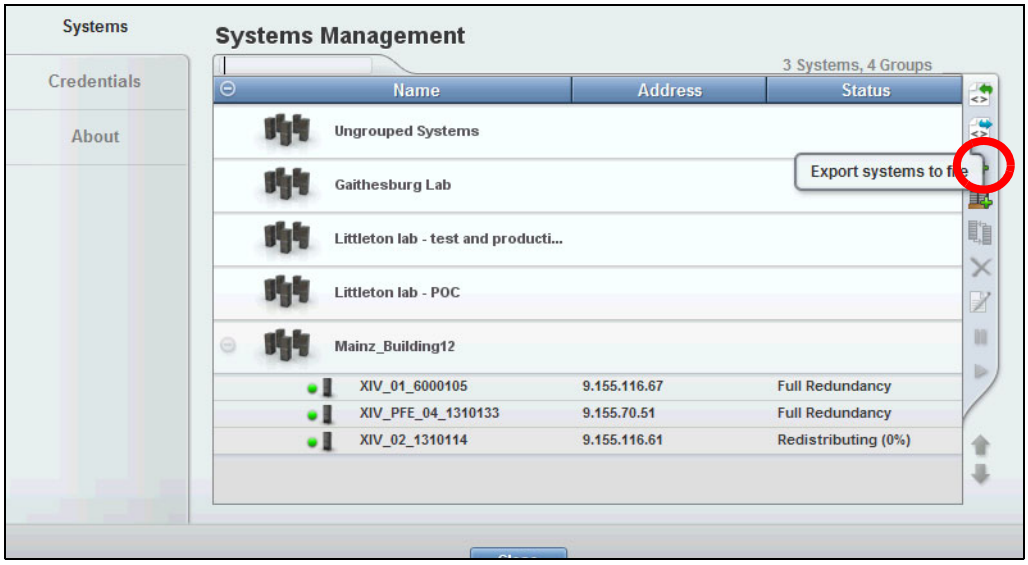

Figure 1-12 Export system file

1.3 XIV Storage Management software usage

Now that the XIV Storage Management Server software has been installed, the XIV Storage Management GUI can be used for further configuration and storage management.

Beginning with the XIV Storage Management GUI, this section describes the following topics:

- Signing into the XIV Storage Management GUI
- Connecting to XIV Storage Systems
- Getting an overview of the management and system views

1.3.1 Using XIV Storage Management GUI in Hyper-Scale Manager mode

The XIV Storage Management GUI is a multi-function tool that can be even more effective when used to access the XIV Storage Hyper-Scale Manager server. By default, the option to access the Hyper-Scale Manager server is not available, and must be configured using the systems preferences menu.

The following scenario assumes that the XIV Storage Management GUI has been installed, and the login was successful. The IP addresses of up to two Hyper-Scale Manager servers can be specified in the systems preferences. This enables the option to log in to either server for access to the XIV Storage Systems defined to those Hyper-Scale Manager servers.

1.3.2 Configuring the IBM Hyper-Scale Manager server

This section explains and illustrates how to use the Hyper-Scale Manager. First, create an inventory of the XIV systems to be managed with Hyper-Scale Manager.

Initializing the inventory

If the direct mode version of the GUI already has a set of XIV systems defined, it is possible to export the systems list to a file in your local file system (to be imported later into the Hyper-Scale Manager server). To do so, follow these steps:

1. Select **File Export Systems File**.
2. Install the new XIV GUI, and configure it to work with the Hyper-Scale Manager.
3. Start the XIV GUI. If it is not configured to use with Hyper-Scale Manager server, you get the default login dialog shown in Figure 1-13.

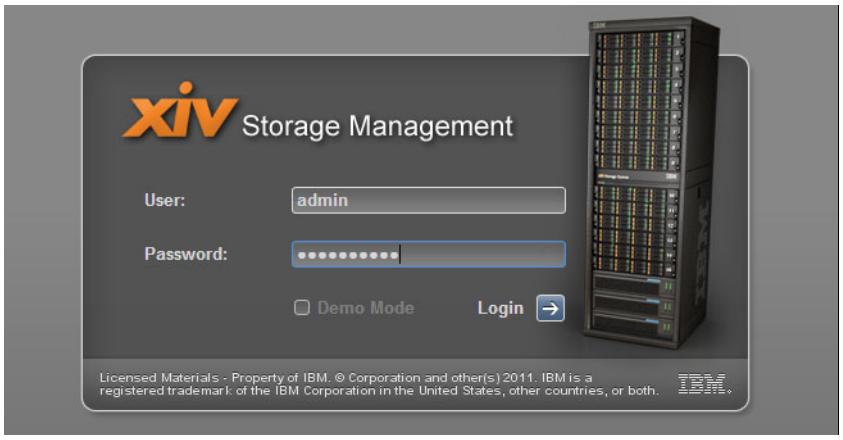

Figure 1-13 Default GUI login dialog

4. To connect the GUI to the Hyper-Scale Manager server, open **Systems** → **Preferences** from the main XIV GUI, as shown in Figure 1-14.

Figure 1-14 Set systems preferences

5. In the dialog that is shown in Figure 1-15, complete the primary server IP address and DNS, and click **OK**.

Figure 1-15 Hyper-Scale Manager server configuration in the XIV GUI

Important: Make sure that the XIV Systems file has been exported before you click **OK** because the GUI will try to log in to the Hyper-Scale Manager server configured and if it succeeds, the local GUI System file will be over written. To get XIV Systems file on to the Hyper-Scale Manager server, you must import the XIV Systems file.

6. There is now an option to log in to the Hyper-Scale Manager server. Log in using `Server Admin` credentials that were configured in the installation process of the Hyper-Scale Manager installation wizard.
7. When the connection to the Hyper-Scale Manager server has been configured, and you click **OK**, the GUI will automatically try to log in to the specified Hyper-Scale Manager server. A certificate warning dialog might appear the first time.

The next time that the XIV GUI is started, the XIV GUI login dialog includes a Mode option, as shown in Figure 1-16. The Mode field enables you to select between **Server** mode, **Direct** mode, and **Demo** mode.

Figure 1-16 Login window with default access

8. Configure a System Machine Account (this is a system user, predefined as *xiv_msms*, that will be used to monitor this IBM XIV Storage System from the Hyper-Scale Manager server). Go to **Systems** → **Manager Configuration** and enter the Hyper-Scale Manager server access code, which can be defined on the Hyper-Scale Manager server, as depicted in Figure 1-17.

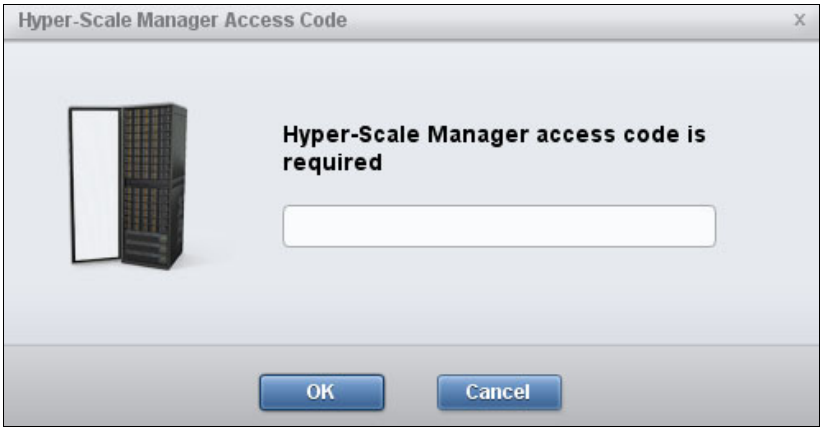

Figure 1-17 Access code

9. Go to **Credentials**, complete the **System Machine Account** password, and click **Update**, as shown in Figure 1-18 on page 15. The access code for the MSM can also be updated on this menu.

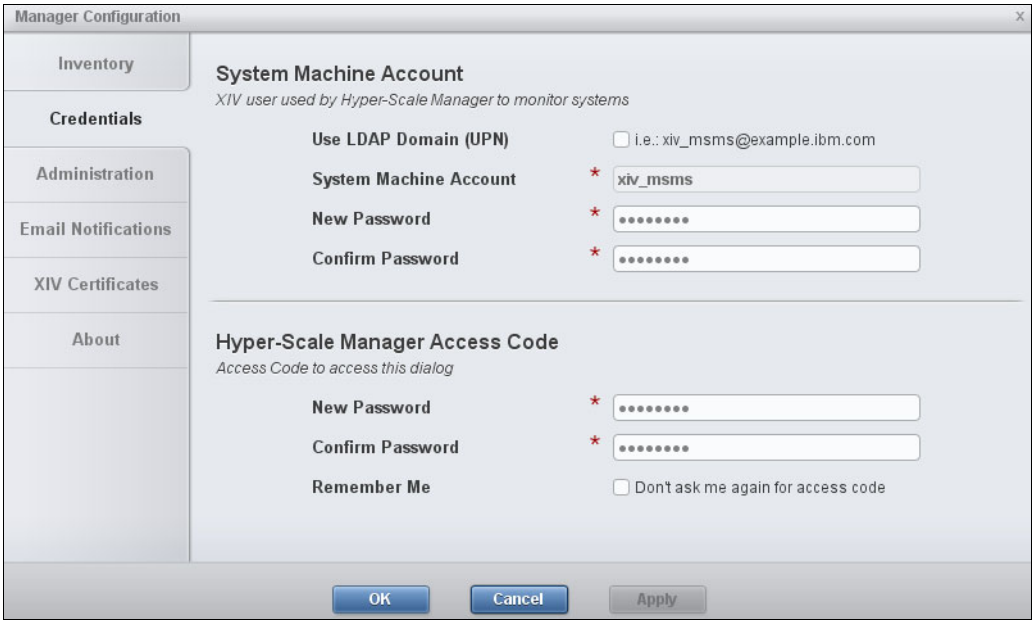

Figure 1-18 System Machine Account window

Restriction: Only a server admin can access this GUI window with the Hyper-Scale Manager server access code.

About the System Machine Account

This System Machine Account (xiv_msms user) authenticates into all XIV systems in order to poll configuration data only:

- This user does not change the configuration
- This user's name is hardcoded: xiv_msms
- This user can be defined in LDAP (make sure that it is added to all XIV storage admin groups in the LDAP)
- This user must have a storage administrator role (similarly to the admin user)
- This user must be defined with the same password on all XIV systems in the IBM Hyper-Scale Manager inventory
- This user must be defined in the IBM Hyper-Scale Manager (through the GUI or CLI)

Managing the system inventory

A list of systems can be imported from a file. This is a way to easily keep multiple Hyper-Scale Manager servers in sync with the same configuration. Systems can also be added individually and can be included in a group.

Monitoring of a selected XIV system can be disabled or enabled as needed, or the system can be removed from the inventory. After a system is added, the IP addresses for the selected system can be modified if needed.

Importing systems from a file

To import systems, follow these steps:

1. Open **Systems** → **Manager Configuration** → **Inventory** and click **Import systems from file**, as shown in Figure 1-19.
2. If a previously exported system XML file is available, browse to the exported XML file and import it.

Figure 1-19 System Import file

3. When the systems were added, if they get an authentication error status, select the non-authenticated systems, right-click, and select **Diagnose/Fix Authentication Error** from the menu.

Adding a system

To add an XIV system to the Hyper-Scale Manager server inventory, follow these steps:

1. Open **Systems** → **Manager Configuration** → **Inventory** and click the **Add System** icon, as shown in Figure 1-20.
2. Enter the XIV system IP addresses and select a group.

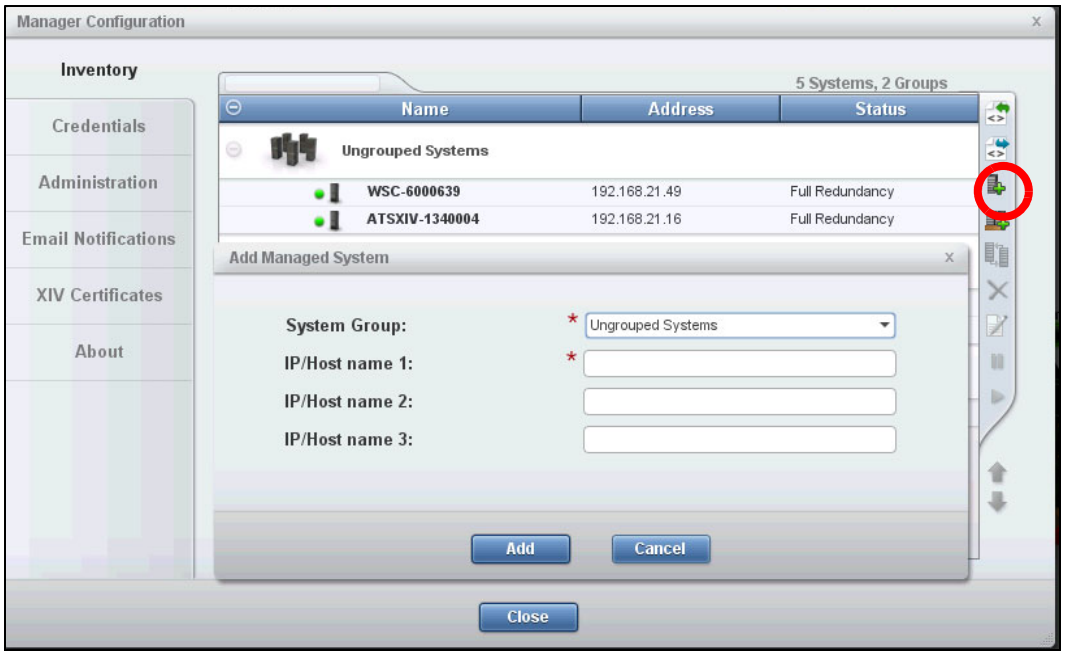

Figure 1-20 Add single system to the Hyper-Scale Manager GUI

Adding a system group

To define a system group, follow these steps:

1. Open **Systems** → **Manager Configuration** → **Inventory** and click the **Add Group** icon.
2. As seen in Figure 1-21, there will be a prompt to enter the new system group name.

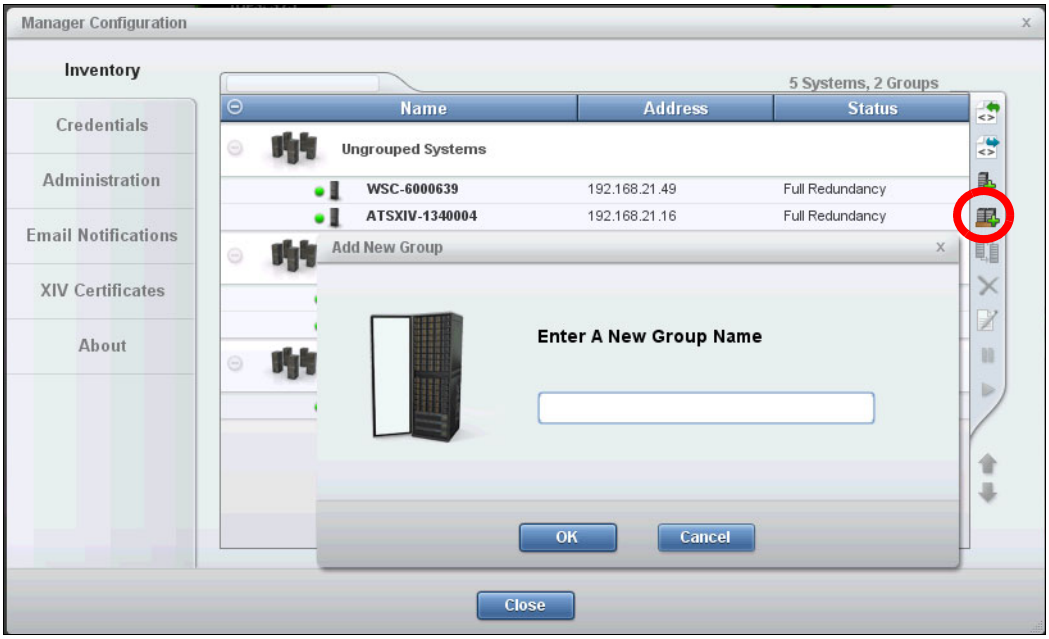

Figure 1-21 Adding a group to the Hyper-Scale Manager GUI

Removing a system from the inventory

To remove an XIV system from the inventory, follow these steps:

1. From the Systems tab in the Multi-System Manager Admin Options window, select the system from the list and click the **Delete** icon, as shown in Figure 1-22 on page 19.
2. There will be a prompt to confirm the delete action, after which the system is removed from inventory.

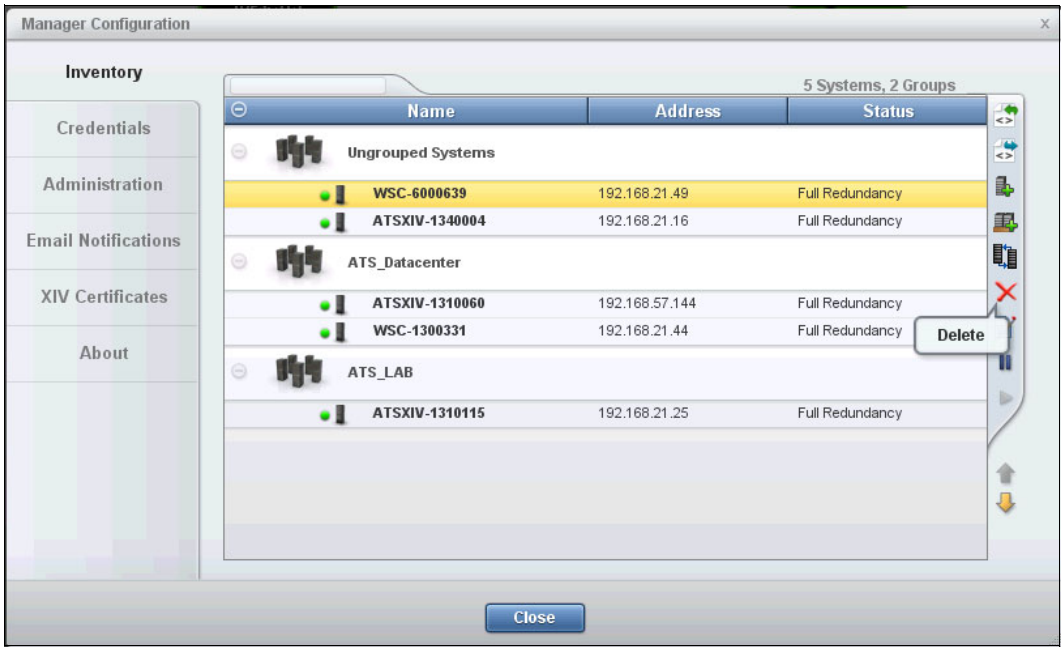

Figure 1-22 Deleting a system from inventory

Modifying a system in the inventory

To modify an XIV system from the inventory, follow these steps:

1. From the **Systems** → **Manager Configuration** → **Inventory** window, select the system from the list, right-click, and select the action that is required, as shown in Figure 1-23.
2. The **Edit** option that is available on the right side of the window enables updating the IP addresses for the system selected.

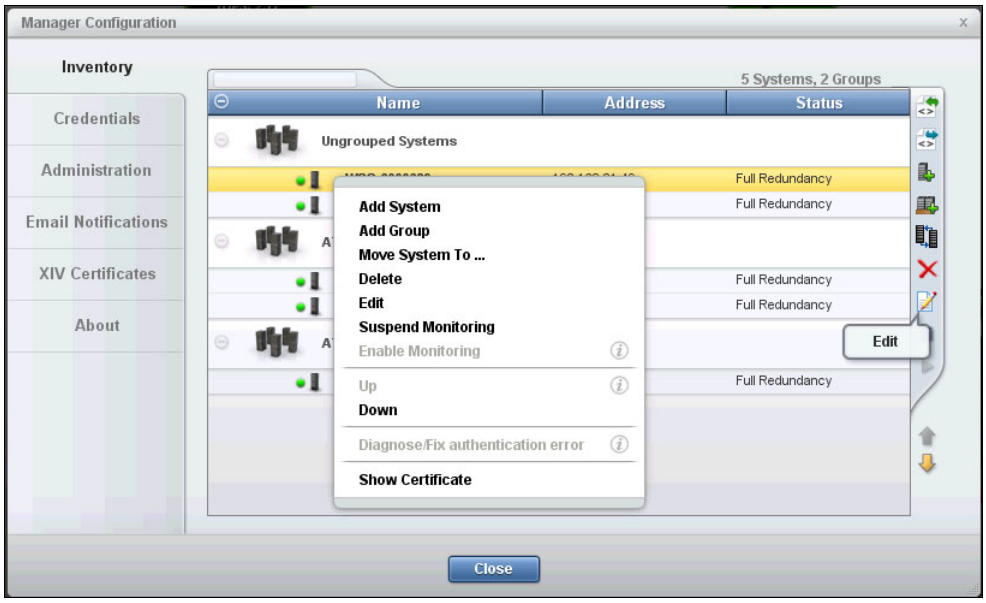

Figure 1-23 Modifying systems to the Hyper-Scale Manager GUI

Suspending or enabling monitoring of an XIV system

To suspend and enable the Hyper-Scale Manager's connection to an XIV system from the inventory, follow these steps:

1. From the **Systems** → **Manager Configuration** → **Inventory** window, select the system from the list.
2. Click either **Suspend** or **Enable,** as shown in Figure 1-24 and Figure 1-25.

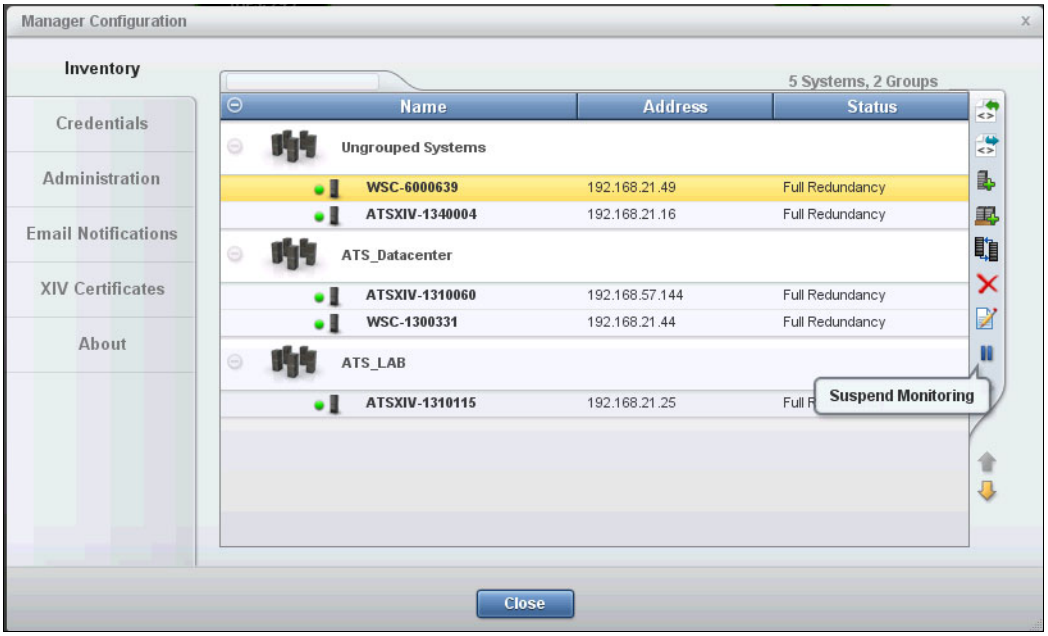

Figure 1-24 Suspend system monitoring to the Hyper-Scale Manager GUI

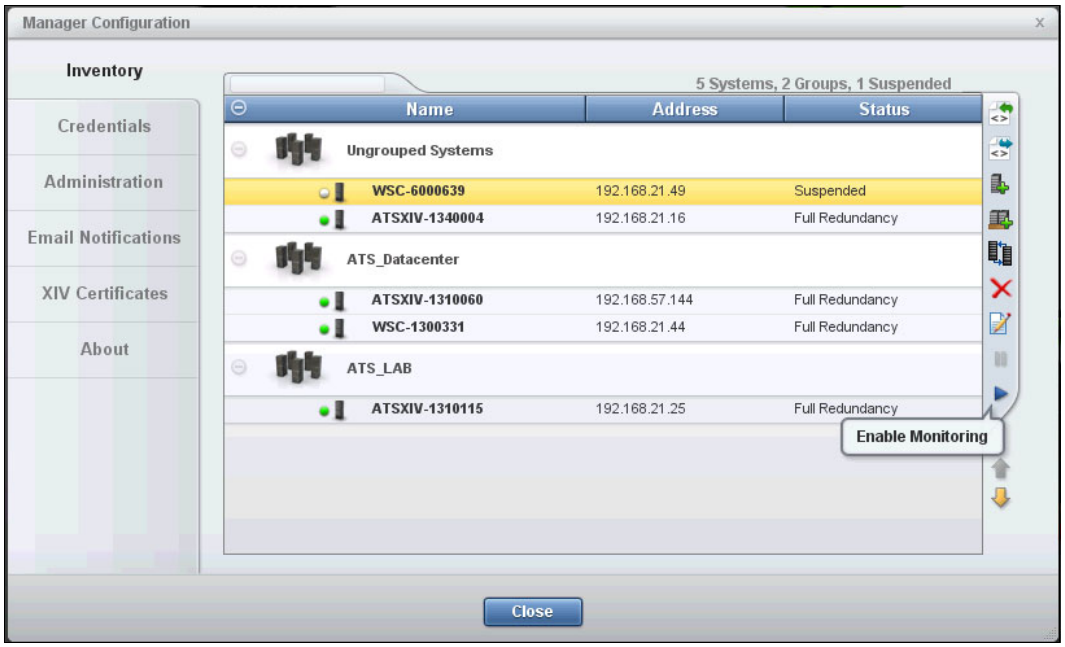

Figure 1-25 Enable system monitoring to the Hyper-Scale Manager GUI

Performing administration tasks

This Administration menu, which is shown in Figure 1-26, provides various administrative tasks:

Download Manager Logs	Download all types of Hyper-Scale Manager logs to your local workstation.
Download Manager Audit Log	Download a detailed audit log of Manager Configurations.
Discover new XIV systems to all users	Authenticate all logged in users on all XIV systems.
Hyper-Scale Manager certificate	Manage the authentication certificate.
Upgrade IBM Hyper-Scale Manager	Upgrade the code version.

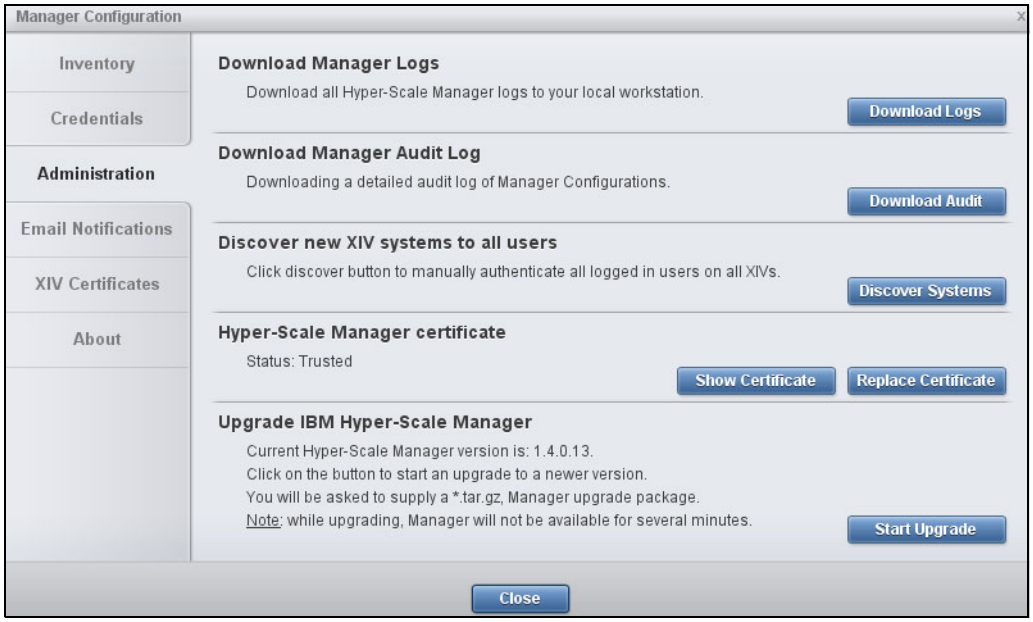

Figure 1-26 Administration tasks

Configuring email notifications

To configure email notifications for monitoring the XIV systems in the inventory, follow these steps:

1. From the **Systems** → **Manager Configuration** → **Email Notifications** window, enter the SMTP server information.

2. Use the icons to add, edit, or remove email addresses, as shown in Figure 1-27.

Figure 1-27 Email Notifications view

Administering XIV certificates

To administer the certificates used to access the XIV systems in the inventory, follow these steps:

1. Using the icons on the right, display the certificate properties, as shown in Figure 1-28.

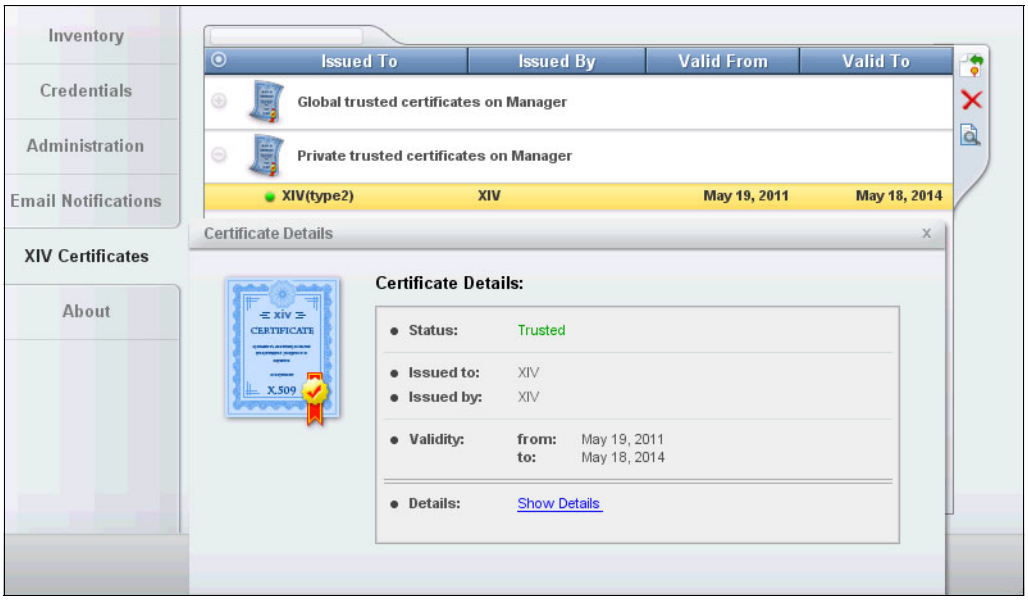

Figure 1-28 XIV certificate properties

2. Using the icons on the right, import or remove certificates, as shown in Figure 1-29.

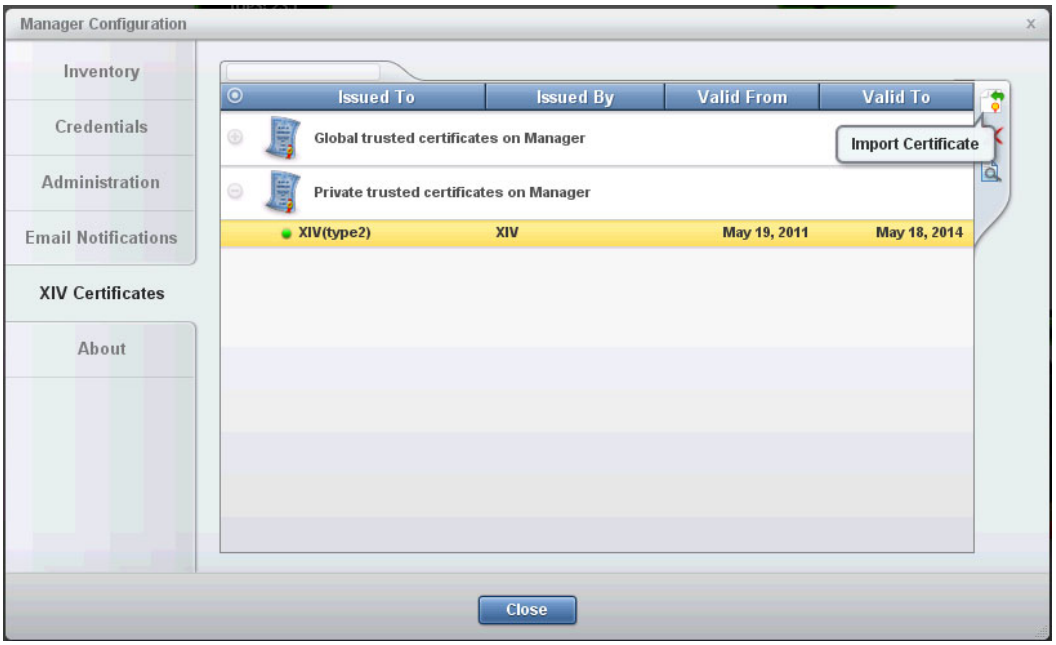

Figure 1-29 XIV Certificates view

1.3.3 Using the IBM Hyper-Scale Manager features

Management of the XIV Storage Systems in the inventory of the Hyper-Scale Manager begins at the main systems view in the XIV GUI. The Hyper-Scale Manager features are enabled when the logon option specifies one of the defined Hyper-Scale Manager servers.

Depending on the user permissions, the interface might differ slightly. There will be consolidated lists of all the XIV systems in the inventory, thereby providing a consolidated view of all of the available resources and events. The powerful search feature operates over the entire inventory, making it easy to find specific resources of interest.

The Hyper-Scale Manager GUI provides several options at the All Systems level:

- Connectivity view
- List view
- Tiles view

Perspective: The figures included in this chapter depict the interface layout from a storage administrator perspective. Not all options are available to the application administrator and read-only permission levels.

These views can be accessed in the upper left portion of the Hyper-Scale Manager GUI from the Connectivity menu, shown in Figure 1-30.

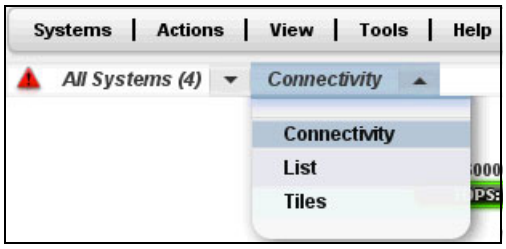

Figure 1-30 Hyper-Scale Manager menu options

Systems can also be viewed and selected with a summary tab that is at the upper left of the Hyper-Scale Manager GUI. This summary tab provides a list of XIV systems that are being managed and visible by Hyper-Scale Manager, in addition to the status and any alerts on the system itself. This view is shown in Figure 1-31. You can manage any particular system by selecting it from the list to then get the standard system view displayed.

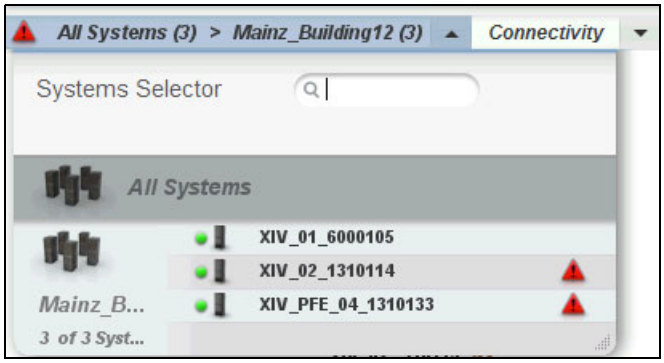

Figure 1-31 All Systems view

Connectivity View

The IBM Hyper-Scale Manager GUI offers the ability to view the systems connectivity. This view, which is shown in Figure 1-32, shows all systems that are connected for replication.

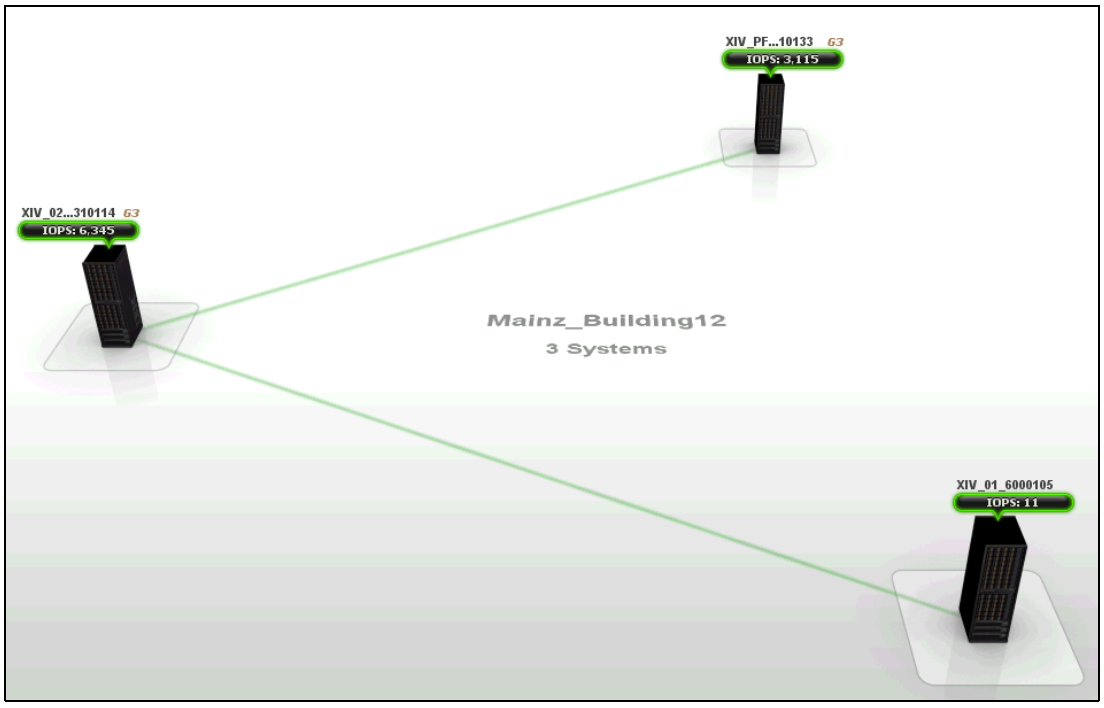

Figure 1-32 Connectivity view

List view

The Hyper-Scale Manager GUI can display all of the managed systems in a list view. The list view displays the overall capacity usage of the systems. For each system, it also displays critical system information, such as system name, system group, capacity available, capacity used, and overall input/output operations per second (IOPS), as shown in Figure 1-33 and Figure 1-34.

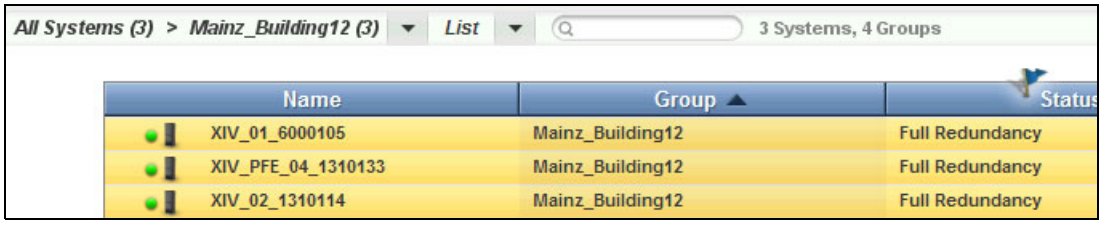

Figure 1-33 Hyper-Scale Manager List view

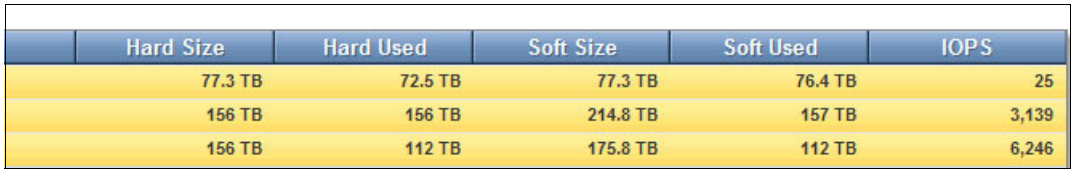

Figure 1-34 Hyper-Scale Manager List view: continuation

Additional summary columns can also be added to the list view, by right-clicking the column header to open the dialog panel that is shown in Figure 1-35.

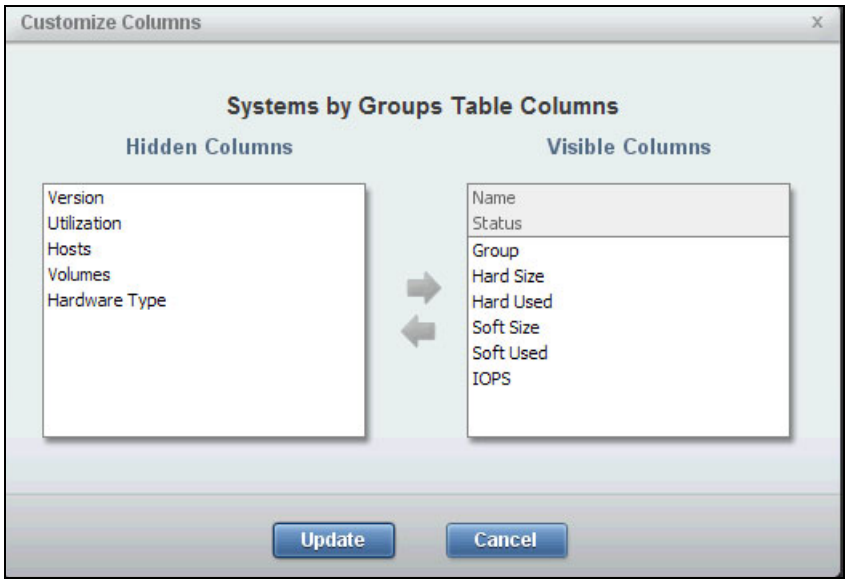

Figure 1-35 Options to add columns

Tiles view

The Hyper-Scale Manager GUI also provides a tiled view of the systems, as shown in Figure 1-36. Each tile shows IOPS and a redundancy summary of the corresponding system.

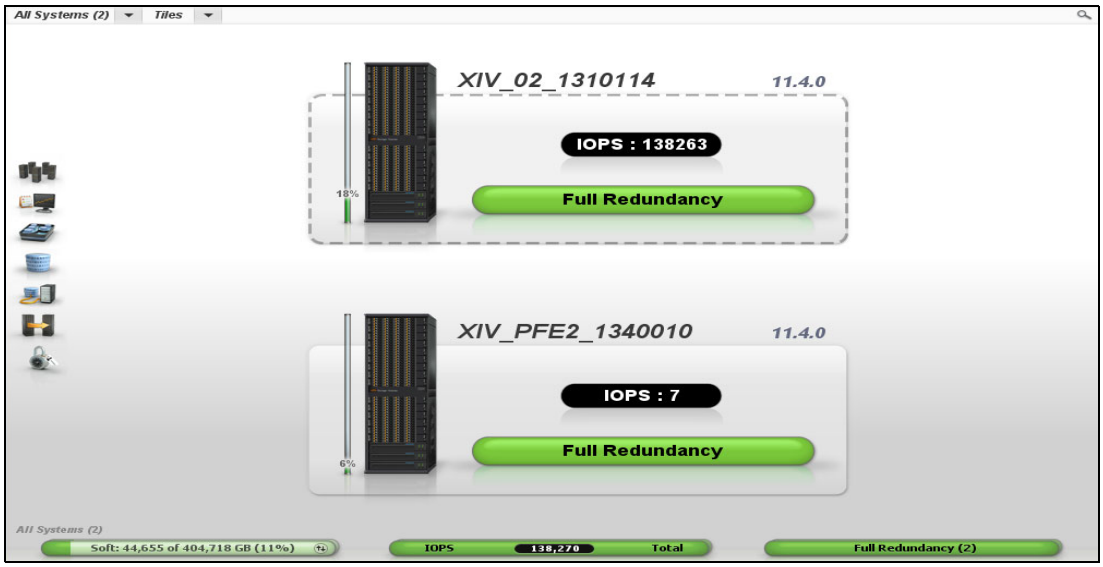

Figure 1-36 Hyper-Scale Manager Tiles view

1.3.4 Mass-configuring XIV systems

Mass configuration enables you to change the configuration on multiple XIV systems with a single click. You can copy system configuration from one system and paste it onto several XIV systems. Mass configuration can be run from the GUI in server mode and in direct mode. Mass configuration requires the appropriate access rights to all of the involved XIV systems.

Mass configuration is available for the following features:

- Lightweight Directory Access Protocol (LDAP) configuration
- Support parameters
- Pool alert thresholds
- Adding and editing a user
- Changing the password for a user
- Adding hosts to multiple XIV Systems
- Adding and editing a cluster

Mass copy-paste configuration

To perform a mass copy-paste operation, follow these steps:

1. From the XIV GUI, right-click the XIV system that you want to copy the configuration from, and then select **Copy System Configuration** from the menu, as shown in Figure 1-37.

   The menu closes.

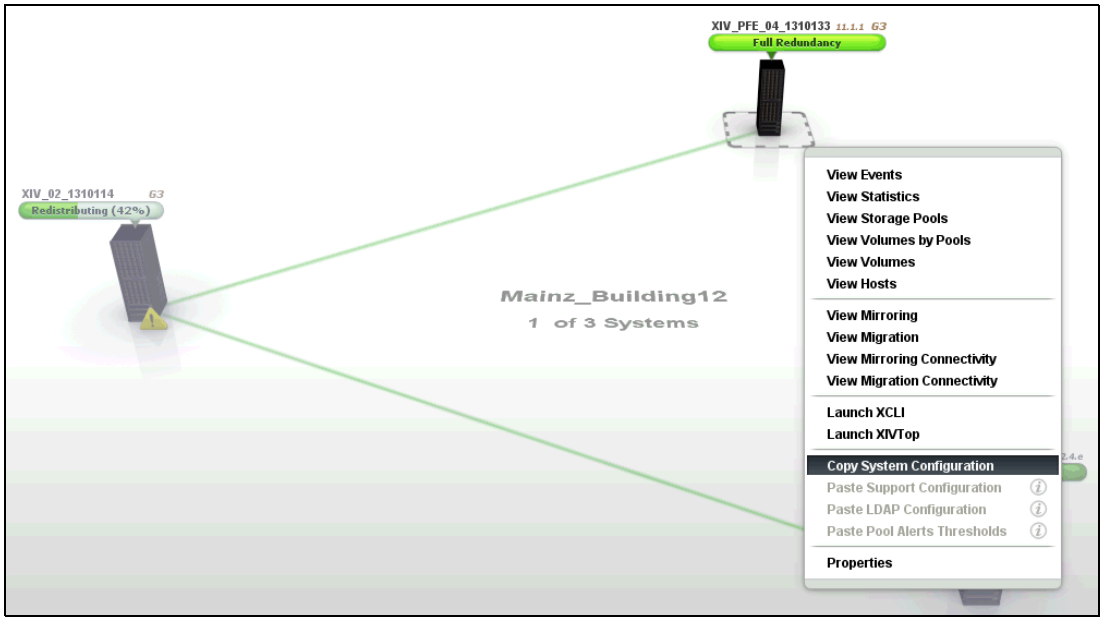

Figure 1-37 Select XIV system to copy from

2. Now, select one or several systems to which you want to paste the configuration:
   a. Right-click and select **Paste** → **Configuration** from the menu.
   b. In this example, Paste Support Configuration is selected.
   c. The Mass Pool Alerts Thresholds Configuration dialog opens, as shown in Figure 1-38.

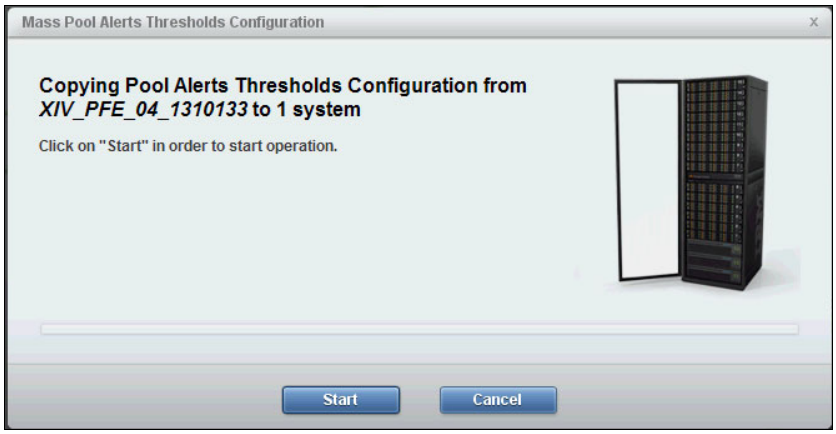

Figure 1-38 Mass support window

> **Tip:** Whenever a paste option is unavailable, hover over the option to display a tooltip that explains the reason. In this example, Paste LDAP Configuration is unavailable and the tooltip says that the target system's version is not compatible with the source system.

3. Click **Start**.

   A progress bar is displayed on the copy dialog panel, as shown in Figure 1-39. Clicking **Cancel** at this stage cancels the mass configuration. When the copy operation is done, a summary of the results is displayed. Clicking **Show Results** opens a detailed report window.

> **Important:** Mass configuration does not stop on error. In other words, it tries to configure all systems even though some might fail.

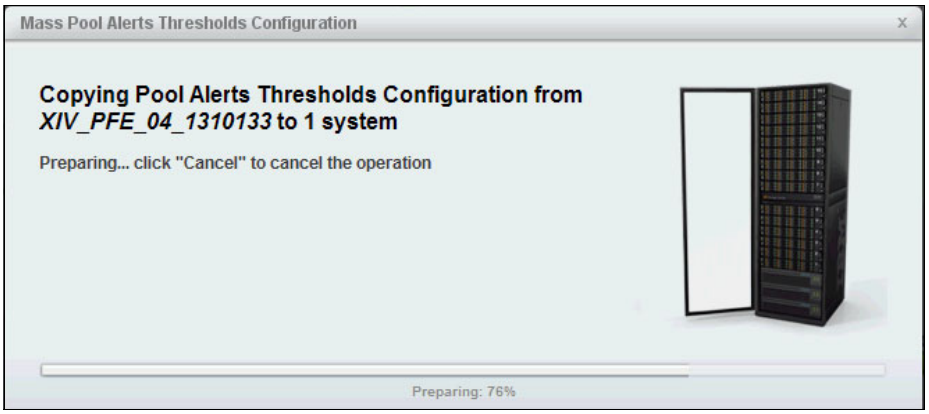

Figure 1-39 Progress bar

After this task, the configuration of one system was deployed on other systems.

Note: Closing the GUI cancels the paste operation and ends the mass configuration in an unknown state. In such a case, it is advisable to go over the systems and see what has already been properly configured, and what still needs to be configured.

Add a user on multiple XIV systems

The ability to run the same process on multiple XIV systems at the same time can be a convenient time saver. The Add User dialog enables the selection of multiple systems on which a new user will be created, as shown in Figure 1-40.

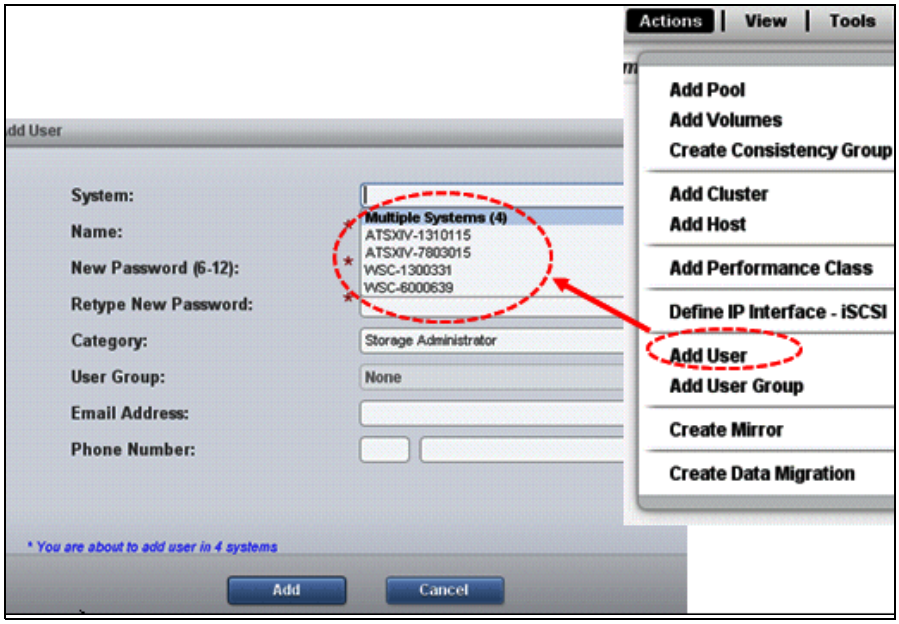

Figure 1-40 Common actions for multiple XIV systems

Adding a host to multiple XIV Systems

With the Hyper-Scale Manager, it is possible to add a host definition to several XIV systems at one time, as shown in Figure 1-41.

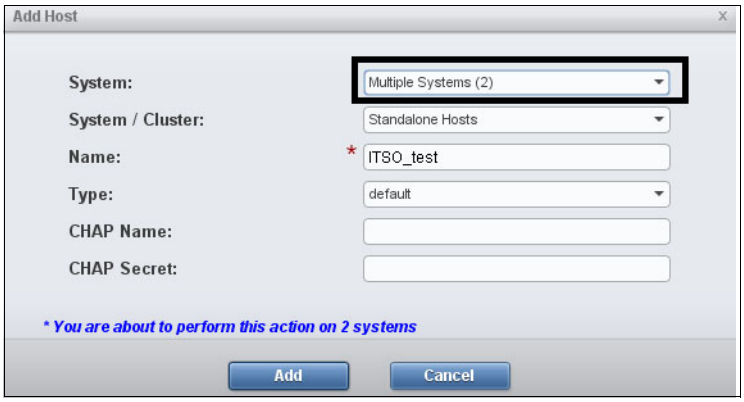

Figure 1-41 Add a host to multiple systems

After adding a host, it is possible to display a log of what has been created, as shown in Figure 1-42.

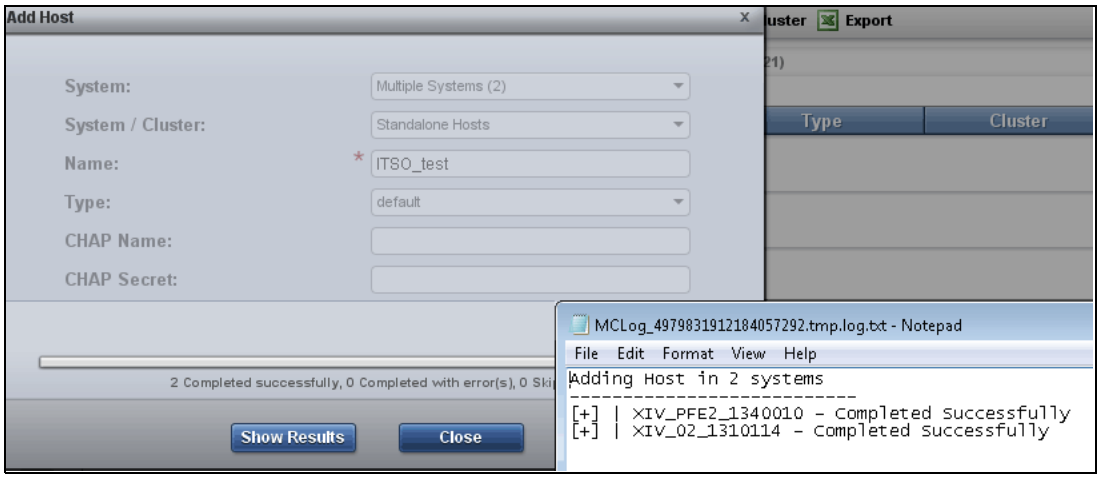

Figure 1-42 Log showing the result of adding a host to multiple systems

Adding a cluster to multiple XIV Systems

Adding a cluster definition is done in the same way, as shown in Figure 1-43.

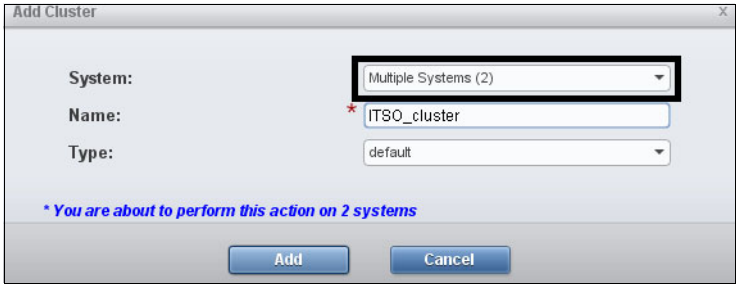

Figure 1-43 Adding a cluster definition to multiple XIV systems

1.4 General XIV Storage System operations

The configuring and monitoring functions available in the XIV Storage System GUI using Direct mode are also available when using the Hyper-Scale Manager Server mode. The Hyper-Scale Manager server, however, gives views of all of the defined XIV systems. It also provides an overall view of the current configurations.

1.5 XIV Capacity planning with Hyper-Scale Manager

The XIV Management Tools v4.4 provides an infographic-style reporting to ease capacity planning across multiple XIV systems. The intuitive report provides information about capacity growth, utilization, and trends. It empowers administrators to understand and respond quickly to storage needs. The report is only available in a Hyper Scale Manager configuration and assumes connectivity to the XIV systems where the capacity monitoring is planned for.

Note: Capacity Planning report is available to all XIV releases (generation 2 and Gen3 systems) through Hyper-Scale Manager v1.5 included in XIV management Tool v4.4. Statistical data is collected and kept by the Hyper-Scale Manager and requires at least 30 days of data to generate a useful report.

The IBM Hyper-Scale Manager collects usage statistics to calculate a forecast of the future use of XIV systems and pools. Statistical data is also exported to a CSV file in order to be used by common analytical tools. Figure 1-44 shows an architectural overview.

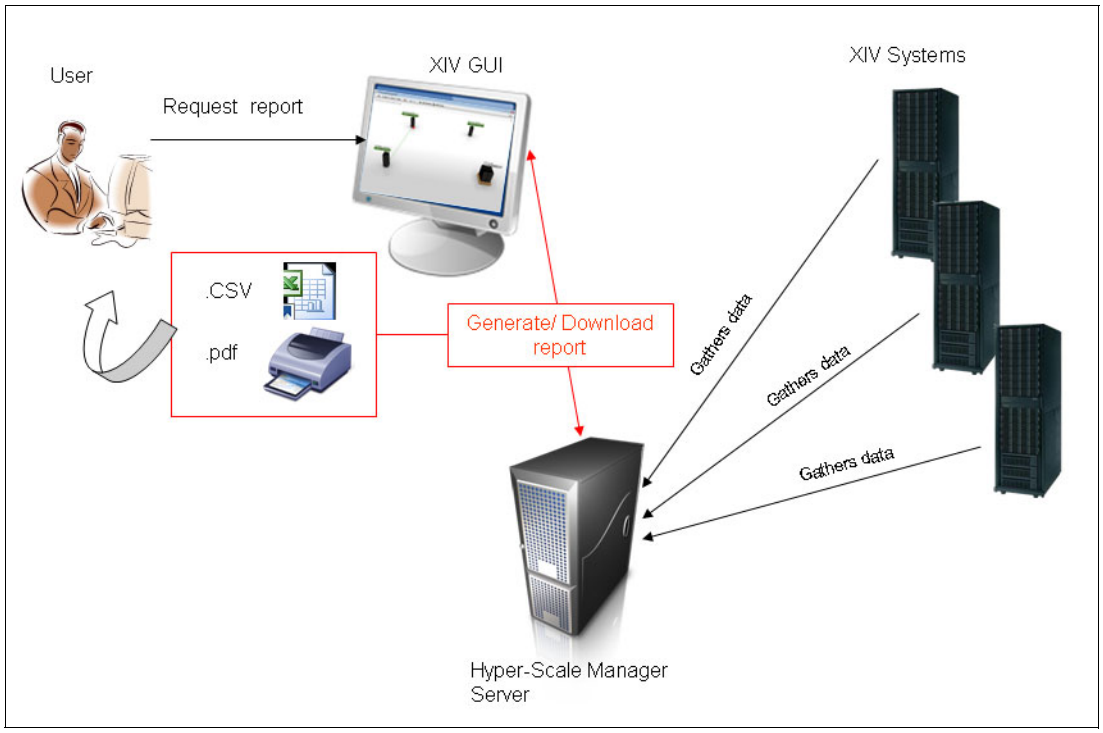

Figure 1-44 Hyper-Scale capacity reporting architecture

The capacity data for the systems included in the inventory is collected daily by the Hyper-Scale Manager server and stored in a database on the server. Since capacity data is available in the database, there is no interaction with the XIV systems when the administrator asks for a report to be generated.

The capacity data must meet several criteria for the forecast to be calculated. Capacity forecast is not calculated for the following reasons:

- User has insufficient access rights (role must be storage administrator or read-only) for all of the XIV systems included in the inventory.
- Insufficient number of samples. The forecast is not calculated if the number is less than 30.
- Utilization is too low. If the system or a storage pool utilization is less than 10%, no forecast is calculated.
- Trend cannot be calculated on pools that have no available space for volumes allocation.
- No trend. That is, capacity is fluctuating or flat, or space utilization is decreasing.

The IBM Hyper-Scale Manager collects capacity data for XIV systems that are listed on the inventory (see “Managing the system inventory” on page 16). Removing a system from the inventory implies stopping the data collection. However, to overcome situations in which the system was mistakenly removed from the inventory, or removed from the inventory for a short period of time, the IBM Hyper-Scale Manager applies the following rules on collecting capacity data for systems that are removed from the inventory:

- As long as the system is listed on the inventory, the IBM Hyper-Scale Manager collects and keeps its capacity data.
- Whenever the system is removed from the inventory, its capacity data is not immediately deleted. It is kept until the next time slot on which the data is collected from the machine.
- If the system is returned to the inventory before arriving to the next collection time slot, the capacity data and its continuity are kept.
- If the system is removed from the inventory, it is impossible to reset its capacity data. To reset the capacity data, the system must be listed on the inventory.
- If the user chooses to reset capacity data for all systems, even non-monitored systems capacity data will be reset.

The capacity report is generated from the XIV GUI. Instructions on how to generate the report are provided in the following sections.

1.5.1 Generating a capacity report

To generate a capacity report, you must be connected to the XIV GUI (at a level distributed with the XIV Management Tools v4.4 or later) and in Hyper-Scale Manager mode (that is, not in direct GUI mode). Refer to the illustration given in Figure 1-45. The Hyper-Scale server that is selected is used as a proxy to connect to XIV systems configured in its inventory.

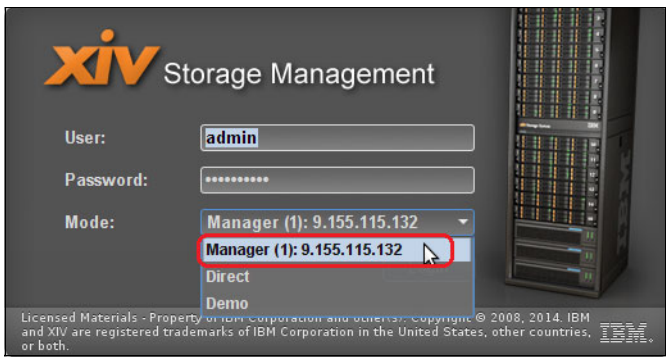

Figure 1-45 Capacity planning - login into Hyper Scale Manager

To invoke generating the capacity report, you have two options:

- You can select the Generate Capacity Report from the Tools drop-down menu, as shown in Figure 1-46.

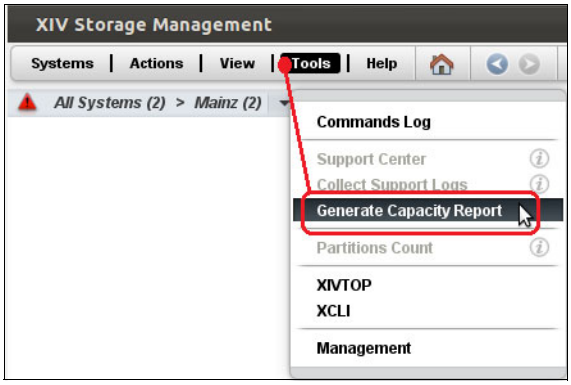

Figure 1-46 Capacity report - option 1

- Alternatively, you can right-click any system, in the All Systems view, and select Generate Capacity Report from the pop-up menu as shown in Figure 1-47.

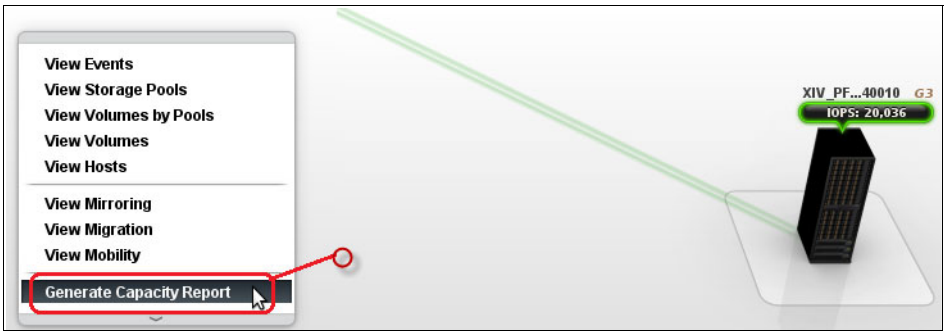

Figure 1-47 Capacity report - option 2

Using either of the above options, displays a dialog window, as shown in Figure 1-48, where you can specify a file name for the capacity report (the file type must be .zip), or just accept the predefined name. The predefined name of the capacity report file is: XIV_capacity_report_<# of>_systems_yyyy-mm-dd_hhmm.zip.

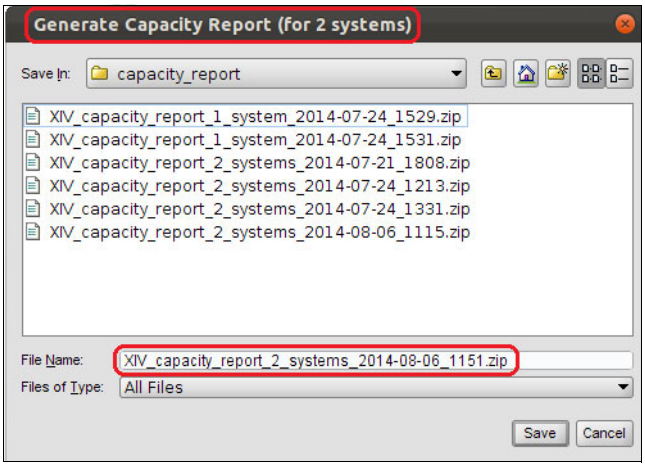

Figure 1-48 Capacity report - compressed report files location

Click **Save** to start generating the capacity report. When successful, a dialog like shown in Figure 1-49 is displayed. Click OK to close the dialog.

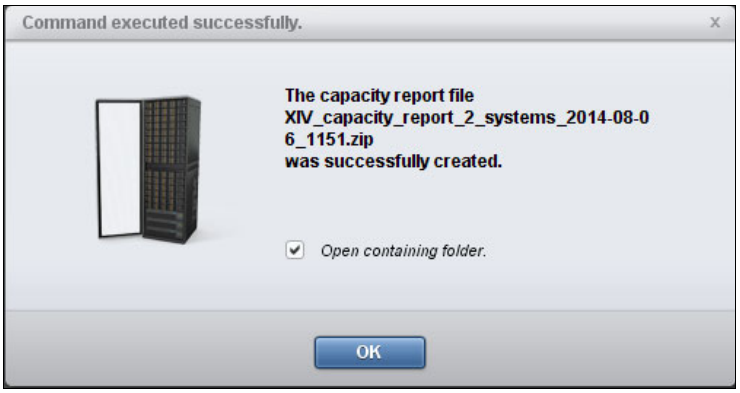

Figure 1-49 Capacity report - successful creation message

The generated .zip file contains a CSV and PDF file using a similar naming convention to the .zip file - XIV_capacity_report_<# of>_systems_yyyy-mm-dd_hhmm.<N>.zip -. Refer to Figure 1-50.

Figure 1-50 Capacity report - compressed report files

1.5.2 Capacity report structure

The PDF file contains a formatted report with summarized informational graphics and graphs for each monitored XIV system, domain, and pool.

The report, as illustrated in Figure 1-51, is divided into three parts:

- Part 1 - System Allocation
- Part 2 - Domain Usage and Allocation
- Part 3 - Pool Usage

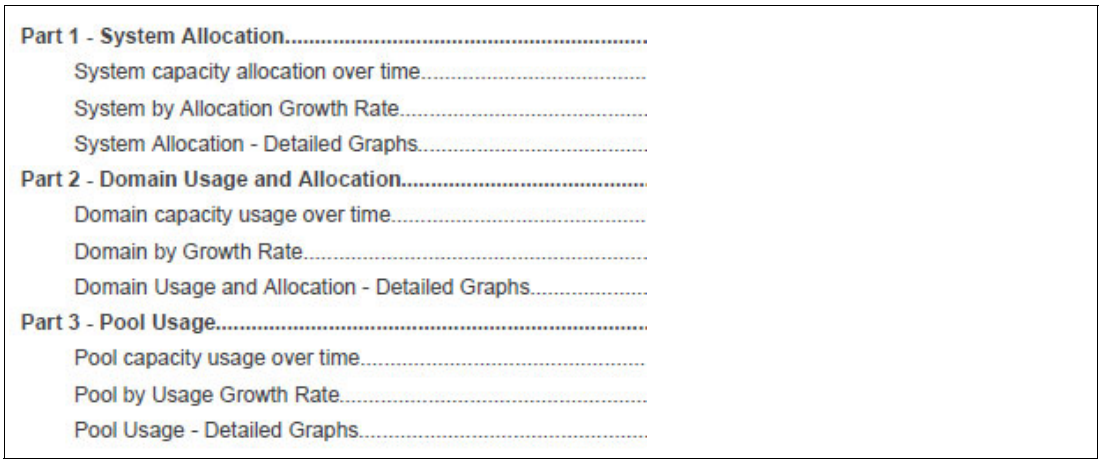

Figure 1-51 Capacity report structure

Part 1: System Allocation

This part shows the system capacity allocation over time.

A first diagram, as illustrated in Figure 1-52 shows an overall view for all of the XIV systems in the inventory and a capacity forecast covering the next 12 months.

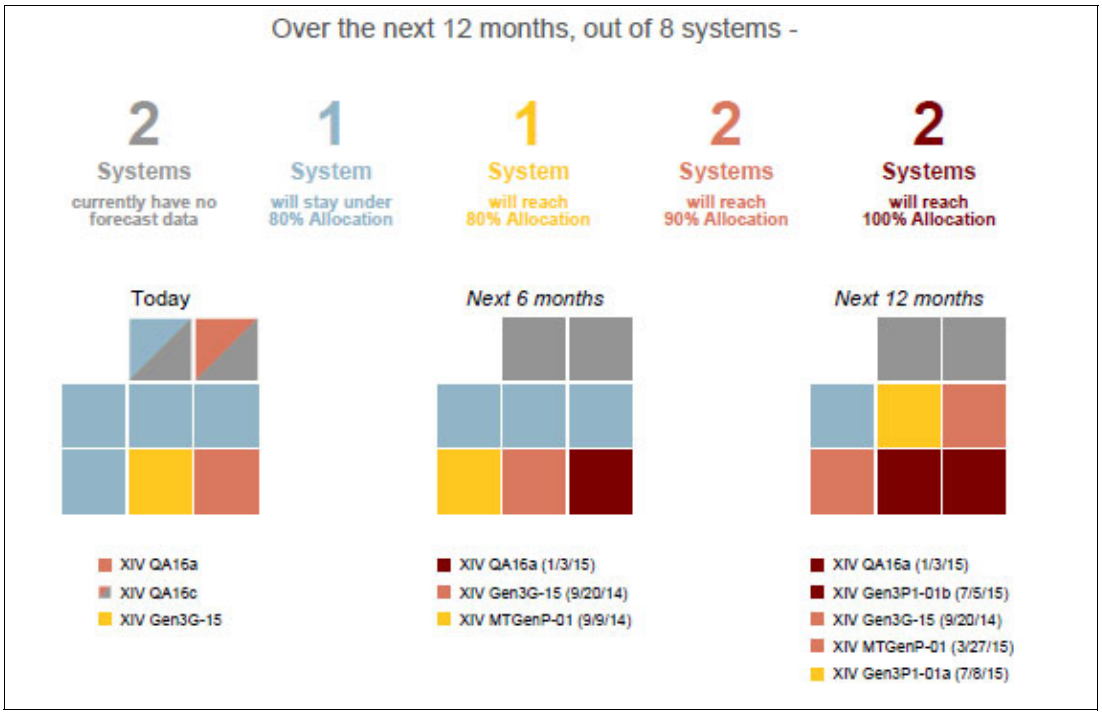

Figure 1-52 Capacity report - XIV system overview, capacity forecast for the next 12 months

The next graph, which is shown in Figure 1-53, ranks the system per capacity growth rate.

Figure 1-53 Capacity report - XIV system capacity graph for one system

Next, for each XIV Storage System, a graph is provided showing the progression of the system capacity allocation growth. Refer to Figure 1-54 (several systems are grouped) and Figure 1-55 showing the details of a single graph.

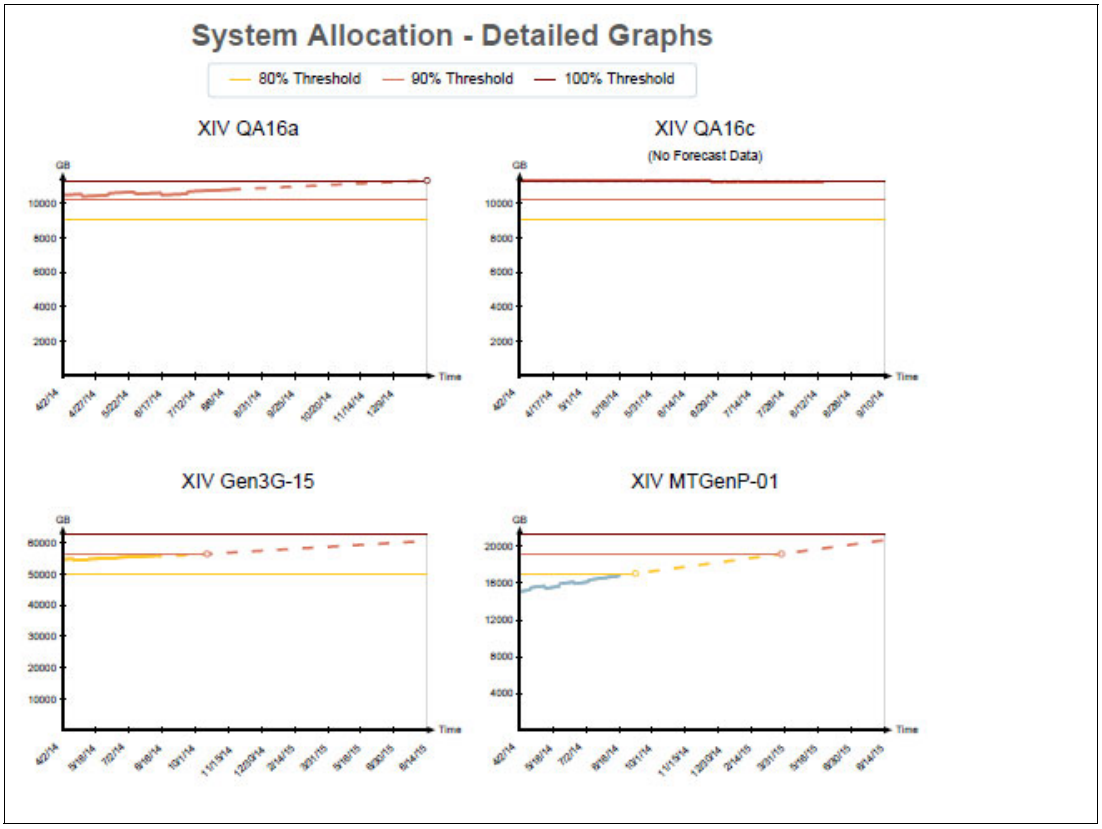

Figure 1-54 Capacity report - XIV system capacity graph overview

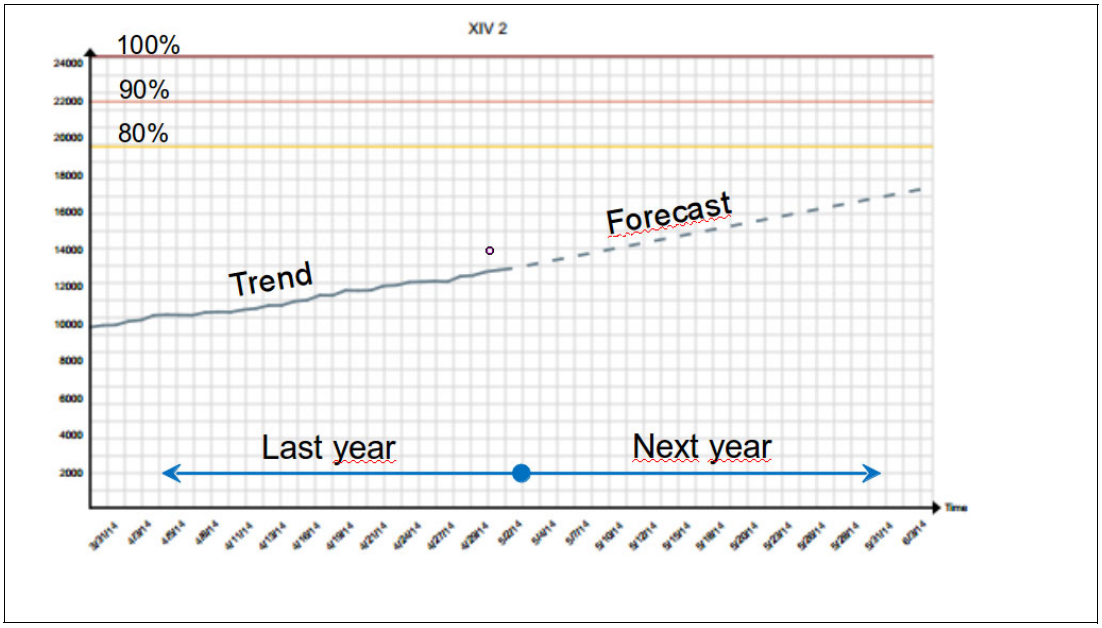

Figure 1-55 Capacity report - System allocation - detail

Part 2: Domain Usage and Allocation

This section provides a summary for each of the domains, according to the user domain association. A global administrator sees all of the system domains, regardless of the access policy. A domain administrator sees only the specific domains for which they are authorized.

The summary view is shown in Figure 1-56. The next diagram ranks the domains by growth rate. An illustration is given in Figure 1-57.

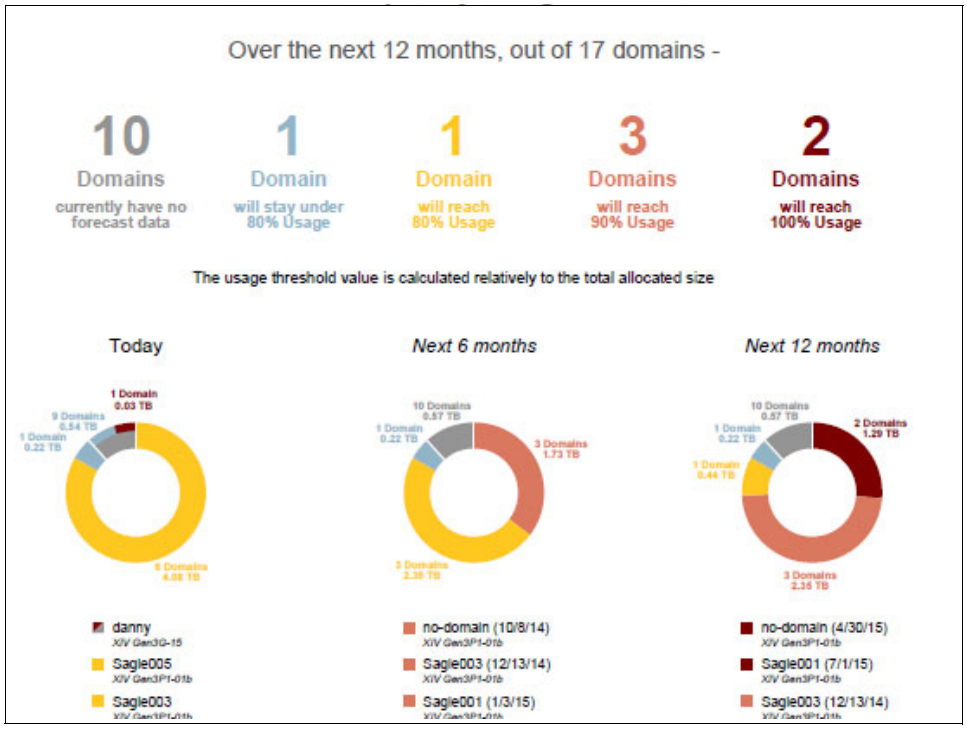

Figure 1-56 Domain capacity usage summary

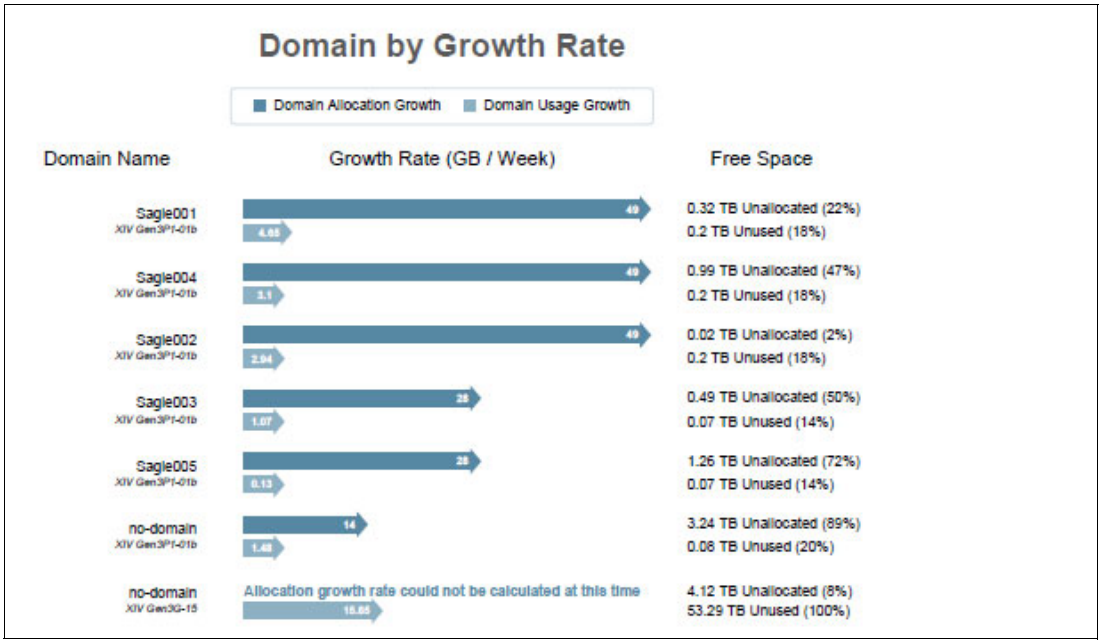

Figure 1-57 Capacity report - XIV domain growth rate in GB by week

The domain wise growth rate is very useful for data storage providers where the growth rate is shown and calculated as forecast for each domain.

In addition, for each domain a graph is provided showing the capacity growth progression for allocated and used capacity. Refer to Figure 1-58.

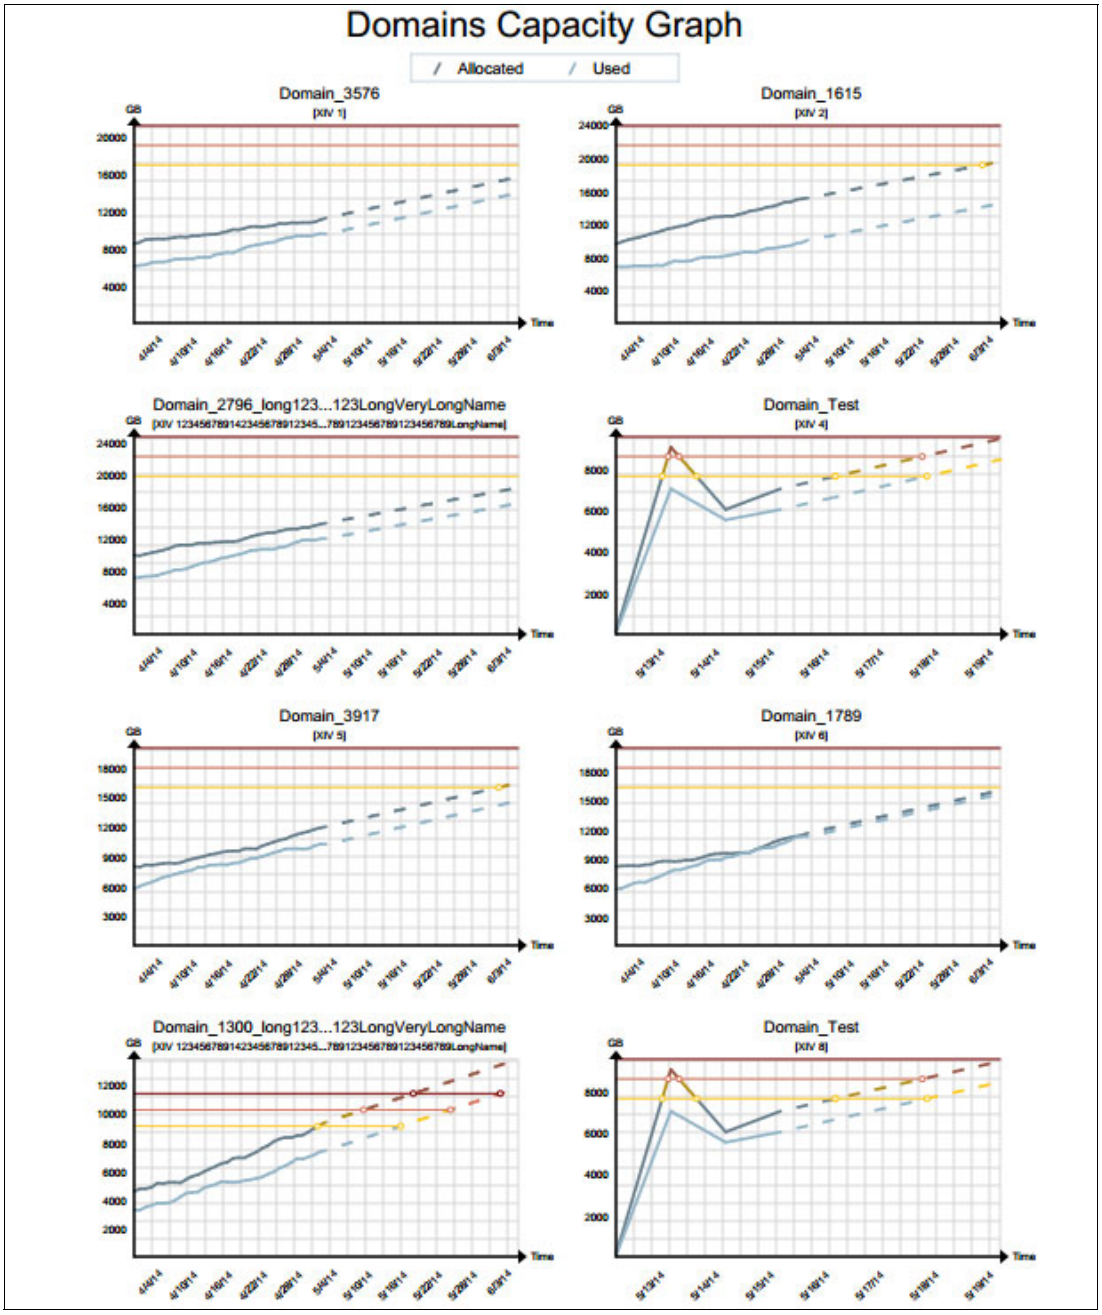

Figure 1-58 Capacity report - XIV domain capacity graphs

Part 3: Pool Usage

This section provides a summary for each of the storage pools capacity information collected for each of the monitored XIV systems. Refer to Figure 1-59.

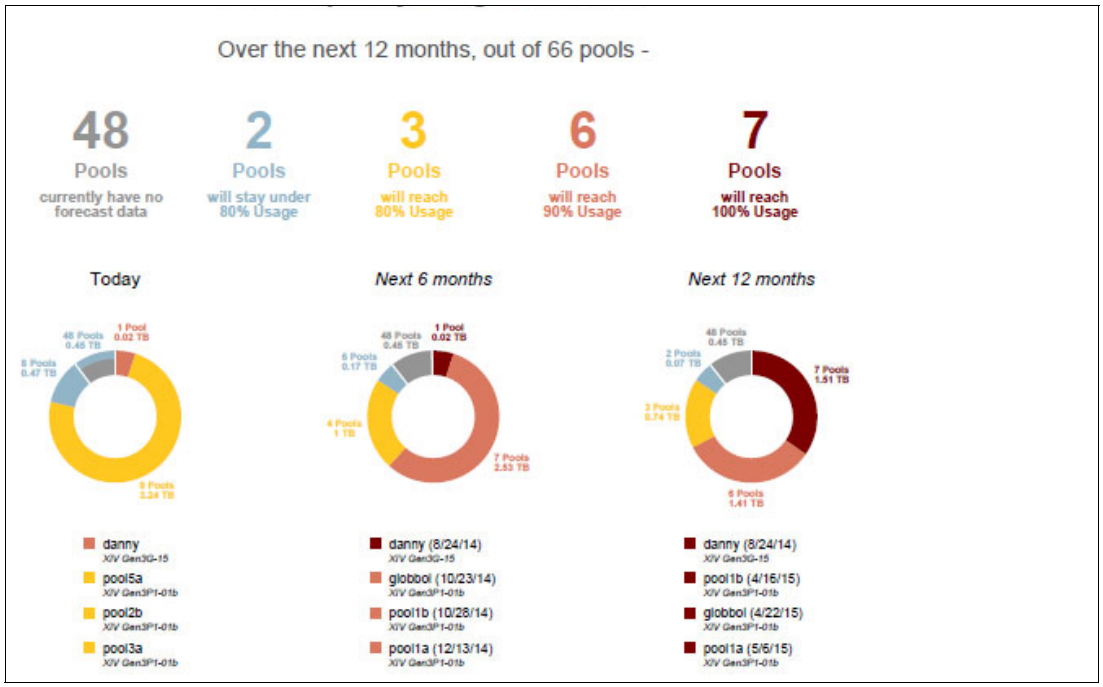

Figure 1-59 Pool capacity usage over time

This information provides further granularity to help you identify data growth areas in your XIV environment. Detailed graphs, per pool, are also provided. An illustration is shown in Figure 1-60.

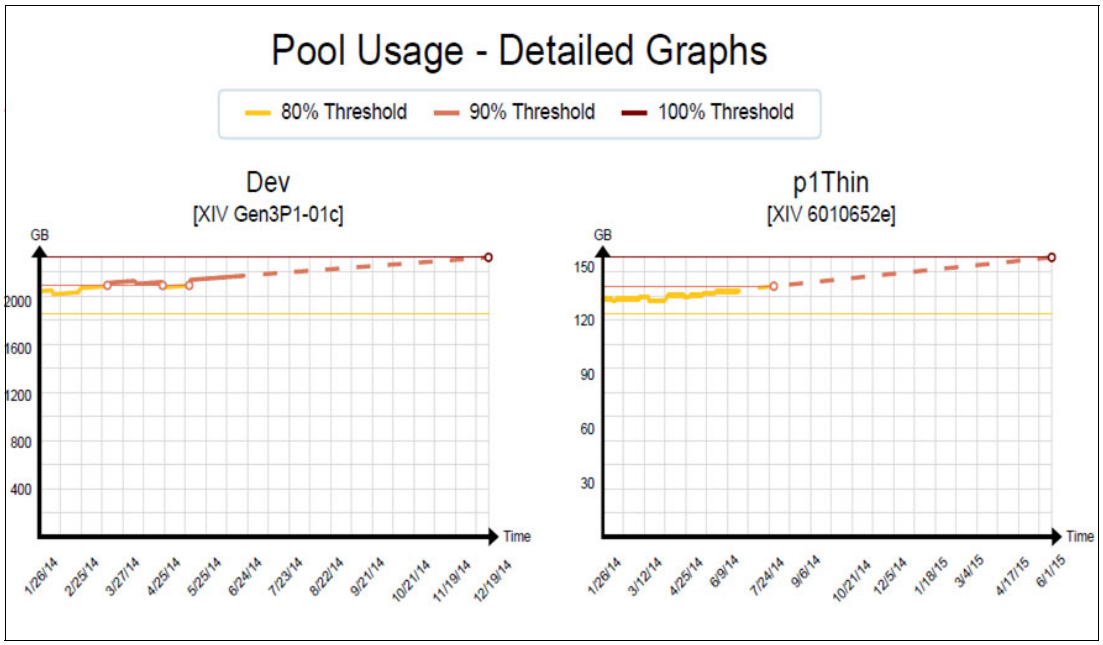

Figure 1-60 Capacity planning report - pool usage

1.5.3 Format of the capacity statistics gathered

In addition to the .pdf report, a comma separate variable (CSV) file is also available. It contains the raw data, collected over time by the Hyper-Scale Manager for each XIV Storage System inventoried. This raw data can be used for further analysis or to create your own, customized reports.

For details about the format and contents of the .CSV file, refer to the *IBM XIV Storage System Management Tools Version 4.4 Operations Guide*, SC27-5986-01, available at:

http://www-01.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/docs/XIV_MT4_4_Operations_Guide.pdf?lang=en-us

1.5.4 Moving the capacity data among Manager instances

Moving the capacity data from one IBM Hyper-Scale Manager to another is done by tasks invoked from the IBM Hyper-Scale Manager management menu.

Access the Hyper-Scale Manager management menu on the Hyper-Scale Manager server (usually, start a telnet session with the server, issue **cd /home/msms/hyperscale** and run the **./management_menu.sh** procedure as shown in Figure 1-61).

```
[root@localhost /]# cd /home/msms/hyperscale
[root@localhost hyperscale]# ./management_menu.sh
```

Figure 1-61 Invoking the management menu on the Hyper-Scale Manager server

1. From the Hyper-Scale Manager User Menu, select Option 1, IBM Hyper-Scale Manager Inventory options, as shown in Figure 1-62.

```
------------------------------------------------------------------
------------ IBM Hyper-Scale Manager v1.5.0.56 ---------------
------------------------------------------------------------------
User Menu
---------------
IBM Hyper-Scale Manager is running

 1) IBM Hyper-Scale Manager Inventory Options
 2) Show Audit Log
 3) Gather IBM Hyper-Scale Manager Logs
 4) Backup/Restore Options
 5) Change Manager Access Code
 6) Replace IBM Hyper-Scale Manager Certificate
 7) Change Network Settings
 8) SMTP Configuration
 9) Start/Stop Hyper-Scale Manager Service
10) Exit
Your Selection>1
```

Figure 1-62 Hyper-Scale Manager main menu

2. Next, select Option 10, Manage Capacity Planning Data, as shown in Figure 1-63 on page 42.

```
------------------------------------------------------------------
------------ IBM Hyper-Scale Manager v1.5.0.56 ---------------
------------------------------------------------------------------

Manager Inventory Options
-------------------------
IBM Hyper-Scale Manager is running.

 1) System Inventory List
 2) Add System
 3) Modify System
 4) Remove System
 5) System Monitoring Suspend
 6) System Monitoring Resume
 7) Re-authenticate All Users
 8) Change System Machine Account
 9) Manage System Certificates
10) Manage Capacity Planning Data
11) Back to previous menu
Your Selection> 10
```

Figure 1-63 Invoking the Manage Capacity Planning Data menu

3. The Manage Capacity Planning Data menu, shown in Figure 1-64, is displayed.

```
------------------------------------------------------------------
------------ IBM Hyper-Scale Manager v1.5.0.56 ---------------

Manage Capacity Planning Data
-----------------------------
1) Import Capacity Data
2) Export Capacity Data
3) Reset Capacity Data
4) Back to previous menu
Your Selection>
```

Figure 1-64 Manage Capacity Planning Data menu

4. From the menu shown in Figure 1-64, you can either export, import, or reset capacity data:

- Exporting the raw capacity data

  The raw data is exported in order to make it available for import to another IBM Hyper-Scale Manager. The file is exported in the same method other files (that is, backups, logs, and more) are exported.

  From the Manage Capacity Planning Data menu shown in Figure 1-64, select option 2, Export Capacity Data. The Capacity Data is exported to a file as shown in Figure 1-65 on page 43.

```
Manage Capacity Planning Data
-----------------------------
1) Import Capacity Data
2) Export Capacity Data
3) Reset Capacity Data
4) Back to previous menu
Your Selection>2
Capacity data was exported successfully to
/home/msms/hyperscale/files/export/XIV_capacity_samples_data_2014-08-13_1921.da
t
Press any key to continue
```

Figure 1-65 Exporting capacity data

Press any key to finish the export task.

The .dat file in the export directory can then be copied or transferred, for import in another Hyper-Scale Manager server:

- Importing the raw capacity data

  A capacity data file that was created on one IBM Hyper-Scale Manager and was exported can be used by another IBM Hyper-Scale Manager in order to maintain the continuity of XIV systems history. The exported capacity data file must be copied into the upload folder (`home/msms/hyperscale/files/upload`) of the targeted Hyper-Scale Manager.

  On the target Hyper-Scale Manager, navigate to the Manage Capacity Planning Data menu, and invoke option1, Import Capacity Data, as shown in Figure 1-66.

```
Manage Capacity Planning Data
-----------------------------
1) Import Capacity Data
2) Export Capacity Data
3) Reset Capacity Data
4) Back to previous menu
Your Selection>1
Please upload the capacity data file (*.dat) to the
(/home/msms/hyperscale/files    /upload) folder and press any key...
Press any key to continue
```

Figure 1-66 Importing the capacity data

  The IBM Hyper-Scale Manager unifies the imported data according to the following continuity rules:

  - Data of XIV systems that are not managed by both IBM Hyper-Scale Manager instances is no longer tracked
  - Data for XIV systems that were already tracked by both IBM Hyper-Scale Manager instances will be overridden, in order to avoid duplicates
  - Data for systems that are currently tracked and whose data was not imported remains unchanged

- Resetting the raw capacity data

  In order to clear the XIV system history from irregularities (that is, machine repurposing), and to allow for collecting raw data from scratch, you can clear the machine history from the previously collected raw data. Simply select option 3, Reset Capacity Data from the Manage Capacity Planning Data menu, which is shown in Figure 1-66 on page 43.

2

IBM Hyper-Scale Mobility

IBM Hyper-Scale Mobility is a powerful function for moving volumes between XIV Storage System systems in a manner that is transparent to host applications. The first section of this chapter provides a general description of the IBM Hyper-Scale Mobility feature, and its architecture and design.

Next, the chapter outlines the prerequisites and conditions of usage for IBM Hyper-Scale Mobility. It highlights functions of both the graphical user interface (GUI) and XIV Storage System command-line interface (XCLI), which support IBM Hyper-Scale Mobility, and then provides step-by-step illustrated examples.

2.1 Introduction to IBM Hyper-Scale Mobility

Introduced with the XIV Storage System Software V11.3, IBM Hyper-Scale Mobility is a feature of the XIV system. IBM Hyper-Scale Mobility enables clients to move a volume from one XIV Gen3 system to another (within the same physical site) in a manner that is indiscernible to host applications before, during, and after the volume migration. IBM Hyper-Scale Mobility, as with all XIV software features, is included in the XIV base software.

No-charge: IBM Hyper-Scale Mobility is a no-charge feature that is included with XIV Storage System Software V11.3.

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 XIV scalability, and directly addresses several client storage-related issues:

- Managing storage growth
- Providing more flexibility in capacity forecasting
- Managing costs
- Dealing with storage performance problems

IBM Hyper-Scale Mobility capabilities also provide the XIV Storage System with an even better total cost of ownership (TCO) position, and improved service continuity.

IBM Hyper-Scale Mobility enables the following actions:

- Balancing workloads across systems without service downtime, for better performance
- Retiring systems gracefully, and upgrading to the latest XIV technology
- Detaching storage and host maintenance cycles

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

- XIV Storage System online volume migration is useful when combined with XIV Storage System thin provisioning.

  When an XIV Storage System is overprovisioned, and the total system hard space is approaching system-wide depletion, Hyper-Scale provides a way to move volumes to another system, freeing up needed hard pool space.

  Alternatively, the volumes in the XIV thin pool that are unexpectedly using more hard space than planned can be moved to an XIV Storage System with plenty of hard pool space, with no effect on the application host severs.
- IBM Hyper-Scale Mobility is equally useful in a regularly provisioned XIV Storage System that is running out of physical storage space.
- Volume migration using Hyper-Scale Mobility can be used in a situation where a client has two XIV Storage System systems with different performance characteristics. For example, if an application that was on a 12-module XIV Storage System configuration needs more performance capacity, Hyper-Scale can be used to move those volumes to a 15-module XIV Storage System with solid-state drives (SSDs).
- IBM Hyper-Scale Mobility can also be used to rebalance the workload, away from a system that is heavily used to a system that is underused. It does not matter if the capacity imbalance was performance-related or data storage-related. IBM Hyper-Scale Mobility can help resolve both situations.

2.2 IBM Hyper-Scale Mobility design considerations

IBM Hyper-Scale Mobility is designed to provide a process to move volumes between IBM XIV Storage Systems with little host effect. Moving volumes usually requires defining a volume on the destination (or target) XIV Storage System, porting data, and disrupting host applications while 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 XIV Gen3 Storage System while 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 (proxy) from the source XIV Storage System to the destination XIV Storage System 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, should 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 XIV GUI or the XCLI, see 2.7, “Using IBM Hyper-Scale Mobility” on page 52.

2.3 IBM Hyper-Scale Mobility requirements

The integration of IBM Hyper-Scale Mobility in the XIV Storage System provides initial support for hosts that are running IBM AIX®, Linux, and various Windows Server versions.

To cause minimal disruption to the host, and to enable uninterrupted I/O operations, several 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 XIV systems must be accessible to the host system concurrently. This might require changes to the zoning in use by the host that is involved. The two XIV Storage System systems also need zoning to enable connections for data migration.

- IBM XIV Storage System systems.

  Obviously, there must be two XIV systems, one being the source and the other the 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 should also be given for any possible increase in snapshot space that might be required on the destination system.

2.4 IBM Hyper-Scale Mobility process description

From a design standpoint, the IBM Hyper-Scale Mobility process that is used to move a volume between XIV 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 2-1.

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.

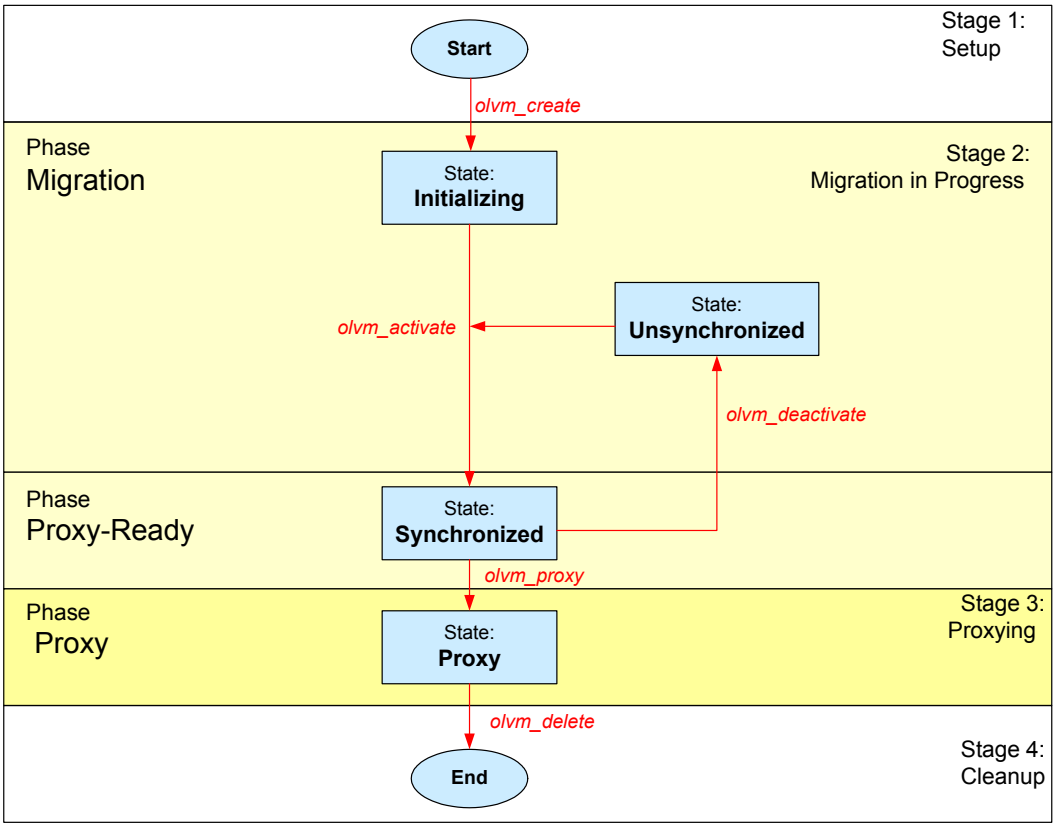

Figure 2-1 IBM Hyper-Scale Mobility flow diagram

Stage 1: Setup

The first stage of the IBM Hyper-Scale Mobility is to create the volume on the destination XIV Storage System, 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 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 storage system 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.

2.5 Usage prerequisites and considerations

In addition to the general requirements for using IBM Hyper-Scale Mobility that are described in 2.3, “IBM Hyper-Scale Mobility requirements” on page 47, there are additional prerequisites and considerations.

2.5.1 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 should also be given for any possible increase in snapshot space that might be required on the target system.

2.5.2 Management workstation connectivity

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

2.5.3 Connectivity between the source and destination XIV systems

To effectively use IBM Hyper-Scale Mobility, there must be adequate connectivity between the source and destination XIV Storage System 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 you plan IBM Hyper-Scale Mobility activities, you can follow the same guidelines and leading practices regarding connectivity between source and destination XIV Storage System systems as those followed with synchronous remote mirroring. For more information, see the chapters about remote mirroring and synchronous remote mirroring in *IBM XIV Storage System Copy Services and Migration*, SG24-7759.

Note, however, that as made available with Version 11.3 of the IBM XIV Storage System Software, the IBM Hyper-Scale Mobility restricts the source and target XIV systems to be located at a short distance from each other (that is, within the same physical site). Also, the two systems must be connected over Fibre Channel (FC). Internet Small Computer System Interface (iSCSI) is not supported at this time.

2.5.4 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 XIV Storage System has the appropriate access to the destination XIV Storage System:

- 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.

2.5.5 Other current restrictions

In addition to the volumes and pool considerations previously mentioned, the following current limitations apply:

- Using Hyper-Scale Mobility to mobilize volumes that are mapped to UNIX or Linux hosts, and are backed up using IBM Tivoli® Storage FlashCopy® Manager version 3.2 or earlier, is not supported.
- Using Hyper-Scale Mobility to mobilize volumes that are mapped to Windows Server hosts, and are backed up using Microsoft Volume Shadow Copy Service (VSS), with IBM XIV Provider for Microsoft Windows VSS version 2.3.2 or earlier, is not supported. To use Hyper-Scale Mobility in this case, IBM XIV Provider for Microsoft Windows VSS version 2.4.0 or later must be installed.

Restriction: IBM Tivoli Storage FlashCopy Manager for Windows is relying on Microsoft VSS. To use Hyper-Scale Mobility with IBM Tivoli Storage FlashCopy Manager for Windows, IBM XIV Provider for Microsoft Windows VSS version 2.4.0 or later must be installed.

- Using Hyper-Scale Mobility to mobilize volumes that are mapped to Windows Server hosts, and are protected by IBM Storage Enabler for Windows Failover Clustering version 1.1.0 or earlier, is not supported. To use Hyper-Scale Mobility in this case, IBM Storage Enabler for Windows Failover Clustering version 1.2.0 or later must be installed. Additional information may be available from the respective release notes of the products mentioned in the immediately preceding restrictions note.

  For IBM Tivoli Storage FlashCopy Manager for UNIX and Linux, refer to the release notes of version 4.1 or later.

  For IBM XIV Provider for Microsoft Windows VSS release notes, refer to the following website:

  http://pic.dhe.ibm.com/infocenter/strhosts/ic/topic/com.ibm.help.strghosts.doc/vss-homepage.html

  For IBM Storage Enabler for Windows Failover Clustering release notes, refer to the following website:

  http://pic.dhe.ibm.com/infocenter/strhosts/ic/index.jsp?topic=%2Fcom.ibm.help.strghosts.doc%2Fmscs-homepage.html

2.6 Management software support for IBM Hyper-Scale Mobility

XIV System Management Software V4.2 introduces support for the IBM Hyper-Scale Mobility feature through both the GUI and XCLI. Either tool can be used to perform all the tasks that are related to IBM Hyper-Scale Mobility.

2.6.1 XIV Management GUI features

In support of the IBM Hyper-Scale Mobility feature, the XIV Management GUI was updated with a new Volume Mobility view, which is available from the Remote menu in the left pane, as shown in Figure 2-2.

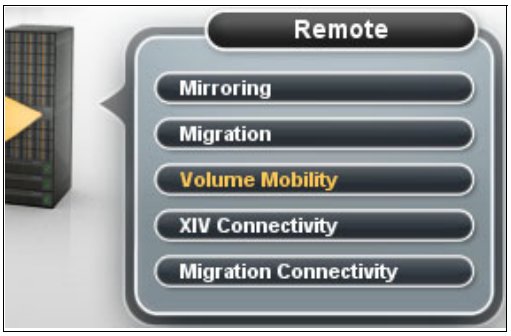

Figure 2-2 Volume Mobility menu option to support IBM Hyper-Scale Mobility

The Volume Mobility view can also be accessed in the View section of the Main Menu, as shown in Figure 2-3.

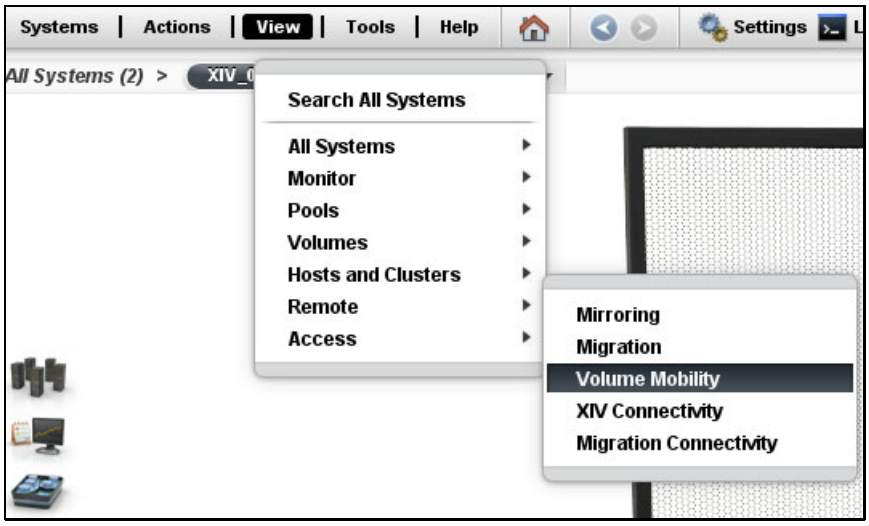

Figure 2-3 Online Migration in the View menu

Other views and GUI panels were updated to reflect support for IBM Hyper-Scale Mobility. The Multiple Storage Management System was updated to accommodate IBM Hyper-Scale Mobility.

2.6.2 XCLI functions that support IBM Hyper-Scale Mobility

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

`olvm_list`	Displays details for IBM Hyper-Scale Mobility relationships on a local system.
`olvm_create`	Creates an IBM Hyper-Scale Mobility process for a volume.
`olvm_activate`	Activates (starts) a previously created IBM Hyper-Scale Mobility process.
`olvm_proxy`	The volume in the proxy-ready state enters a proxy state.
`olvm_deactivate`	Deactivates the IBM Hyper-Scale Mobility process for a specified volume.
`olvm_abort`	Aborts an IBM Hyper-Scale Mobility relationship.
`olvm_delete`	Deletes an IBM Hyper-Scale Mobility relationship.
`olvm_delete_all_ref_to_src`	Delete all IBM Hyper-Scale Mobility relationships to a specified source target.

2.7 Using IBM Hyper-Scale Mobility

In the following examples, you port an online volume by using the XIV GUI and the XCLI. In each example, the volume that is being ported is mapped over iSCSI to a Red Hat Enterprise Linux (RHEL) host, and the host is actively generating I/Os to the volume throughout the migration.

Here are the details of the migration example:

- The source XIV Storage System is named XIV_02_1310114.
- The destination XIV Storage System is named XIV_PFE2_1340010.
- The GUI volume to be ported is name HSM_002.
- The XCLI volume to be ported is named HSM_001.
- The AIX host with LUN mapping is named ITSO_LPAR1.
- The number of physical paths from the host to each XIV Storage System is two.

2.7.1 Using the XIV GUI to use IBM Hyper-Scale Mobility

This section illustrates using IBM Hyper-Scale Mobility by using the XIV GUI. We have decomposed the process in to a series of six tasks.

Task 1: Planning and preparation

To complete planning and preparation, follow these steps:

1. First, confirm that both the source and destination XIV systems are configured and accessible in the GUI, and that there is FC connectivity between these systems. Figure 2-3 on page 52 shows the source and destination systems with confirmed connectivity between the two systems.

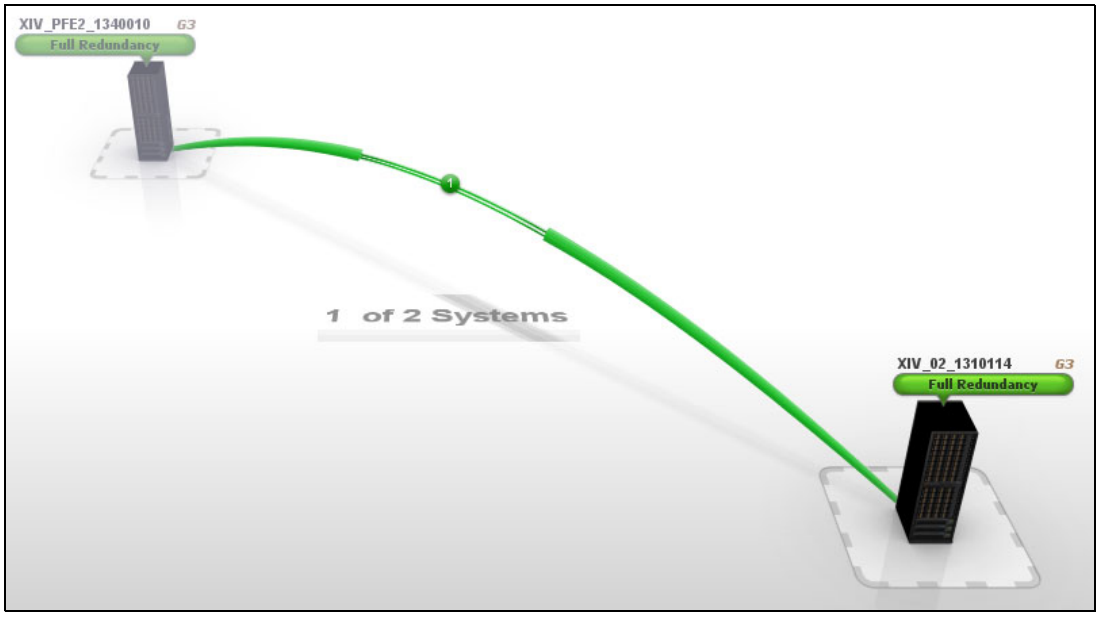

Figure 2-4 GUI Connectivity view of IBM Hyper-Scale source and destination systems

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

Example 2-1 Check multipath connectivity

```
[p7-770-02v13:root:/home/root:] xiv_devlist

XIV Devices
------------------------------------------------------------------------
Device       Size (GB)  Paths  Vol Name  Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------
/dev/hdisk1  103.2      6/6    HSM_001   121     1310114  ITSO_LPAR1
------------------------------------------------------------------------
```

```
/dev/hdisk2  103.2      6/6    HSM_002   122     1310114  ITSO_LPAR1
-----------------------------------------------------------------------


Non-XIV Devices
-----------------------------
Device       Size (GB)  Paths
-----------------------------
/dev/hdisk0  54.8       N/A
-----------------------------
```

Note the Vol ID of 122 and XIV ID of 1310114 of volume HSM_002, indicating that the volume is presented to the host from the XIV_02_1310114 XIV Storage System.

3. Lastly, you can validate that the host is actively sending I/Os to the source volume from the Performance Chart view of that volume in the XIV Top tool, as shown in Figure 2-5.

Figure 2-5 Initial XIV Top Performance Chart view of source volume

As part of the preparation, you can also define the host to the target XIV Storage System. This action can also be completed while the migration is taking place.

Next, you will create the host definition and port definition for your host on the destination XIV Storage System.

4. Click the **Hosts and Clusters** link in the **Hosts and Clusters** menu on the destination XIV Storage System, as shown in Figure 2-6.

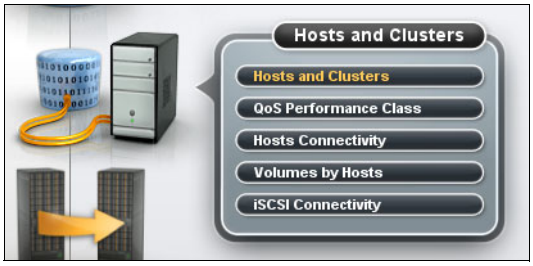

Figure 2-6 Destination XIV Storage System Hosts and Clusters menu

This illustration applies to an FC-connected host. For an iSCSI-connected host, see the procedure that is explained in *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904.

5. Next, right-click in the window and select the **Add Host** option, as shown in Figure 2-7.

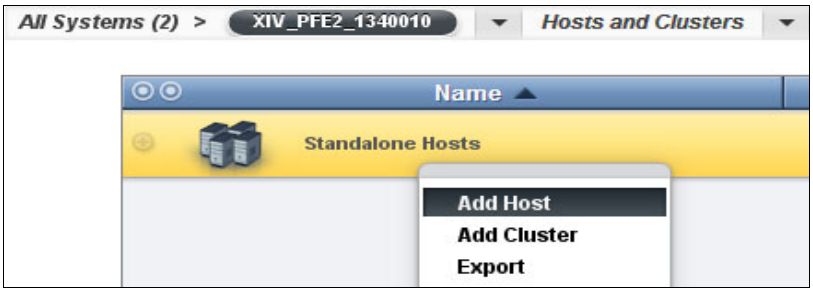

Figure 2-7 Add Host on the destination XIV Storage System

6. In the Add Host window, enter the details for the host and click **Add**, as shown in Figure 2-8. Ensure that the details match the corresponding host details from the source XIV Storage System. In this example, the Name of the host is `ITSO_LPAR1`, the Type is default, and the Challenge Handshake Authentication Protocol (CHAP) values are empty.

Figure 2-8 Add host details about the destination XIV Storage System

7. After the host is added to the destination XIV Storage System, add the ports for the host. Right-click the host entry and select the **Add Port** option, as shown in Figure 2-9.

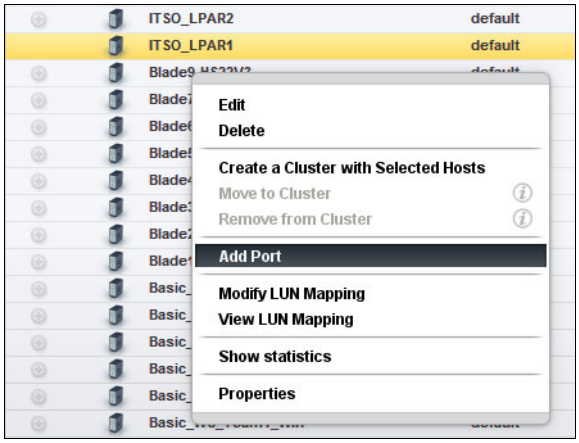

Figure 2-9 Add host port on the destination XIV Storage System

8. In the Add Port window, enter the appropriate port details, as shown in Figure 2-10, and click **Add** to add the port to the host.

   Ensure that the details match the corresponding port details from the source XIV Storage System. In this example, you choose FC for the Port Type, and the Port Name is 10000000C99987D9.

Figure 2-10 Enter port details about the destination XIV Storage System

9. Add other FC ports to the hosts as well, as demonstrated in Figure 2-9 on page 55 and Figure 2-10 with the first port.
10. The new host and port definitions should now be visible in the Hosts and Clusters view on the destination XIV Storage System, as shown in Figure 2-11.

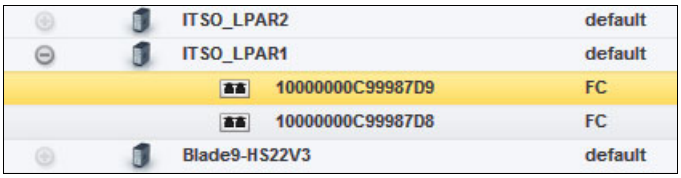

Figure 2-11 Newly defined host and port on the destination XIV Storage System

Task 2: Setup

To complete the Setup task, complete the following steps:

1. To create the IBM Hyper-Scale Mobility relationship, click the **Volume Mobility** link in the **Remote** menu, as shown in Figure 2-12.

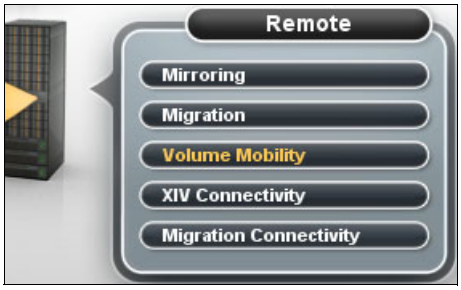

Figure 2-12 Volume Mobility link in the Remote menu

2. From the Volume Mobility view, click the **Move Volume to Another System** link in the toolbar, as shown in Figure 2-13.

Figure 2-13 Move Volume to Another System link in the toolbar

Alternatively, from the Volumes and Snapshots view, right-click the source volume and select the **Move Volume to Another System** menu option, as shown in Figure 2-14. Note that the Size of the source volume is 103 gigabytes (GB), and its Used capacity is 96 GB.

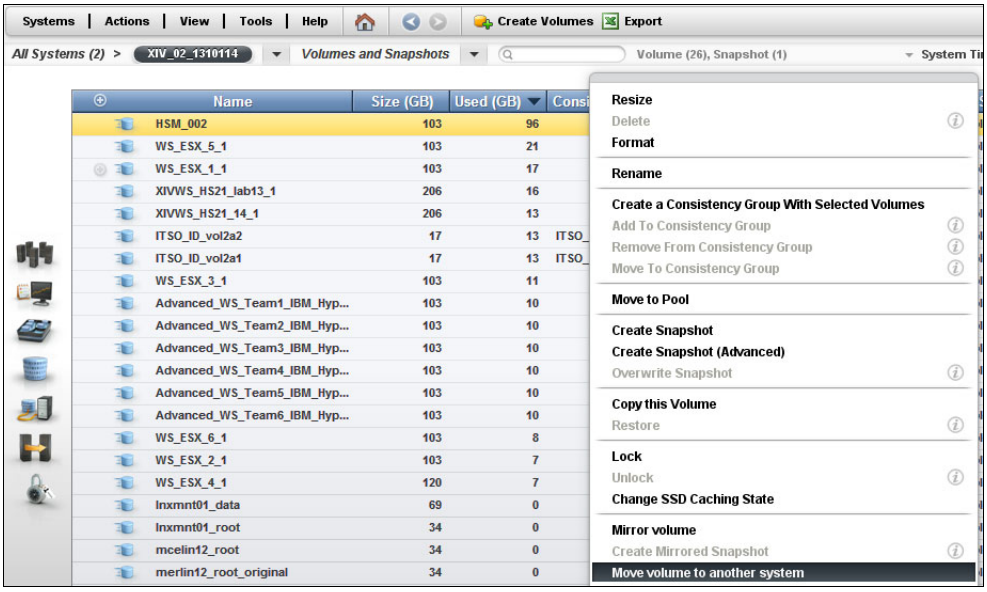

Figure 2-14 Move Volume to Another System menu option in Volumes and Snapshots view

3. Either method opens the Create Online Volume Mobility window, as shown in Figure 2-15.

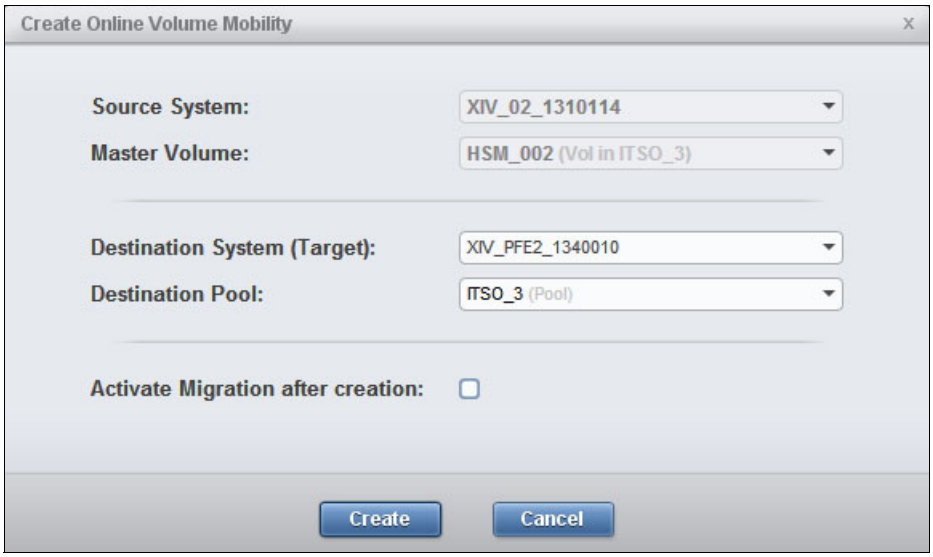

Figure 2-15 Create Online Volume Mobility window

4. Use the drop-down menus that are shown in Figure 2-15 on page 57 to select the values for **Source System**, **Master Volume**, **Destination System (Target)**, and **Destination Pool**.

   In this example, the Source System is XIV_02_1310114 and the Master Volume is HSM_002. The Destination System is XIV_PFE2_1340010, and the Destination Pool is ITSO_3.

5. To immediately activate the IBM Hyper-Scale Mobility relationship upon creation, select the box. In this example, you clear the box, as you will explicitly activate the relationship in the next step.

   Click **Create** to create the relationship.

   Tip: If you opened the Create Online Volume Mobility panel by right-clicking a volume, the values for Source System and Master Volume are pre-populated and you cannot edit them.

   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.

6. The Volume Mobility view for the source XIV Storage System now displays the new IBM Hyper-Scale Mobility relationship, as shown in Figure 2-16. Notice that the Phase and Status are both Inactive.

Figure 2-16 Newly created, inactive IBM Hyper-Scale Mobility relationship on the source XIV

The Volume Mobility view for the destination XIV Storage System also displays the new IBM Hyper-Scale Mobility relationship, as shown in Figure 2-17. Notice that the Phase and Status are also both Inactive.

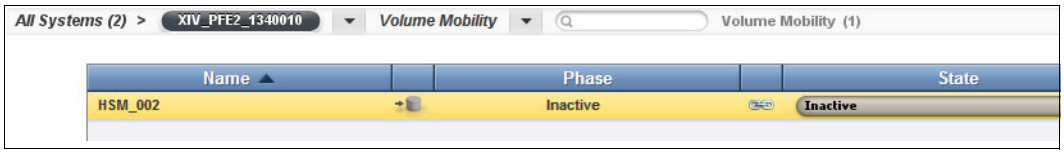

Figure 2-17 Newly created, inactive IBM Hyper-Scale Mobility relationship on the source XIV

7. Additionally, you can verify that the destination volume was created by right-clicking the IBM Hyper-Scale Mobility relationship and selecting the **Show Destination Volume** menu option, as shown in Figure 2-18.

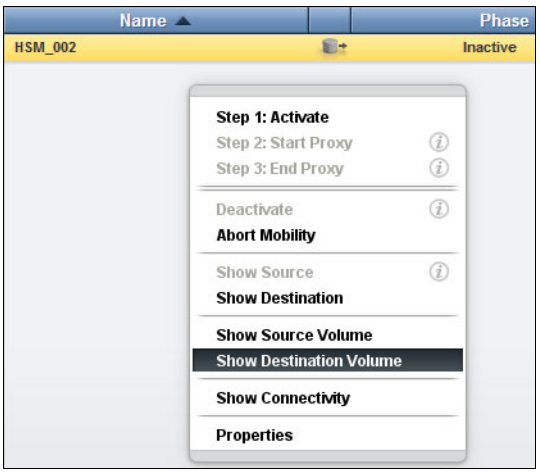

Figure 2-18 Menu for IBM Hyper-Scale Mobility relationships

Tip: The menu for IBM Hyper-Scale Mobility relationships offers useful shortcuts that are relevant to IBM Hyper-Scale Mobility tasks. With this menu, you can quickly and easily navigate between source and destination IBM Hyper-Scale Mobility systems and volumes.

8. In Figure 2-19, you can see that the destination volume was created on the destination XIV Storage System. The size of the destination volume matches the size of the source volumes, but the Used capacity of the volume is 0 GB. This is expected, because you have not yet activated the IBM Hyper-Scale Mobility relationship, and the synchronization of the destination volume has not yet begun.

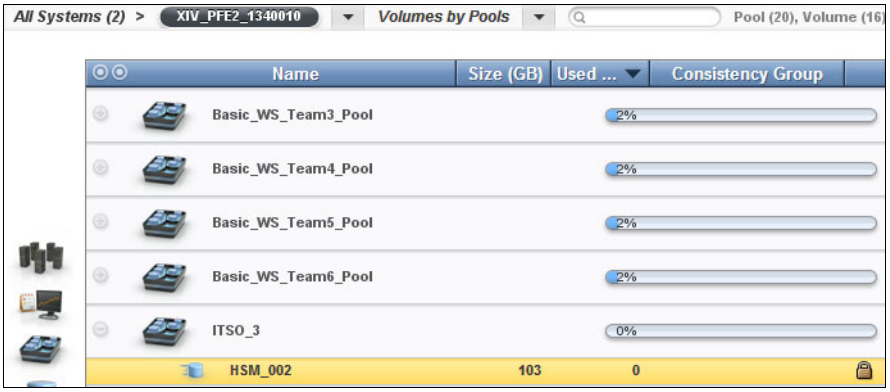

Figure 2-19 Destination volume created on the destination XIV

Note: Although the destination volume has been created, you cannot yet map the volume 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.
- Change SSD Caching State.
- View Volume Mapping.
- Show Statistic.
- Properties.

Task 3: Migration

To complete the Migration task, perform the following steps:

1. In the source XIV Storage System (Alba), right-click the row for the IBM Hyper-Scale Mobility relationship that you created and select the **Step 1: Activate** option, as shown in Figure 2-20.

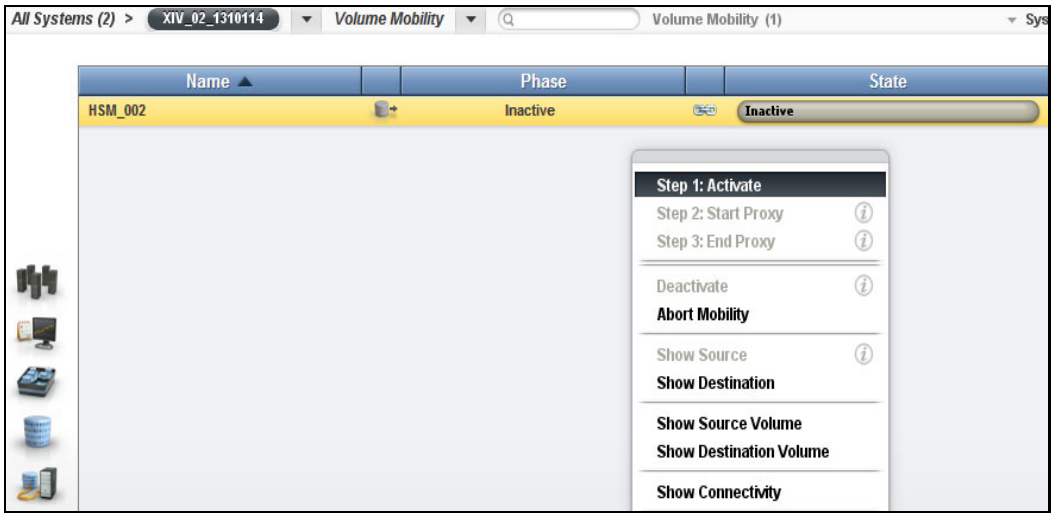

Figure 2-20 Activate the IBM Hyper-Scale Mobility relationship

This action activates the relationship and queues the process to synchronize the destination volume on the destination XIV Storage System.

Remember: Only one IBM Hyper-Scale Mobility process at a time is synchronized from the source XIV Storage System.

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 monitored by watching the status bar, as shown in Figure 2-21. The phase is Migration and the status is Initialization.

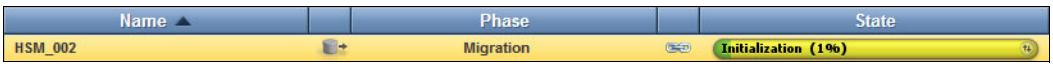

Figure 2-21 IBM Hyper-Scale Mobility status at the source XIV Storage System

2. After the synchronization is complete, the Phase changes to Proxy Ready and the Status changes to Synchronized when you view them on the source XIV Storage System, as shown in Figure 2-22.

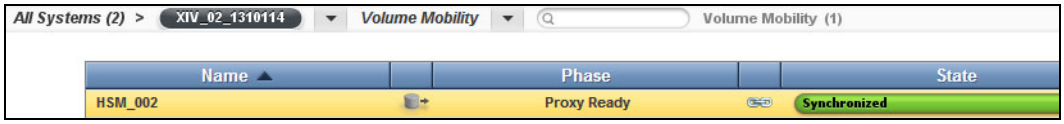

Figure 2-22 Proxy Ready and Synchronized at the source XIV Storage System

3. Right-click the IBM Hyper-Scale Mobility relationship from the destination XIV Storage System and select **Show Destination**. You can see that the Phase, when you view it from the destination XIV Storage System, is Proxy Ready and the status is Consistent, as shown in Figure 2-23.

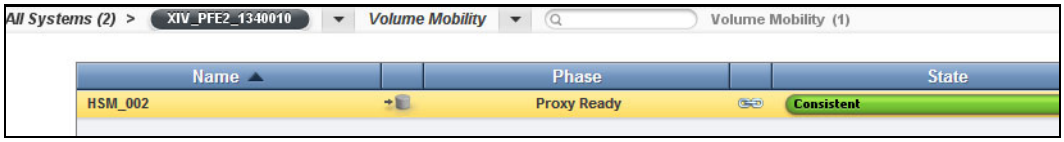

Figure 2-23 Proxy Ready and Consistent at the destination XIV Storage System

Now, the destination XIV Storage System contains a consistent copy of the source volume, in addition to the necessary host and port definitions. You are ready to enter the actual migration stage.

Task 4: Proxying

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

1. Right-click the row for the Hyper-Scale Mobility relationship in the Volume Mobility view on the source XIV Storage System and select **Step 2: Start Proxy**, as shown in Figure 2-24.

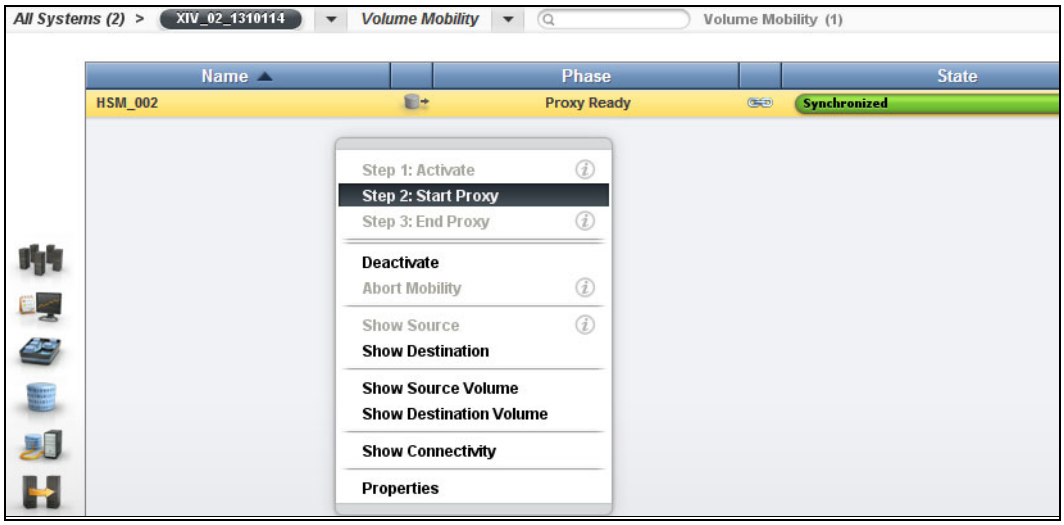

Figure 2-24 Start IBM Hyper-Scale Mobility Proxy

2. As shown in Figure 2-25, a warning message is displayed, indicating that any source volume snapshots that exist will be deleted if the proxy is started.

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

Figure 2-25 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 XIV Storage System, and any host I/O is still synchronized between both the source and destination volumes.

Clicking **OK** in this window triggers the source XIV Storage System to proxy all subsequent host I/O to the volume on the destination XIV Storage System only, and the source volume becomes inconsistent and unavailable. In fact, 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. You can verify that the Proxy phase was initiated from the Volume Mobility view on the source XIV Storage System, as shown in Figure 2-26.

   Note that the Phase is now Proxy and the Status is now Proxy-Source.

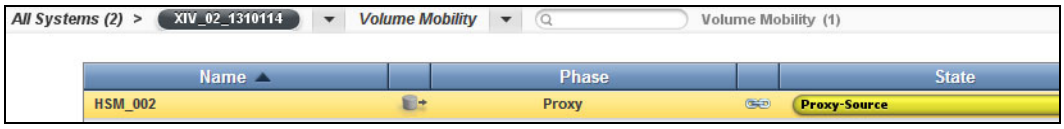

Figure 2-26 Proxy phase viewed from the source XIV Storage System

4. You can further validate the Proxy phase from the Online Migration view on the destination XIV Storage System, as shown in Figure 2-27.

   Note that the Phase is now Proxy and the Status is now Proxy-Destination.

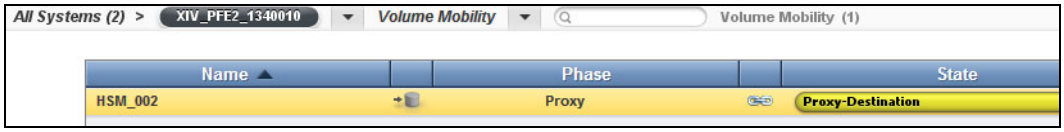

Figure 2-27 Proxy phase viewed from the destination XIV Storage System

5. Now that the Proxy phase is successfully initiated, you are ready to map the newly created volume to the host on the destination XIV Storage System. Go to the destination volume by right-clicking the IBM Hyper-Scale Mobility relationship and selecting the **Show Destination Volume** option, as shown in Figure 2-28.

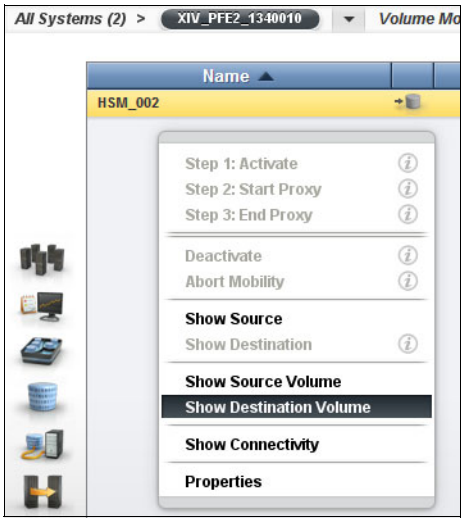

Figure 2-28 Show Destination Volume from the destination IBM Hyper-Scale Mobility relationship

6. Next, right-click the volume and select **Map Selected Volumes**, as shown in Figure 2-29.

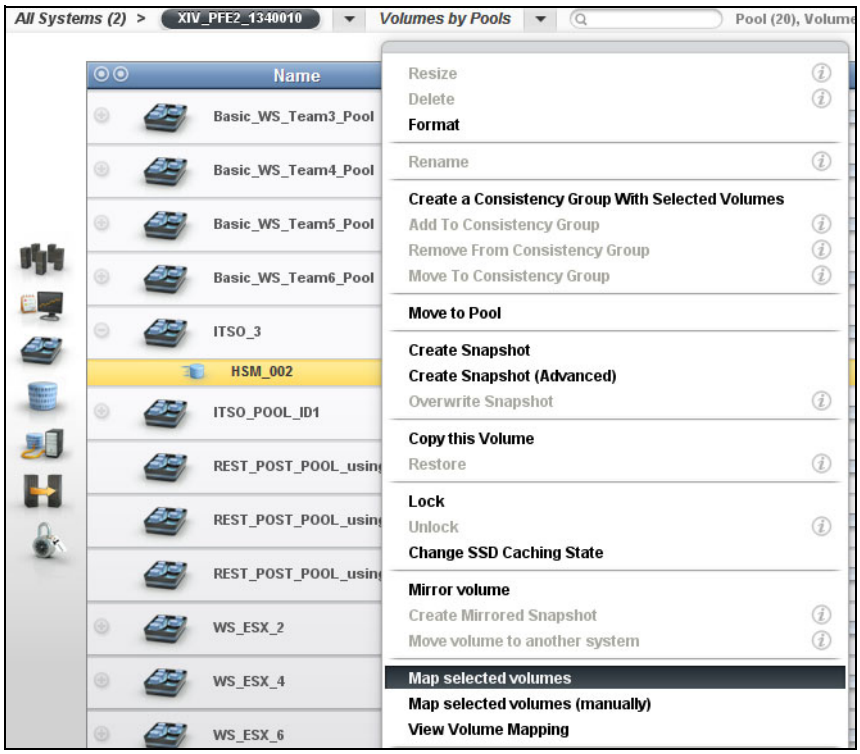

Figure 2-29 Map the migrated LUN to the host on the destination XIV Storage System

7. In the window that opens, select the new host from the menu and click **OK** to map the volume, as shown in Figure 2-30.

Figure 2-30 Select the host to map

8. After the destination volume is mapped to the host, you need to rescan the devices from the host 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 2-2.

Example 2-2 Host rescan

```
[p7-770-02v13:root:/home/root:] xiv_fc_admin -R
[p7-770-02v13:root:/home/root:] xiv_devlist

XIV Devices
------------------------------------------------------------------------
Device       Size (GB)  Paths  Vol Name  Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------
/dev/hdisk1  103.2      6/6    HSM_001   121     1310114  ITSO_LPAR1
------------------------------------------------------------------------
/dev/hdisk2  103.2      12/12  HSM_002   122     1310114  ITSO_LPAR1
------------------------------------------------------------------------


Non-XIV Devices
----------------------------
Device       Size (GB)  Paths
----------------------------
/dev/hdisk0  54.8       N/A
----------------------------
```

Note the 12 paths (six additional) for the destination volume, HSM_002, in the bold row.

Task 5: Cleanup

Now that you have validated that the host has connectivity to the volume through the new paths to the destination XIV Storage System, you are ready to unmap the volume from the host on the source XIV Storage System, and remove the paths to the source volume from the host.

To accomplish this task, complete the following steps:

1. In the GUI, go to the Volumes by Hosts view of the source XIV Storage System.

   Right-click the source volume and select the **Unmap** option, as shown in Figure 2-31.

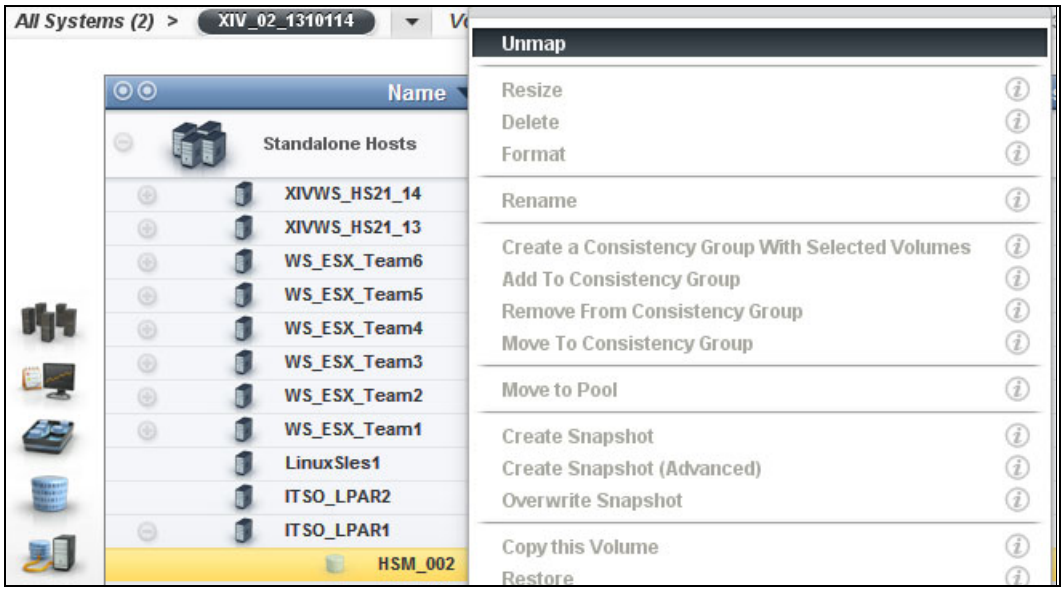

Figure 2-31 Unmap the source volume from the host

2. A confirmation window opens and prompts you to confirm that you want to unmap the selected volume, as shown in Figure 2-32. Click **OK** to unmap the source volume from the host.

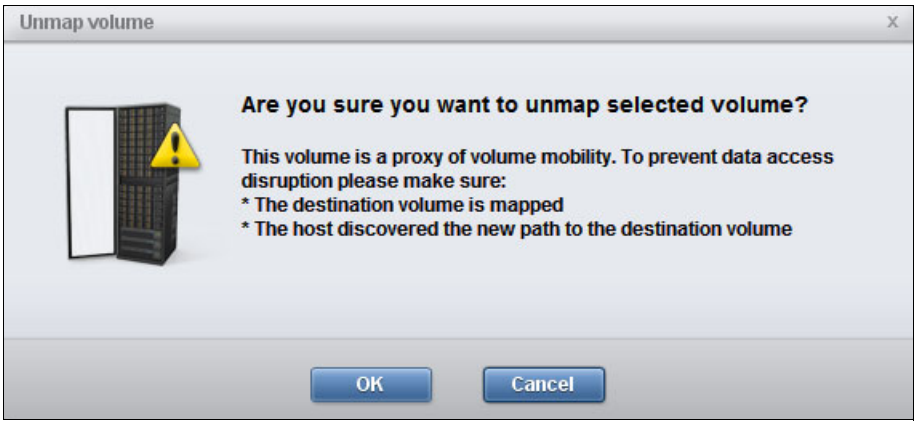

Figure 2-32 Confirm the unmapping of the source volume from the host

3. Now that you have unmapped the volume from the host, you have only six enabled paths to the volume left from the host. Run the **xiv_fc_admin -R** command, and then run the **xiv_devlist** command from the host, as shown in Example 2-3.

Example 2-3 Half the paths enabled

```
[p7-770-02v13:root:/home/root:] xiv_fc_admin -R
[p7-770-02v13:root:/home/root:] xiv_devlist

XIV Devices
------------------------------------------------------------------------
Device       Size (GB)  Paths  Vol Name  Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------
/dev/hdisk1  103.2      6/6    HSM_001   121     1310114  ITSO_LPAR1
------------------------------------------------------------------------
/dev/hdisk2  103.2      6/12   HSM_002   122     1340010  ITSO_LPAR1
------------------------------------------------------------------------


Non-XIV Devices
----------------------------
Device       Size (GB)  Paths
----------------------------
/dev/hdisk0  54.8       N/A
----------------------------
```

Note that the two paths for the destination volume, HSM_002, in the bold row. Also, note that the Vol ID is still 122 and the XIV ID has changed to 1340010, which indicates that the host is connected *only* to the destination volume on the destination XIV Storage System. The destination XIV Storage System keeps the same Vol ID, if possible.

Remove the failed paths in AIX. A sample script for removal of the failed paths is shown Example 2-4.

Example 2-4 Removing the paths

```
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
```

After the removal, six of six paths are available, as shown in Example 2-5.

Example 2-5 Failed paths removed

```
[p7-770-02v13:root:/home/root:] xiv_devlist

XIV Devices
------------------------------------------------------------------------
Device       Size (GB)  Paths  Vol Name  Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------
/dev/hdisk1  103.2      6/6    HSM_001   121     1310114  ITSO_LPAR1
------------------------------------------------------------------------
```

```
/dev/hdisk2  103.2      6/6    HSM_002   122     1340010  ITSO_LPAR1
-------------------------------------------------------------------


Non-XIV Devices
----------------------------
Device       Size (GB)  Paths
----------------------------
/dev/hdisk0  54.8       N/A
----------------------------
```

Task 6: Post-cleanup

You have successfully used IBM Hyper-Scale Mobility to move the volume HSM_002 from XIV Storage System 1310114 to XIV Storage System 1340010 without incurring any downtime on the host.

The final step is to end the proxy. To accomplish this task, complete the following steps:

1. In the XIV GUI, go to the Volume Mobility view of the source XIV Storage System.

   Right-click the volume relationship and select the **Step 3: End Proxy** option, as shown in Figure 2-33.

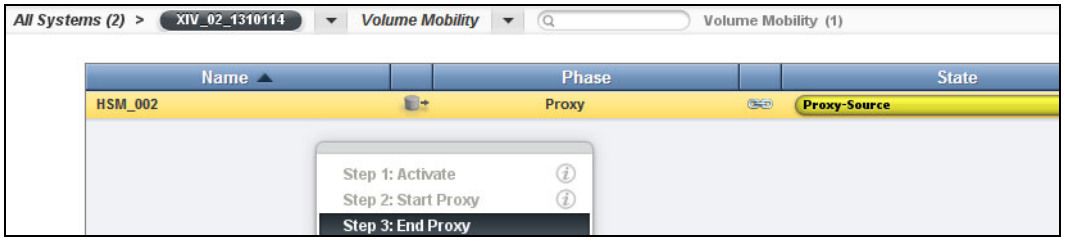

Figure 2-33 End the IBM Hyper-Scale Mobility proxy

2. A confirmation window opens, as shown in Figure 2-34, to ensure that you mapped the host to the destination volume and confirmed access to it. This was confirmed in the Cleanup stage.

   Click **OK** to end the proxy.

Figure 2-34 Confirm the ending of the IBM Hyper-Scale Mobility proxy

3. The proxy has been removed, and you can confirm the expected end state of the volumes on both the source and destination XIV systems.

   On the source XIV Storage System, you can see from the Volumes by Pools view that the volume HSM_002 no longer exists, as shown in Figure 2-35.

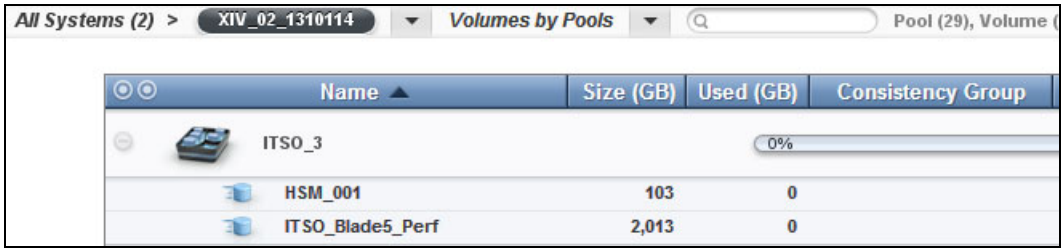

Figure 2-35 Confirm that the source volume is no longer on the source XIV Storage System

4. On the destination XIV Storage System, you can see from the Volumes by Pool view that the volume HSM_002 exists, as shown in Figure 2-36. Note that the Size of 103 GB and Used capacity of 96 GB each match the corresponding values from the source volume that were noted in “Task 2: Setup” on page 56.

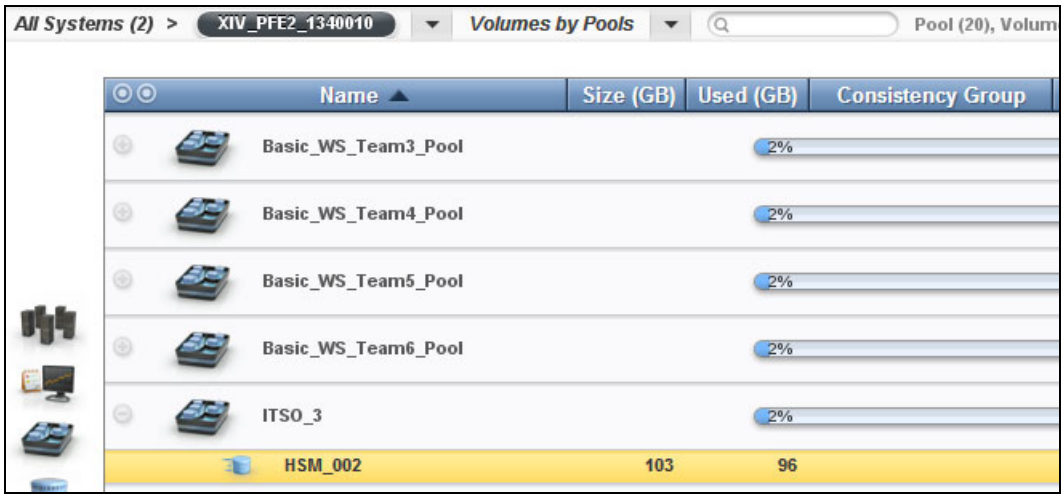

Figure 2-36 Confirm that the destination volume exists on the destination XIV Storage System

5. If you right-click the volume and select **Show Statistics**, you can see that the host I/Os are still active on the destination volume, as shown in Figure 2-37.

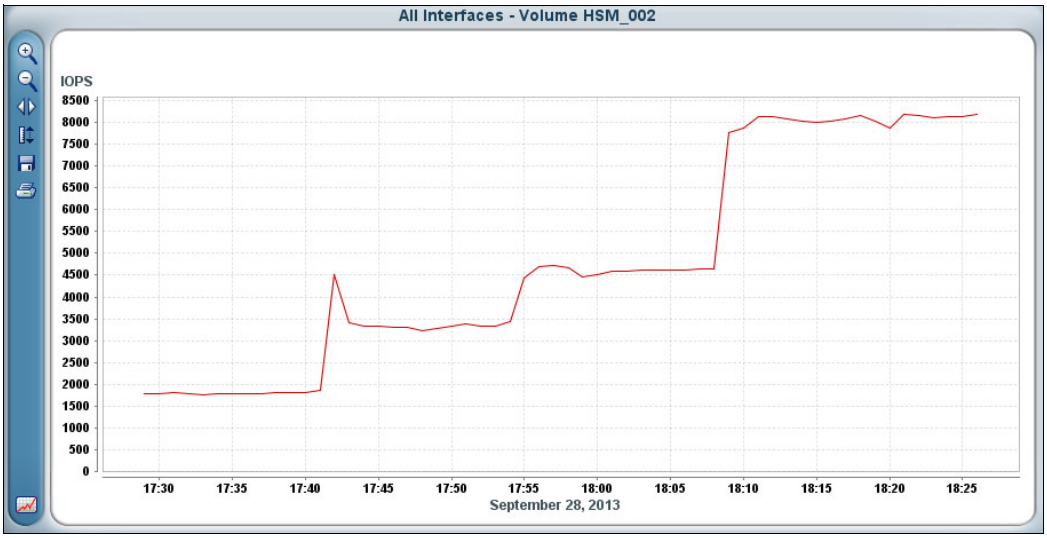

Figure 2-37 Statistics view of the destination volume after IBM Hyper-Scale Mobility is complete

2.7.2 Using the XIV command-line interface to use IBM Hyper-Scale Mobility

This section illustrates using the XCLI to use IBM Hyper-Scale Mobility. The process consists of the same six tasks that are illustrated in 2.7.1, “Using the XIV GUI to use IBM Hyper-Scale Mobility” on page 53.

Task 1: Planning and preparation

To perform an online migration of an XIV volume by using the XCLI, follow these steps:

1. First, confirm that both the source and destination XIV systems are configured and accessible through the XCLI, and that there is connectivity between these systems.

   Example 2-6 shows the **`target_connectivity_list`** command that is run on the source XIV Storage System to confirm the connectivity. The output that is shown in Example 2-6 shows that the destination XIV Storage System is connected through three active FC connections.

Example 2-6 The target_connectivity_list command

```
XIV_02_1310114>>target_connectivity_list
Target Name      Remote Port        FC Port         IP Interface   Active   Up
XIV_04_1340008   500173809C480171   1:FC_Port:9:4                  yes      yes
XIV 7820501      9.11.235.67                        M7P1_Tucson    yes      yes
EMC_VMAX         50000974C0082518   1:FC_Port:9:4                  yes      yes
XIV_PFE2_1340010 500173809C4A0182   1:FC_Port:8:3                  yes      yes
XIV_PFE2_1340010 500173809C4A0162   1:FC_Port:6:3                  yes      yes
XIV_PFE2_1340010 500173809C4A0142   1:FC_Port:4:3                    yes      yes
```

2. In addition, confirm that any host that has mappings to the source volume on the source XIV Storage System has the appropriate access to the destination XIV Storage System.

Consider the following items:

- If you are migrating an iSCSI LUN, ensure that the host can access the destination system over the IP network.
- If you are migrating an FC LUN, ensure that the host is correctly zoned to the destination system.

3. Now, you can validate the host multipath connectivity to the source volume by running the **xiv_devlist** command of the HAK, as shown in Example 2-7.

Example 2-7 The xiv_devlist command

```
[p7-770-02v13:root:/home/root:] xiv_devlist

XIV Devices
------------------------------------------------------------------------
Device       Size (GB)  Paths  Vol Name  Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------
/dev/hdisk1  103.2      6/6    HSM_001   121     1310114  ITSO_LPAR1
------------------------------------------------------------------------
/dev/hdisk2  103.2      6/6    HSM_002   122     1340010  ITSO_LPAR1
------------------------------------------------------------------------


Non-XIV Devices
-----------------------------
Device       Size (GB)  Paths
-----------------------------
/dev/hdisk0  54.8       N/A
-----------------------------
```

Note the two paths for the source volume, HSM_001, in the bold row. Also, note the Vol ID of 121 and XIV ID of 1310114, which indicates that the volume is presented to the host from the source XIV Storage System XIV_02_1310114.

Example 2-8 shows the output of a **vol_list** command for the source and destination XIV systems before you start IBM Hyper-Scale Mobility.

Example 2-8 The vol_list command by pool

```
#Source:
XIV_02_1310114>>vol_list pool=ITSO_3
Name               Size (GB)   Master Name   Consistency Group   Pool   Creator   Used Capacity (GB)
ITSO_Blade5_Perf   2013                                          ITSO_3 itso      0
HSM_001            103                                           ITSO_3 itso      0

#Destination:
XIV_PFE2_1340010>>vol_list pool=ITSO_3
Name      Size (GB)   Master Name   Consistency Group   Pool   Creator   Used Capacity (GB)
HSM_002   103                                           ITSO_3           96
```

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 2-9. The required parameters are **vol=**, **remote_pool=**, and **target=**.

> **Attention:** Any existing snapshot for the IBM Hyper-Scale Mobility source volume will be *deleted without explicit warning* from XCLI.

Example 2-9 The olvm_create command

```
#Source:
XIV_02_1310114>>olvm_create vol=HSM_001 remote_pool=ITSO_3 target=XIV_PFE2_1340010
Command executed successfully.
XIV_02_1310114>>olvm_list
Volume name   Role   Remote System      Active   Phase       State          Link Up
HSM_001       source XIV_PFE2_1340010   no       migration   Initializing   yes

#Destination:
XIV_PFE2_1340010>>olvm_list
Volume name   Role          Remote System   Active   Phase       State          Link Up
HSM_001       destination   XIV_02_1310114  no       migration   Initializing   yes
```

2. As seen in Example 2-9, 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 an OLVM relationship between the source XIV Storage System and the destination XIV Storage System

   The **olvm_list** 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 2-10.

Example 2-10 The olvm_activate command

```
#Source:

XIV_02_1310114>>olvm_activate vol=HSM_001
Command executed successfully.
XIV_02_1310114>>olvm_list
Volume name   Role   Remote System      Active   Phase       State          Link Up
HSM_001       source XIV_PFE2_1340010   yes      migration   Initializing   yes

#After initialization is complete:

XIV_02_1310114>>olvm_list
Volume name   Role   Remote System      Active   Phase   State          Link Up
HSM_001       source XIV_PFE2_1340010   yes      ready   Synchronized   yes

#Destination:
```

```
XIV_PFE2_1340010>>olvm_list
Volume name   Role          Remote System   Active   Phase       State          Link Up
HSM_001       destination   XIV_02_1310114  yes      migration   Initializing   yes

#After initialization is complete:

XIV_PFE2_1340010>>olvm_list
Volume name   Role          Remote System   Active   Phase   State        Link Up
HSM_001       destination   XIV_02_1310114  yes      ready   Consistent   yes
```

2. The **olvm_activate** command initializes the synchronization of the source and destination volumes. Note that the State 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.

Task 4: Proxying

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

1. After the synchronization is complete, the Phase changes to Proxy_Ready and the Status changes to Synchronized, as shown in Example 2-10 on page 71.

   While the migration takes place, you can create the host and host port definitions by running the commands shown in Example 2-11. Note that these definitions can be done before, and independently of, the migration process.

Example 2-11 The host_define, host_add_port, and map_vol commands

```
#Destination:
XIV_PFE2_1340010>>host_define host=ITSO_LPAR1
Command executed successfully.
XIV_PFE2_1340010>>host_add_port host=ITSO_LPAR1 fcaddress=10000000C99987D8
Command executed successfully.
XIV_PFE2_1340010>>host_add_port host=ITSO_LPAR1 fcaddress=10000000C99987D9
Command executed successfully.
```

2. Run the XCLI **olvm_proxy** command to start the Proxy phase, as shown in Example 2-12 on page 73.

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 XIV Storage System, and any host I/O is still synchronized between both source and destination volumes.

Answering `y` to confirm the command execution shown in Example 2-12 triggers the source XIV Storage System to proxy all subsequent host I/O to the volume on the destination XIV Storage System only, and the source volume becomes inconsistent and unavailable. In fact, the source volume is now just a logical placeholder for the original volume. It size is 0 GB, and its hard (total usable) capacity is returned to the storage pool.

Example 2-12 The olvm_proxy command

```
#Source:

XIV_02_1310114>>olvm_proxy vol=HSM_001

Warning:   ARE_YOU_SURE_YOU_WANT_TO_OLVM_PROXY y/n:
Command executed successfully.
XIV_02_1310114>>olvm_list
Volume name   Role   Remote System      Active   Phase   State   Link Up
HSM_001       source XIV_PFE2_1340010   yes      proxy   Proxy   yes

#Destination:

XIV_PFE2_1340010>>olvm_list
Volume name   Role          Remote System   Active   Phase   State   Link Up
HSM_001       destination   XIV_02_1310114  yes      proxy   Proxied yes
```

3. You can verify the start of the Proxy phase by running the **olvm_list** command, as shown in Example 2-12. 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.
4. The Proxy phase is now successfully initiated, and you are ready to map the destination volume to the host by running the **map_vol** command, as shown in Example 2-13.

Example 2-13 The map_vol command

```
#Destination:

XIV_PFE2_1340010>>map_vol  vol=HSM_001 host=ITSO_LPAR1 lun=2
Command executed successfully.
```

5. After mapping, rescan the devices from the host to pick up the new paths to the volume on the destination XIV Storage System. Run the **xiv_fc_admin -R** command, and then run the **xiv_devlist** command from the host, as shown in Example 2-14.

Example 2-14 The xiv_fc_admin -R and xiv_devlist commands

```
[p7-770-02v13:root:/home/root:] xiv_fc_admin -R
[p7-770-02v13:root:/home/root:] xiv_devlist

XIV Devices
------------------------------------------------------------------------
Device       Size (GB)  Paths  Vol Name  Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------
/dev/hdisk1  103.2      12/12  HSM_001   121     1310114  ITSO_LPAR1
------------------------------------------------------------------------
/dev/hdisk2  103.2      6/6    HSM_002   122     1340010  ITSO_LPAR1
------------------------------------------------------------------------

Non-XIV Devices
----------------------------
Device       Size (GB)  Paths
----------------------------
/dev/hdisk0  54.8       N/A
----------------------------
```

Note the 12 paths (six additional) for the destination volume, HSM_001, in the bold row.

Task 5: Cleanup

To complete the Cleanup task, perform the following steps:

1. Now that you have validated that the host has connectivity to the destination volume through the new paths to the destination XIV Storage System, you are ready to unmap the source volume on the source XIV Storage System from the host, as shown in Example 2-15.

Example 2-15 The unmap_vol command

```
#Source:
XIV_02_1310114>>unmap_vol vol=HSM_001 host=ITSO_LPAR1
Command executed successfully.
XIV_02_1310114>>olvm_list
Volume name   Role   Remote System      Active   Phase   State   Link Up
HSM_001       source XIV_PFE2_1340010   yes      proxy   Proxy   yes

#Destination:
XIV_PFE2_1340010>>olvm_list
Volume name   Role          Remote System   Active   Phase   State   Link Up
HSM_001       destination   XIV_02_1310114  yes      proxy   Proxied yes
```

2. Now that you have unmapped the volume from the host, you have only six enabled paths to the volume remaining from the host. Run the **xiv_fc_admin -R** command, and then run the **xiv_devlist** command from the host, as shown in Example 2-16.

Example 2-16 Half the paths enabled

```
[p7-770-02v13:root:/home/root:] xiv_fc_admin -R
[p7-770-02v13:root:/home/root:] xiv_devlist
XIV Devices
------------------------------------------------------------------------
Device       Size (GB)  Paths  Vol Name  Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------
/dev/hdisk1  103.2      6/12   HSM_001   121     1340010  ITSO_LPAR1
------------------------------------------------------------------------
/dev/hdisk2  103.2      6/6    HSM_002   122     1340010  ITSO_LPAR1
------------------------------------------------------------------------
Non-XIV Devices
-----------------------------
Device       Size (GB)  Paths
-----------------------------
/dev/hdisk0  54.8       N/A
-----------------------------
```

Note the two paths for the destination volume, HSM_001, in the bold row. Also, note the Vol ID is still 121 and the XIV ID has changed to 1340010, which indicates that the host is connected *only* to the destination volume on the destination XIV Storage System. The destination XIV Storage System keeps the same Vol ID, if possible.

3. Remove the failed paths in AIX. A sample script for removal of the failed paths is shown Example 2-4 on page 66.

After the removal, six of six paths are available, as shown in Example 2-17.

Example 2-17 Failed paths removed

```
[p7-770-02v13:root:/home/root:] xiv_devlist
XIV Devices
------------------------------------------------------------------------
Device       Size (GB)  Paths  Vol Name  Vol ID  XIV ID   XIV Host
------------------------------------------------------------------------
/dev/hdisk1  103.2      6/6    HSM_001   121     1340010  ITSO_LPAR1
------------------------------------------------------------------------
/dev/hdisk2  103.2      6/6    HSM_002   122     1340010  ITSO_LPAR1
------------------------------------------------------------------------
Non-XIV Devices
---------------------------
Device       Size (GB)  Paths
---------------------------
/dev/hdisk0  54.8       N/A
---------------------------
```

Task 6: Post-cleanup

You have successfully used IBM Hyper-Scale Mobility to move the HSM_001 volume from XIV Storage System XIV_02_1310114 to XIV Storage System XIV_PFE2_1340010, without incurring any downtime on the host. The final step is to end the proxy and delete the relationship:

1. Run the **olvm_delete** command, as shown in Example 2-18.

Example 2-18 The olvm_delete command

```
#Source:
XIV_02_1310114>>olvm_delete vol=HSM_001
Warning:  ARE_YOU_SURE_YOU_WANT_TO_DELETE_OLVM_RELATIONSHIP_IN_THIS_PHASE y/n:
Command executed successfully.
XIV_02_1310114>>olvm_list
No olvms match the given criteria
#Destination:
XIV_PFE2_1340010>>olvm_list
No olvms match the given criteria
```

2. The proxy was removed, and you can confirm the expected end state of the volumes on both the source and destination XIV Storage Systems by running **vol_list** commands from each system. Example 2-19 shows the output of the **vol_list** commands that were run on the source and destination XIV systems after the migration is complete.

Example 2-19 The vol_list by pool command

```
#Source:
XIV_02_1310114>>vol_list pool=ITSO_3
Name               Size (GB)   Master Name   Consistency Group   Pool   Creator   Used Capacity (GB)
ITSO_Blade5_Perf   2013                                          ITSO_3 itso      0

#Destination:
XIV_PFE2_1340010>>vol_list pool=ITSO_3
Name      Size (GB)   Master Name   Consistency Group   Pool   Creator   Used Capacity (GB)
HSM_002   103                                           ITSO_3           96
HSM_001   103                                           ITSO_3           48
```

3

IBM Hyper-Scale Consistency

IBM Hyper-Scale Consistency was introduced in version 11.4 of the IBM XIV Storage System Software, and enables crash-consistent snapshots to be taken across multiple systems.

IBM Hyper-Scale Consistency enables a client to establish a volume consistency group spanning multiple XIV systems. Such a consistency group is called a cross-system Consistency Group (XCG). Hyper-Scale Consistency enables administrators to use XCGs to coordinate the creation of snapshots across XIV systems.

While you previously had to use XCLI commands to facilitate the creation of coordinated snapshots among multiple systems, XIV software release 11.5 with XIV Management Tools version 4.4, offers a fully automated solution. The solution requires Hyper-Scale Manager, which is now included in XIV Management Tools v4.4.

3.1 The business case for IBM Hyper-Scale Consistency

Managing multiple volumes as one consistent group in a single system is common, especially for database applications where tables and logs are on different LUNs. However, there are cases where storage requirements go beyond a single system. As a result, the need for cross-system snapshot consistency arises.

3.1.1 Data consistency

This section provides information about data consistency:

- Data is said to be *crash-consistent* when it is comparable to the state of data when a host crashes. There could be valid application data remaining on the host cache that was not written because the host or application crashed. Because block storage is not aware of information that is still on the host, when a snapshot is taken without some form of interaction with that host, it is called crash-consistent.
- Conversely, data is said to be *application-consistent* when some external process ensures that all data is flushed from the host cache to the volumes. This enables the storage to store exactly what the application has written, without the possibility of data remaining in cache.
- When all volumes are not quiesced (paused) simultaneously, a write-dependent crash consistency issue arises. Taking snapshots of all of the volumes when they are all paused does ensure that dependent writes are ordered. Therefore, a given piece of data appears on the storage if, and only if, the data on which it is dependent (even if it was written to a different volume) is also on the storage.

  The write-dependent crash consistency limiting qualification implies that more recent writes can be on the storage, although earlier writes written to different volumes might be missing because the respective volume's I/O was paused later.
- Within a single XIV, the I/O is paused for all volumes in a consistency group (CG) simultaneously when a CG snapshot is taken. Therefore, crash consistency is virtually assured.

3.1.2 Load balancing across multiple storage systems

Spreading the data across multiple systems can sometimes help in load balancing. When the application's volumes are spread across multiple systems, there is a need for crash consistency across those systems.

3.1.3 Provisioning space bigger than is possible by one system

Sometimes an application's required space is spread across multiple systems. Therefore, all volumes across the systems must be managed as one logical unit (LU). Therefore, the cross-system crash consistency is needed.

3.2 IBM Hyper-Scale Consistency implementation

The IBM Hyper-Scale Consistency extends the individual system crash consistency, by allowing an administrator to establish a volume consistency group spanning multiple XIV systems (an XCG).

3.3 Creating a cross-system consistency group by using GUI

There are two methods of creating a cross-system consistency group:

- Create the cross-consistency group with selected consistency groups at once
- Create the cross-consistency group and subsequently add the consistency groups

3.3.1 Create the cross-consistency group with consistency groups at once

Using the XIV GUI 4.4 or higher version in Hyper-Scale Manager mode, select **Volumes**, then from the pop-up menu, select **X-Consistency Group** as shown in Figure 3-1.

Figure 3-1 Selecting X-Consistency Groups

The cross system Consistency Groups windows open. Select existing consistency groups from unassigned consistency group sets that need to be added to the cross consistency group. To select multiple consistency groups, hold down the Ctrl key to select/deselect the appropriate ones from the various XIV systems. After the consistency groups are selected, right-click any selected CG to open an operations menu. From the operations menu, select **Create a XCG with selected CGs** as shown in Figure 3-2.

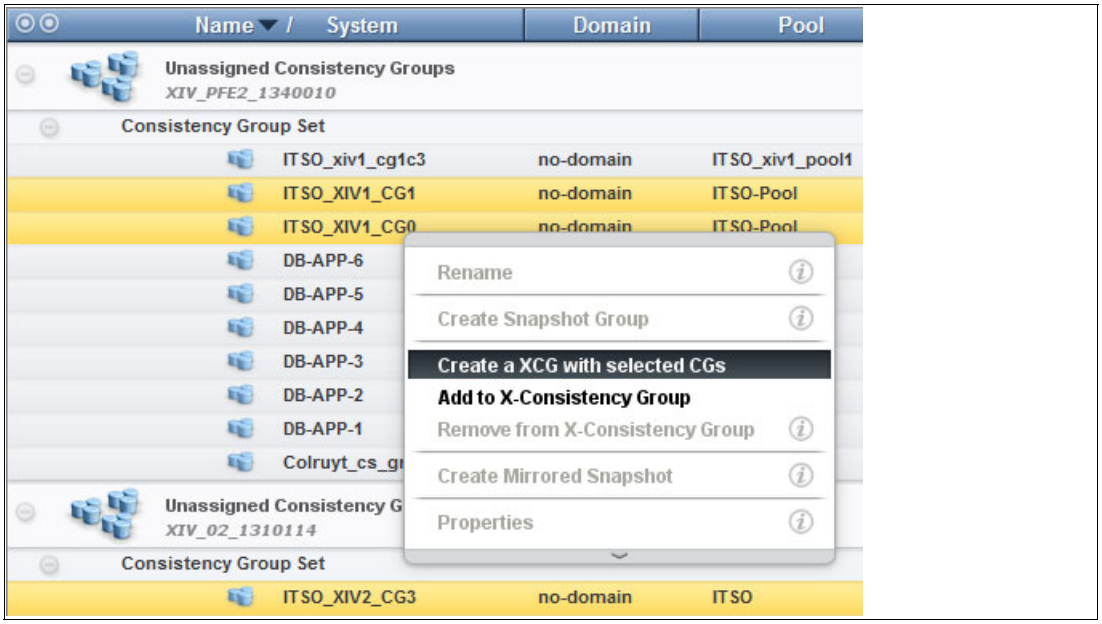

Figure 3-2 Create an XCG with selected CGs

A dialog window opens. Enter a name for the cross consistency group and click Enter as shown in Figure 3-3. Note that the name must be unique (otherwise said, it must not exist in the system yet).

Figure 3-3 Dialog window for creating cross consistency group

After cross consistency group creation, you can check details about its associated consistency groups, by selecting the newly created XCG (ITSO_XCG1 in our example) in the Cross-Consistency Groups view, as shown in Figure 3-4.

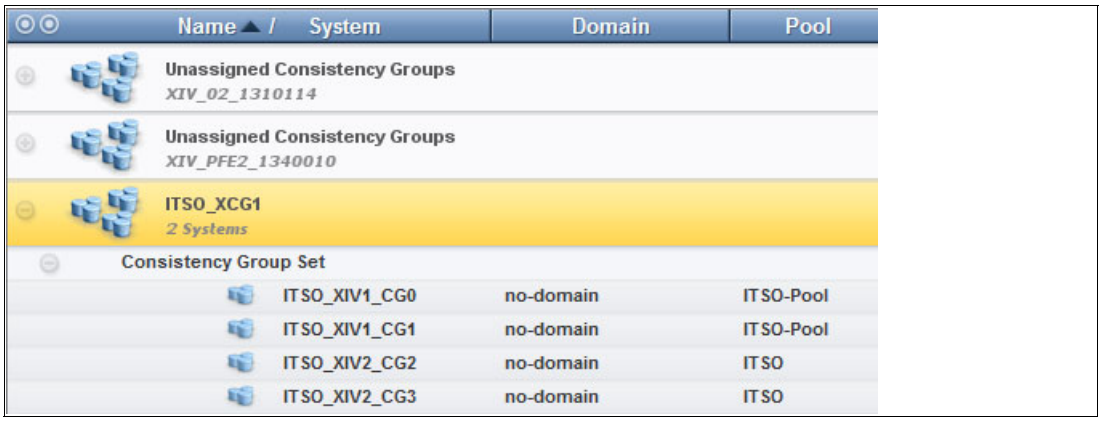

Figure 3-4 List XCG and its related consistency groups

3.3.2 Create the cross-consistency group and subsequently add the consistency groups

From the **Volumes** → **X-Consistency Groups** panel, you can create a cross-consistency group (XCG) without initially adding consistency groups (CG) to it. In Figure 3-5, you can see a Create X-Consistency Group icon in the menu bar at the top of the window.

Click the **Create X-Consistency Group** icon. A creation dialog box opens, as shown in Figure 3-3. Enter a unique name for a new X-Consistency Group.

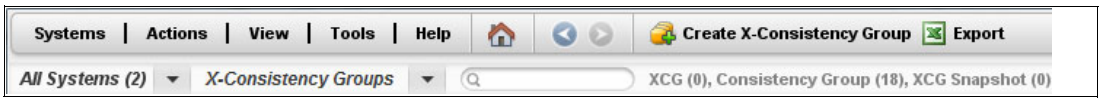

Figure 3-5 Create a new XCG without adding CGs

When created, the XCG appears in the X-Consistency Groups view. The new XCG does not have any CGs associated as illustrated in Figure 3-6 on page 81.

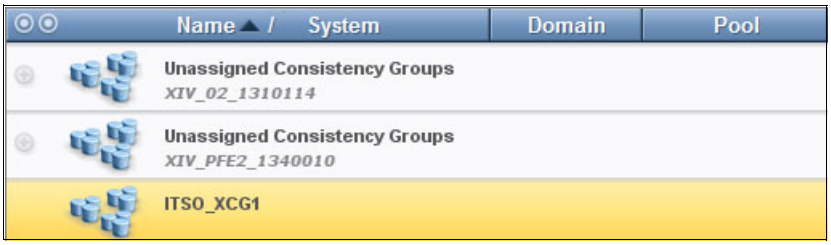

Figure 3-6 Validate the new XCG

Using the same **Volumes** → **X-Consistency Groups** view in the GUI, select the consistency groups to add to the XCG. After selecting the wanted consistency groups, right-click any of them and select **Add to X-Consistency Group** from the pop-up action menu. Figure 3-7 shows an illustration of four highlighted consistency groups cross in two XIV systems being added to the XCG.

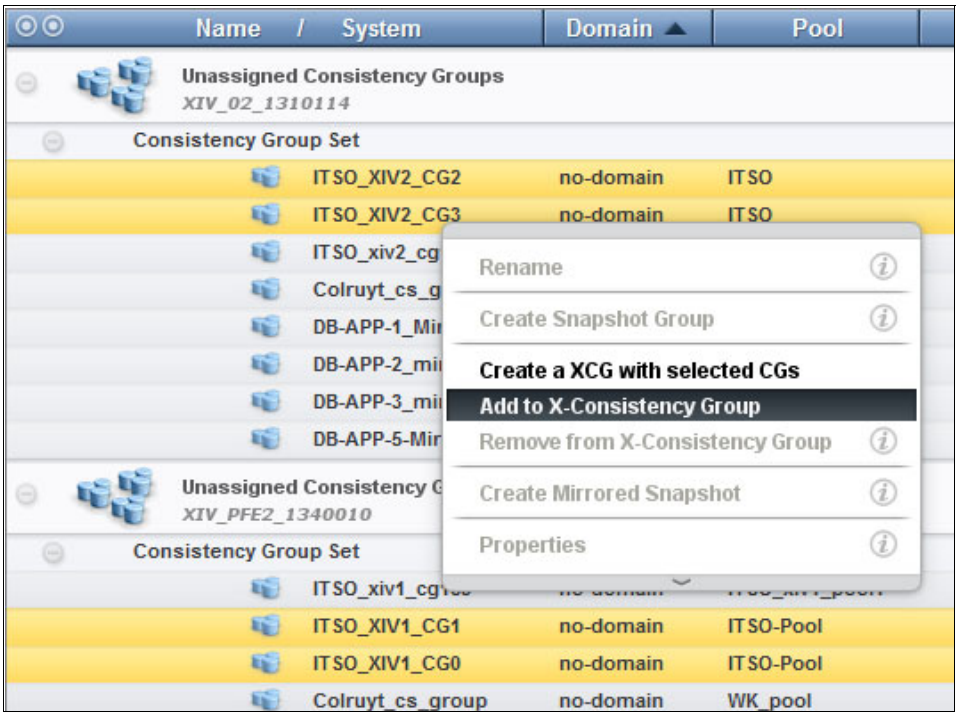

Figure 3-7 Adding CGs to XCG

A dialog window, which is shown in Figure 3-8, prompts you for the name of an existing XCG. In our illustration, we select the ITSO_XCG1 cross consistency group that was previously created. Click **Add** to complete the operation.

Figure 3-8 Selecting an XCG for adding consistency groups

3.3.3 Creating an XCG snapshot group

When the cross consistency group (XCG) is created and the consistency groups (CG) added, you can create an XCG snapshot group. As shown in Figure 3-9, right-click the appropriate XCG and select **Create XCG Snapshot Group** from the pop-up menu. The system immediately creates a snapshot group.

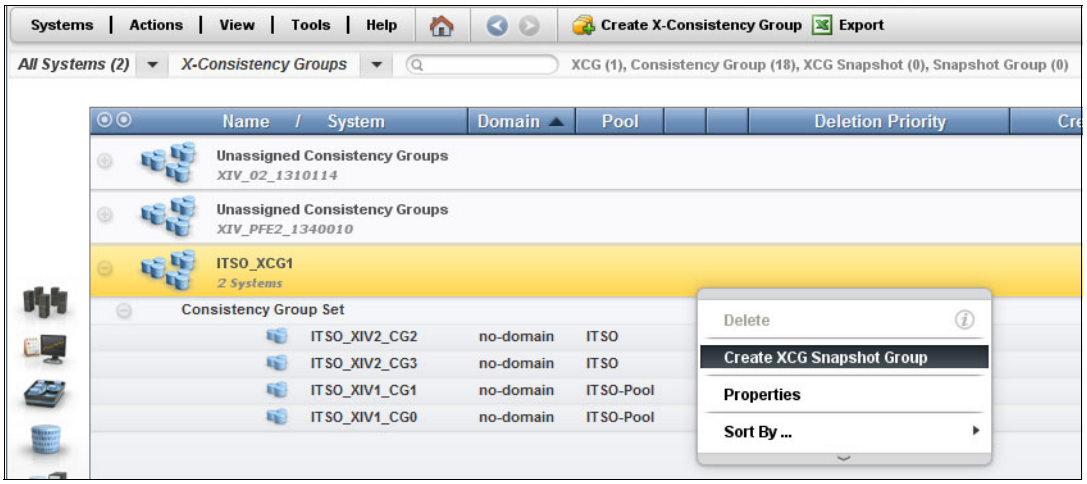

Figure 3-9 Creating a snapshot for XCG

The new snapshots are created and displayed underneath the Consistency Group Set in the **X-Consistency Groups** view (Figure 3-10). These snapshots all have the same creation date and time. To make sure that all systems show the same time, it is best for time synchronization to configure a Network Time Protocol (NTP) server as part of your XIV system settings.

Each snapshot is locked upon creation and has the same default characteristics as a regular snapshot. The snapshots are contained in a group structure (called a sn*apshot group*) that allows all the snapshots to be managed by a single operation such as delete, rename, overwrite, duplicate, and unlock.

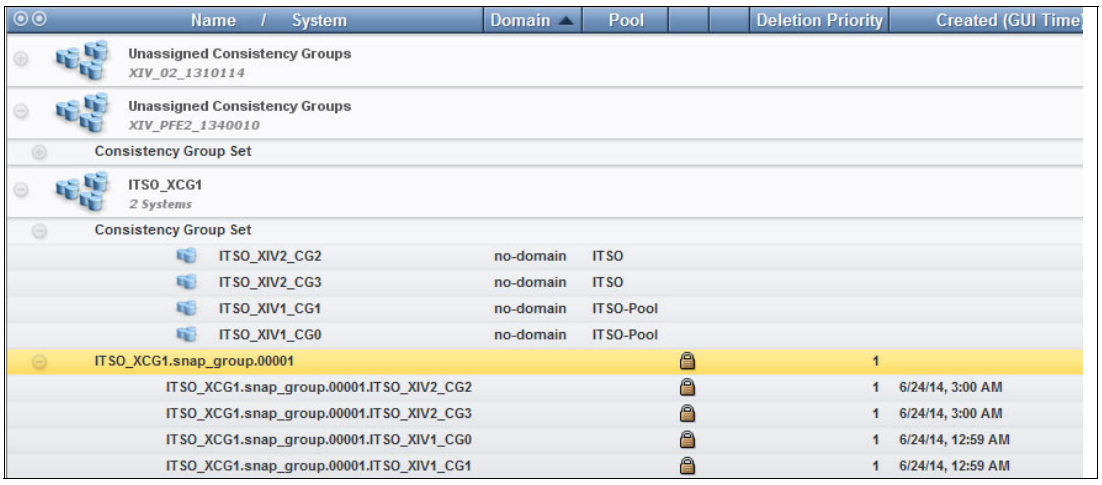

Figure 3-10 Validating the new snapshots in the XCG

The cross-system consistency group does not prevent you from creating a single consistency group snapshot. If the single consistency group snapshot is created, it is not displayed in the XCG view. The system takes you automatically to the Consistency Group view for details.

3.3.4 Managing a cross consistency group

When snapshots have been created within a cross-system consistency group, you can take various actions on the relevant snapshots. Those actions are the same ones as apply to regular volume snapshots. The deletion priority is modifiable, and the snapshot group can be unlocked and locked, as well as it can be restored or overwritten. For specific details about performing these operations, refer to the IBM Redbooks publication, *XIV Storage System Business Continuity Functions*, SG24-7759.

In addition to the snapshot functions, you can remove a consistency group from the XCG. To remove a consistency group from the XCG, right-click the individual consistency group. From the pop-up menu, select **Remove From X-Consistency Group** and confirm the removal. Figure 3-11 provides an illustration of removing the `ITSO_XIV2_CG2` consistency group from the XCG.

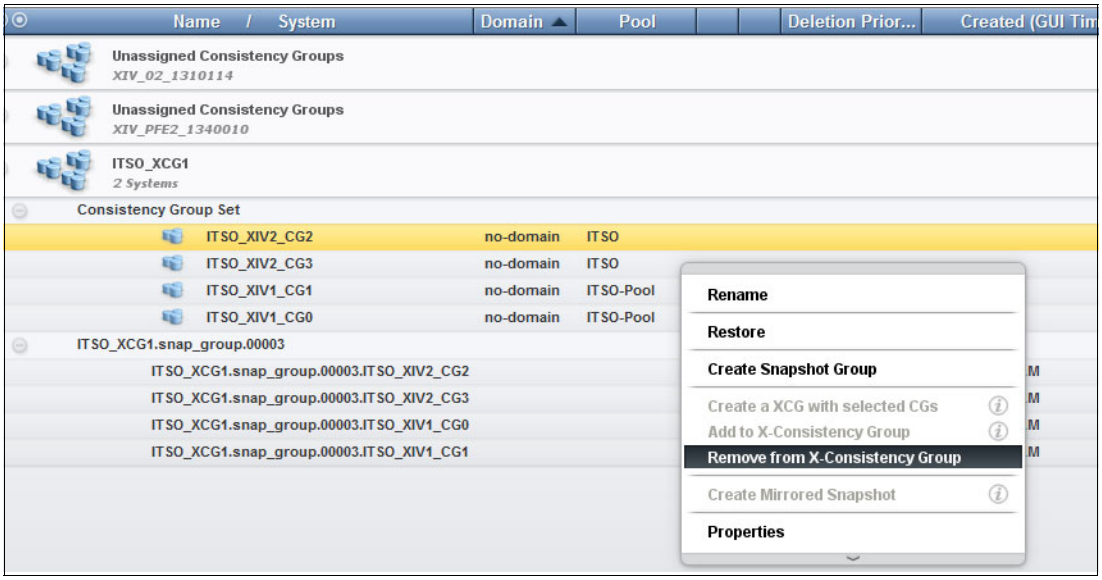

Figure 3-11 Removing consistency group from XCG

The **Snapshots Group Tree** view from the **Volumes** function menu is not available (it works for single system only).

You can navigate to the Consistency Group view to see the removed consistency group and its related snapshot details. The groups can be unlocked or locked, restored, or overwritten. All the operations available for simple snapshots are available with the **`snap_group`** operations.

3.3.5 Deleting a cross consistency group

Before a cross consistency group (XCG) can be deleted, the associated consistency groups must be removed from the XCG as shown in Figure 3-12.

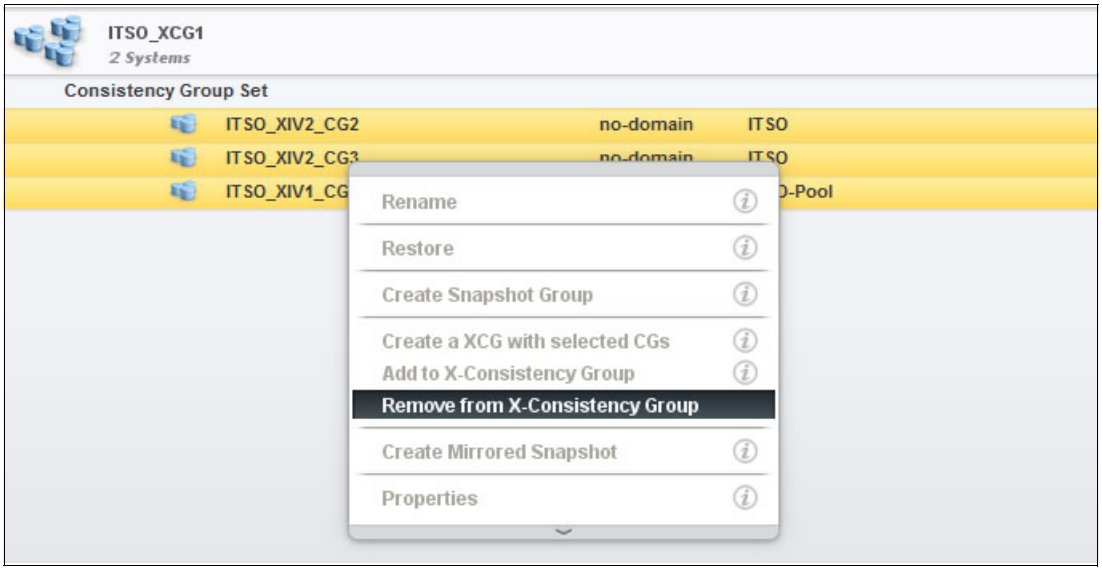

Figure 3-12 Remove consistency groups from XCG

On deletion of a cross consistency group, the snapshots become independent snapshots and remain tied to their respective consistency group. To delete the cross consistency group, right-click the group and select **Delete**. Validate the operation by clicking **OK**. Figure 3-13 provides an example of deleting the cross consistency group called `ITSO_XCG1`.

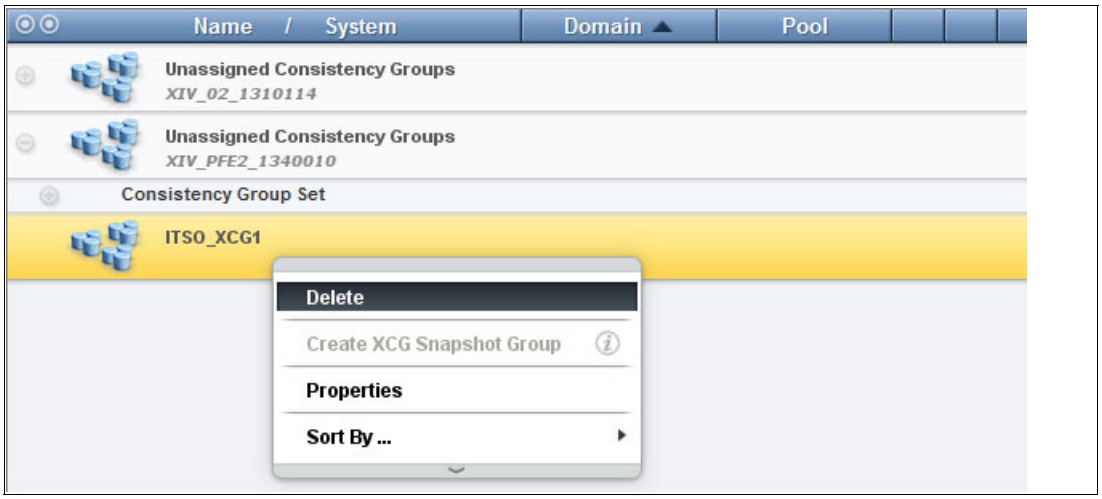

Figure 3-13 Deleting a cross consistency group

3.4 Managing the IBM Hyper-Scale Consistency with XCLI

Hyper-Scale Consistency can be managed through a set of XCLI commands against the XCG object, including the ability to stop and resume I/O to an XCG.

Important: Using the XIV GUI to manage XCG is the preferred method. As explained in 3.3, “Creating a cross-system consistency group by using GUI” on page 79, the GUI automates operations (XCLI commands) that would have to be manually executed or combined in a script, when using the XCLI.

When using the XCLI, to have crash-consistent data, all I/O to all relevant volumes (within their respective CGs) have to be paused.

The XCLI **io_pause** command is used to that end. The command is run on each of the XIV systems. The I/O is paused on each CG for at most 10 seconds. Within the configured pause time window, all consistency groups across all systems have to be paused. Having all I/O paused on all systems almost assures the consistency. However, there can be a small window where host writes could still get through before the I/O of a given CG is paused.

After all relevant CGs are paused, their respective snapshots are taken. All CGs’ snapshots together make the *one* cross-system consistency snapshot.

After the creation of the snapshots, the paused I/Os are resumed. There are three ways to resume the I/Os:

- Let the **io_pause** commands expire (at most, timeout is 10 seconds, default is 5 seconds).
- Follow the **cg_snapshot _create** commands with the **io_resume** command.
- Calling the **cg_snapshot_create** command with the **auto_resume=yes** parameter.

The following XCLI commands are used to manage the IBM Hyper-Scale Consistency:

- The **io_pause** command
- The **io_resume** command
- The **io_pause_list** command
- The **cg_snapshots_create** command

The following XCLI commands provide additional flexibility in managing the consistency group:

xcg_create	Create a cross_system_cg object.
xcg_add_CG	Add a local CG to the cross_system_cg.
xcg_remove_CG	Remove a local CG to the cross_system_cg.
xcg_get_local_CGs	Retrieve the CGs consisting of XCG.
xcg_add_remote_system	Add a remote system name to the cross_system_cg.
xcg_remove_remote_system	Remove a remote system name to the cross_system_cg.
xcg_get_remote_systems	Retrieve the remote systems name consisting of the Cross-system Consistency Group (XCG).
xcg_delete	Delete a cross_system_cg object.

We illustrate below how to take cross consistency group snapshots using the **auto_resume** parameter. We show all the setup steps, starting from the XCG creation but assuming that the selected CGs already exist on the different XIV systems.

Example 3-1 illustrates the commands that are required for creating an XCG and adding consistency groups and remote systems.

Example 3-1 Creating XCG and adding consistency groups with the XCLI

```
XIV_PFE2_1340010>>
xcg_create xcg=ITSO_XCG1
xcg_add_cg xcg=ITSO_XCG1 cg=ITSO_XIV1_CG0
xcg_add_cg xcg=ITSO_XCG1 cg=ITSO_XIV1_CG1
xcg_add_remote_system xcg=ITSO_XCG1 remote_system="XIV_02_1310114"
XIV_02_1310114>>
xcg_create xcg=ITSO_XCG1
xcg_add_cg xcg=ITSO_XCG1 cg=ITSO_XIV2_CG2
xcg_add_cg xcg=ITSO_XCG1 cg=ITSO_XIV2_CG3
xcg_add_remote_system xcg=ITSO_XCG1 remote_system="XIV_PFE2_1340010"
```

To display the XCG in the system, execute the XCLI **xcg_list** command as illustrated in Example 3-2.

Example 3-2 Lists XCGs

```
XIV_PFE2_1340010>>xcg_list
Name        Num Of CGs   Num Of Remote Systems
ITSO_XCG1   2            1
XIV_02_1310114>>xcg_list
Name        Num Of CGs   Num Of Remote Systems
ITSO_XCG1   2            1
```

Example 3-3 shows the snapshot group creation for XCG, step-by-step using XCLI commands. IOs to volumes in the consistency groups must be stopped for a short time period (**mili_seconds_to_resume** parameter specifies a time-out value representing the duration in millisec format). After the **io_pause** command completes, it returns a token ID, which serves as request identifier. The token ID is used to perform **io_resume** or used with the **auto_resume** parameter at snapshot group creation time. The Interface node collects all IOs that were blocked during the io_pause period and processes the collected IOs after io_pause is released.

Example 3-3 Creating a snapshot group for XCG

```
XIV_PFE2_1340010>>
io_pause milli_seconds_to_resume="10000" cg="ITSO_XIV1_CG0"
Command executed successfully.
    token_id=1363313159045129
io_pause milli_seconds_to_resume="10000" cg="ITSO_XIV1_CG1”
Command executed successfully.
    token_id=1361698251341832
cg_snapshots_create snap_group="ITSO_XCG1.snap_group.00001.ITSO_XIV1_CG0"
auto_resume="1363313159045129" cg="ITSO_XIV1_CG0"
cg_snapshots_create snap_group="ITSO_XCG1.snap_group.00001.ITSO_XIV1_CG1"
auto_resume="1361698251341832" cg="ITSO_XIV1_CG1"

XIV_02_1310114>>
io_pause milli_seconds_to_resume="10000" cg="ITSO_XIV2_CG2"
```

```
Command executed successfully.
    token_id=1264816674045958
io_pause milli_seconds_to_resume="10000" cg="ITSO_XIV2_CG3"
Command executed successfully.
    token_id=1265100141887495
cg_snapshots_create snap_group="ITSO_XCG1.snap_group.00001.ITSO_XIV2_CG2"
auto_resume="1264816674045958" cg="ITSO_XIV2_CG2"
cg_snapshots_create snap_group="ITSO_XCG1.snap_group.00001.ITSO_XIV2_CG3"
auto_resume="1265100141887495" cg="ITSO_XIV2_CG3"
```

Example 3-4 illustrates the **snap_group_list** command.

Example 3-4 Listing all consistency groups and their related snapshots

```
XIV_PFE2_1340010>>snap_group_list
Name                                      CG               Snapshot Time         Deletion Priority
ITSO_xiv1_cg1c3.mirror_snapshot_          ITSO_xiv1_cg1c3  2014-06-20 15:29:26   1
ITSO_XCG1.snap_group.00001.ITSO_XIV1_CG0  ITSO_XIV1_CG0    2014-06-24 01:43:57   1
ITSO_XCG1.snap_group.00002.ITSO_XIV1_CG0  ITSO_XIV1_CG0    2014-06-24 02:40:03   1
ITSO_XCG1.snap_group.00003.ITSO_XIV1_CG0  ITSO_XIV1_CG0    2014-06-24 02:42:52   1
ITSO_XCG1.snap_group.00003.ITSO_XIV1_CG1  ITSO_XIV1_CG1    2014-06-24 02:42:52   1
```

By using the **auto_resume** parameter, the I/O can be resumed faster than when the **io_resume** command is used. Also, when the **auto_resume** parameter is used, the cg_snapshots will be taken only if the I/O is paused. This is as opposed to the case when the **auto_resume** parameter is not used, in which the snapshots are created regardless of whether the I/Os to the volumes in the CG are paused or not.

At this point, all work is back to normal, and the snapshots can be backed up. Because the snapshots are crash-consistent as opposed to application-consistent, there might be a need to carry out application-level recovery procedures when the snapshots are restored.

Once the XCG and its snapshots are not relevant anymore, you can delete the cross consistency group. First, you must remove all the consistency groups one at a time as well as remove the remote system association. In Example 3-5, each consistency group in the XCG is removed, then the cross consistency group is available for deletion. Deletion of the XCG does not delete the individual snapshots. They remain tied to their corresponding consistency group and are only removed from the consistency group when you remove the volumes.

Example 3-5 Deleting cross consistency group and related snapshots, CGs with XCLI

```
XIV_PFE2_1340010>>
xcg_remove_cg xcg=ITSO_XCG1 cg=ITSO_XIV1_CG0
Warning:   Are you sure you want to remove cons group 'ITSO_XIV1_CG0' from its cross
Consistency Group? y/n: y
Command executed successfully.
xcg_remove_cg xcg=ITSO_XCG1 cg=ITSO_XIV1_CG1
Warning:   Are you sure you want to remove cons group 'ITSO_XIV1_CG1' from its cross
Consistency Group? y/n: y
Command executed successfully.

xcg_remove_remote_system xcg=ITSO_XCG1 remote_system=XIV_02_1310114
Command executed successfully.

xcg_delete xcg=ITSO_XCG1
Command executed successfully.
```

```
XIV_02_1310114>>
xcg_remove_cg xcg=ITSO_XCG1 cg=ITSO_XIV2_CG2

Warning:   Are you sure you want to remove cons group 'ITSO_XIV2_CG2' from its
cross Consistency Group? y/n: y
Command executed successfully.
xcg_remove_cg xcg=ITSO_XCG1 cg=ITSO_XIV2_CG3

Warning:   Are you sure you want to remove cons group 'ITSO_XIV2_CG3' from its
cross Consistency Group? y/n: y
Command executed successfully.

xcg_remove_remote_system xcg=ITSO_XCG1 remote_system="XIV_PFE2_1340010"
Command executed successfully.

xcg_delete xcg=ITSO_XCG1
Command executed successfully.
```

3.4.1 IBM Hyper-Scale Consistency additional considerations

The following list highlights the current technology boundaries of the IBM Hyper-Scale Consistency:

- The **io_pause** command returns a tokenId. The tokenId is used to eliminate scenarios where the cg_name has been changed or deleted and created again while the **io_pause** command was in effect. The tokenId identifies the CG and not the specific invocation.
- The CG is considered paused until an **io_resume** command is received, or until its timeout has expired.
- Multiple simultaneous **io_pause** commands are enabled to non-overlapping CGs.
- When a CG is already in a paused state from a previous **io_pause** command, no action is performed by the **io_pause** command. The time remaining from the previously issued **io_pause** commands milliseconds_to_resume value is *not* adjusted.
- It is not possible to increase the amount of time that the CG is in the paused state by issuing **io_pause** commands sequentially within the timeout period.
- During a hot upgrade, all paused CGs are resumed automatically because of the time constraints and the need to finish a hot upgrade as soon as possible.
- During an upgrade, the client cannot take snapshots because management operations are restricted at that time. Therefore, the client is not shown inconsistent snapshots taken due to the resume operation.

Related publications

The publications listed in this section are considered particularly suitable for providing more detailed information about the topics covered in this paper.

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 XIV Storage System Multi-site Mirroring*, REDP-5129
- *IBM Hyper-Scale in XIV Storage*, REDP-5053
- *IBM i and IBM System Storage: A Guide to Implementing External Disks on IBM i*, SG24-7120
- *IBM XIV Storage System Gen3: Architecture, Implementation, and Usage*, SG24-7659
- *IBM XIV Storage System with the Virtual I/O Server and IBM i,* REDP-4598
- *RESTful API Support in IBM XIV*, REDP-5064
- *Using the IBM XIV Storage System in OpenStack Cloud Environments*, REDP-4971
- *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904
- *Solid-State Drive Caching in the IBM XIV Storage System*, REDP-4842
- *XIV Storage System in a VMware Environment*, REDP-4965
- *XIV Security With Data-At-Rest Encryption*, REDP-5047
- *XIV Gen3 with IBM System Storage SAN Volume Controller and Storwize V7000*, REDP-5063

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

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- *IBM XIV Storage System Application Programming Interface,* GC27-3916
- *IBM XIV Storage System User Manual*, GC27-3914
- *IBM XIV Storage System: Product Overview*, GC27-3912
- *IBM XIV Storage System Planning Guide, SC27-5412*
- *IBM XIV Storage System XCLI Utility User Manual*, GC27-3915
- *IBM XIV Remote Support Proxy Installation and User's Guide*, GA32-0795
- *Management Tools Operations Guide*, SC27-5986

Online resources

These websites are also relevant as further information sources:

- IBM XIV Storage System Knowledge Center:

  http://www-01.ibm.com/support/knowledgecenter/STJTAG/product_welcome/xiv_kcwelcomepage.html

- IBM XIV Storage System website:

  http://www.ibm.com/systems/storage/disk/xiv/index.html

- System Storage Interoperability Center (SSIC):

  http://www.ibm.com/systems/support/storage/config/ssic/index.jsp

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

IBM Hyper-Scale in XIV Storage

Redpaper™

Powerful new approach to storage management scalability

Volume mobility and consistency across XIV systems

Capacity planning and reporting

Organizations are struggling with an unprecedented surge in data, both user-created and machine-created. With an increasing reliance on this rapidly growing data, effective scaling of data storage represents a key challenge for all enterprises.

IBM Hyper-Scale helps you easily overcome provisioning scenarios that normally challenge traditional systems. IBM Hyper-Scale can accommodate several critical client scenarios for centralized administration, data mobility, load balancing, over-provisioning, and storage system repurposing.

IBM Hyper-Scale introduces the following major concepts and technologies, with the IBM XIV Storage System (Gen3 models):

- IBM Hyper-Scale Manager is a flexible, consolidated multi-system management application that builds upon the originally released Multi-System Manager with XIV Storage Software V11.2. IBM Hyper-Scale Manager is based on, and seamlessly integrated with, the XIV graphical user interface (GUI) and spans multiple XIV systems. It virtually transforms multiple systems into a single system, allowing customers to centrally manage up to 144 XIV systems.
- IBM Hyper-Scale Mobility was released with XIV Storage System V11.3. It is a powerful function for moving volumes between storage containers transparently, with no disruption to host applications.
- IBM Hyper-Scale Consistency was released with the IBM XIV Storage Software Version 11.4. It offers cross-system consistency, enabling coordinated snapshots across independent XIV systems. It helps to ensure data protection across multiple XIV systems.

This IBM Redpaper publication provides a broad understanding of the IBM Hyper-Scale feature. This publication is intended for XIV clients and users who want a practical understanding of IBM Hyper-Scale concepts and usage. For information about the IBM Hyper-Scale concepts, refer to the IBM white paper, *IBM Hyper-Scale and Its Implementation in XIV Storage: A powerful new approach to storage management scalability.*

REDP-5053-01 0738453986

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Clients, and IBM Business Partners from around the world create timely technical information based on realistic scenarios. Specific suggestions are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks