

IBM XIV Storage System

Architecture, Implementation, and Usage

Support provided for 2 TB SATA
disk drives

Detailed system configuration
instructions

LDAP Configuration Wizard
Secure Remote Support



Bertrand Dufrasne
Roger Eriksson
Wilhelm Gardt
Nils Nause
Markus Oscheka
Carlo Saba
Eugene Tsylin
Kip Wagner
Axel Westphal
Ralf Wohlfarth

Redbooks



International Technical Support Organization

**IBM XIV Storage System: Architecture,
Implementation, and Usage**

January 2011

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

Third Edition (January 2011)

This edition applies to Version 10.2.2 of the IBM XIV Storage System Software and Version 2.5 of the IBM XIV Storage System Hardware.

© Copyright International Business Machines Corporation 2008, 2009, 2011. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
Summary of changes	ix
January 2011, Third Edition	ix
Preface	xi
The team who wrote this book	xii
Now you can become a published author, too!	xiv
Comments welcome	xv
Stay connected to IBM Redbooks publications	xv
Chapter 1. IBM XIV Storage System overview	1
1.1 Introduction	2
1.2 Total Cost of Ownership (TCO)	2
1.3 System models and components	3
1.4 Key design features	4
1.4.1 Massive parallelism	4
1.4.2 Workload balancing	4
1.4.3 Self-healing	5
1.4.4 True virtualization	5
1.4.5 Thin provisioning	5
1.4.6 Processing power	5
1.5 The XIV Storage System software	5
1.5.1 IBM XIV Storage Management Tools	7
1.5.2 IBM XIV Storage Management GUI	8
1.5.3 IBM XIV Storage Management online monitoring with XIVTop	8
1.5.4 IBM XIV Storage Management XCLI	9
1.6 Host support	10
Chapter 2. IBM XIV logical architecture and concepts	13
2.1 Architecture overview	14
2.1.1 Features	14
2.1.2 Hardware elements	14
2.2 Parallelism	16
2.2.1 Hardware parallelism and grid architecture	16
2.2.2 Software parallelism	17
2.3 Full storage virtualization	18
2.3.1 IBM XIV Storage System virtualization design	18
2.3.2 Logical system concepts	20
2.3.3 System usable capacity	23
2.3.4 Storage pool concepts	24
2.3.5 Capacity allocation and thin provisioning	26
2.4 Reliability, availability, and serviceability	35
2.4.1 Resilient architecture	35
2.4.2 Rebuild and redistribution	38
2.4.3 Minimized exposure	43
Chapter 3. IBM XIV physical architecture, components, and planning	47

3.1 IBM XIV Storage System models 2810-A14 and 2812-A14	48
3.1.1 Fully populated configurations	49
3.1.2 Partially populated configurations	50
3.2 IBM XIV hardware components	50
3.2.1 Redundancy	51
3.2.2 Rack and UPS modules	52
3.2.3 Data Modules and Interface Modules	55
3.2.4 SATA disk drives	61
3.2.5 Patch panel	63
3.2.6 Interconnection and switches	64
3.2.7 Support hardware	64
3.3 Hardware planning overview	66
3.3.1 Ordering IBM XIV Storage System hardware	67
3.3.2 Physical site planning	67
3.3.3 Basic configuration planning	69
3.3.4 IBM XIV physical installation	78
3.3.5 System power-on and power-off	80
Chapter 4. Configuration	83
4.1 IBM XIV Storage Management software	84
4.1.1 XIV Storage Management user interfaces	84
4.1.2 XIV Storage Management software installation	85
4.2 Managing the IBM XIV Storage System	89
4.2.1 The XIV Storage Management GUI	91
4.2.2 Logging on to the system with XCLI	99
4.3 Storage pools	102
4.3.1 Function of storage pools	102
4.3.2 Managing storage pools with the XIV GUI	104
4.3.3 Managing storage pools with XCLI	111
4.4 Volumes	113
4.4.1 Managing volumes with the XIV GUI	114
4.4.2 Managing volumes with XCLI	123
4.5 Host definition and mappings	124
4.5.1 Assigning LUNs to a host using the GUI	124
4.5.2 Assigning LUNs to a host using the XCLI	129
4.6 Scripts	130
Chapter 5. Security	131
5.1 Physical access security	132
5.2 Native user authentication	132
5.2.1 XIV local credential repository	132
5.2.2 Managing user accounts with XIV GUI	136
5.2.3 Managing user accounts using XCLI	143
5.2.4 Password management and resets	145
5.2.5 Managing multiple systems	146
5.3 LDAP managed user authentication	147
5.3.1 Introduction to LDAP	147
5.3.2 LDAP directory components	148
5.3.3 LDAP product selection	149
5.3.4 LDAP login process overview	150
5.3.5 LDAP role mapping	151
5.4 Defining LDAP on the XIV system	155
5.4.1 Using the XIV GUI LDAP Wizard to configure LDAP	156

5.4.2 Using the XIV GUI directly to configure LDAP	162
5.4.3 Using XCLI	165
5.5 LDAP managed user authentication	166
5.5.1 LDAP repository credential objects	167
5.5.2 Managing LDAP user accounts	173
5.5.3 Managing user groups using XCLI in LDAP authentication mode	178
5.5.4 Active Directory group membership and XIV role mapping	179
5.5.5 SUN Java Directory group membership and XIV role mapping	183
5.5.6 Managing multiple systems in LDAP authentication mode	186
5.6 Securing LDAP communication with SSL	187
5.6.1 Configuring XIV to use LDAP over SSL	188
5.6.2 Maintaining the SSL certificate	190
5.7 XIV audit event logging	191
5.7.1 Viewing events in the XIV GUI	191
5.7.2 Event attributes	192
5.7.3 Viewing events in the XCLI	192
5.7.4 Defining notification rules	195
Chapter 6. Performance characteristics	197
6.1 Performance concepts	198
6.1.1 Full disk resource utilization	198
6.1.2 Caching mechanisms	198
6.1.3 Data mirroring	199
6.1.4 Snapshots	199
6.2 Best practices	200
6.2.1 Distribution of connectivity	200
6.2.2 Host configuration considerations	200
6.2.3 XIV sizing validation	202
Chapter 7. Monitoring	203
7.1 System monitoring	204
7.1.1 Monitoring with the GUI	204
7.1.2 Monitoring with XCLI	209
7.1.3 Using Tivoli Storage Productivity Center	215
7.2 Performance monitoring	226
7.2.1 Using the GUI	226
7.2.2 Using XIVTop	231
7.2.3 Using the XCLI	235
7.3 XIV event notification	237
7.3.1 Configuring the XIV system	237
7.3.2 SNMP-based monitoring	247
7.4 Custom Monitoring	261
7.4.1 Custom monitoring of a volume	261
7.4.2 Custom monitoring of performance attributes	264
7.4.3 Conclusion	267
7.5 Call Home and Remote support	267
7.5.1 Call Home	267
7.5.2 Remote support	267
7.5.3 Repair flow	275
Appendix A. Additional LDAP information	279
Creating user accounts in Microsoft Active Directory	280
Creating user accounts in SUN Java Directory	286
Securing LDAP communication with SSL	293

Windows Server SSL configuration.	293
SUN Java Directory SSL configuration.	299
Certificate Authority setup	305
Related publications	309
IBM Redbooks publications	309
Other publications	309
Online resources	309
How to get IBM Redbooks publications	310
Help from IBM	310
Index	311

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 give 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.

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

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	Redbooks®	Tivoli®
DS6000™	Redbooks (logo)  ®	TotalStorage®
DS8000®	System p®	XIV®
IBM®	System Storage™	
NetView®	System x®	

The following terms are trademarks of other companies:

Disk Magic, and the IntelliMagic logo are trademarks of IntelliMagic BV in the United States, other countries, or both.

Snapshot, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel Xeon, Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

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

Summary of changes

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

Summary of Changes
for SG24-7659-02
for IBM XIV Storage System: Architecture, Implementation, and Usage
as created or updated on January 20, 2011.

January 2011, Third Edition

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

New information

- ▶ Support for 2 TB drives. These disk drives are supported on the IBM® XIV® Storage System Model A14 with IBM XIV Storage System Software V10.2.1, or later, installed.
- ▶ Various hardware enhancements
- ▶ New information for integration with TPC 4.2
- ▶ New information for integration with IBM Systems Director
- ▶ XIV Remote Support Center (XRSC)

Changed information

- ▶ LDAP based authentication
- ▶ Various XIV GUI and XCLI changes
- ▶ Custom monitoring
- ▶ Host attachment and interoperability topics are now covered in a separate publication, *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904.

Preface

This IBM Redbooks® publication describes the concepts, architecture, and implementation of the IBM XIV Storage System (2810-A14 and 2812-A14).

The XIV Storage System is designed to be a scalable enterprise storage system that is based upon a grid array of hardware components. It can attach to both Fibre Channel Protocol (FCP) and IP network Small Computer System Interface (iSCSI) capable hosts. This system is a good fit for clients who want to be able to grow capacity without managing multiple tiers of storage. The XIV Storage System is well suited for mixed or random access workloads, such as the processing of transactions, video clips, images, and email, and industries, such as telecommunications, media and entertainment, finance, and pharmaceutical, as well as new and emerging workload areas, such as Web 2.0.

In the first few chapters of this book, we provide details about many of the unique and powerful concepts that form the basis of the IBM XIV Storage System logical and physical architecture. We explain how the system was designed to eliminate direct dependencies between the hardware elements and the software that governs the system.

In subsequent chapters, we explain the planning and preparation tasks that are required to deploy the system in your environment. We present a step-by-step procedure describing how to configure and administer the system. We provide illustrations about how to perform those tasks by using the intuitive, yet powerful XIV Storage Manager GUI or the Extended Command Line Interface (XCLI).

This edition of the book contains new information about how to integrate the IBM XIV Storage System for authentication in an LDAP environment.

We also discuss the performance characteristics of the IBM XIV Storage System and present options available for alerting and monitoring, including an enhanced Secure Remote Support capability.

This book is intended for those people who want an understanding of the IBM XIV Storage System and also targets readers who need detailed advice about how to configure and use the system.

Note that host operating systems and other integration aspects are now addressed in a separate IBM Redbooks publication, *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904.

The team who wrote this book

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

Bertrand Dufrasne is an IBM Certified Consulting I/T Specialist and Project Leader for System Storage™ disk products at the International Technical Support Organization, San Jose Center. He has worked at IBM in various I/T 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 Masters degree in Electrical Engineering from the Polytechnic Faculty of Mons (Belgium).

Roger Eriksson is an STG Lab Services consultant, based in Stockholm, Sweden and working for the European Storage Competence Center in Mainz, Germany. He is a Senior Accredited IBM Product Service Professional. Roger has over 20 years experience working on IBM servers and storage, including Enterprise and Midrange disk, NAS, SAN, System x®, System p®, and Bladecenters. He has been working with consulting, proof of concepts, and education, mainly with the IBM XIV product line since December 2008, working with both clients and various IBM teams worldwide. He holds a Technical College graduate degree in Mechanical Engineering.



Wilhelm Gardt holds a degree in Computer Sciences from the University of Kaiserslautern, Germany. He has worked as a Software Developer and subsequently as an IT Specialist designing and implementing heterogeneous IT environments (SAP®, Oracle®, AIX®, HP-UX, SAN, and others). In 2001 he joined the IBM TotalStorage® Interoperability Centre (now Systems Lab Europe) in Mainz where he performed customer briefings and Proof Of Concepts on IBM storage products. Since September 2004 he has been a member of the Technical Pre-Sales Support team for IBM Storage (Advanced Technical Support).

Nils Nause is a Storage Support Specialist for IBM XIV Storage Systems working at IBM Mainz, Germany. Nils joined IBM in the summer of 2005, responsible for Proof of Concepts and delivering briefings for various IBM products. In July 2008 he started working for the IBM XIV Post Sales Support, with a special focus on Oracle Solaris attachment, as well as overall security aspects of the IBM XIV Storage System. He holds a degree in Computer Science from the University of Applied Science in Wernigerode, Germany.

Markus Oscheka is an IT Specialist for Proof of Concepts and Benchmarks in the Disk Solution Europe team in Mainz, Germany. His areas of expertise include setup and demonstration of IBM System Storage and TotalStorage solutions in various environments such as AIX, Linux®, Windows®, VMware ESX, and Solaris, working at IBM for nine years. He has performed many Proof of Concepts with Copy Services on DS6000/DS8000/XIV, as well as Performance Benchmarks with DS4000/DS6000/DS8000/XIV. He has written extensively in various IBM Redbooks publications and also acts as the co-project lead for these Redbooks publications, including *DS6000/DS8000® Architecture and Implementation*, *DS6000/DS8000 Copy Services*, and *IBM XIV Storage System: Concepts, Architecture and Usage*. He holds a degree in Electrical Engineering from the Technical University in Darmstadt.

Carlo Saba is a Test Engineer for XIV in Tucson, AZ. He has been working with the IBM XIV Storage System product since shortly after its introduction and is a Certified XIV Administrator. Carlo graduated from the University of Arizona in 2007 with a BSBA in MIS and a minor in Spanish.

Eugene Tsypin is an IT Specialist who currently works for IBM STG Storage Systems Sales in Russia. Eugene has over 15 years of experience in the IT field, ranging from systems administration to enterprise storage architecture. He is working as Field Technical Sales Support for storage systems. His areas of expertise include performance analysis and disaster recovery solutions in enterprises utilizing the unique capabilities and features of the IBM XIV Storage System as well as other IBM storage, server, and software products.

Kip Wagner is an Advisory Product Engineer for XIV in Tucson, Arizona. He has more than 24 years experience in field support and systems engineering and is a Certified XIV Engineer and Administrator. Kip was a member of the initial IBM XIV product launch team who helped design and implement a world wide support structure specifically for XIV. He also helped develop training material and service documentation used in the support organization. He is currently the team leader for XIV product field engineering supporting customers in North and South America. He also works with a team of engineers from around the world to provide field experience feedback into the development process to help improve product quality, reliability and serviceability.



Axel Westphal is working as an IT Specialist for Workshops and Proof of Concepts at the IBM European Storage Competence Center (ESSC) in Mainz, Germany. He joined IBM in 1996, working for Global Services as a System Engineer. His areas of expertise include setup and demonstration of IBM System Storage products and solutions in various environments. Since 2004 he has been responsible for storage solutions and Proof of Concepts conducted at the ESSC with DS8000, SAN Volume Controller, and XIV. He has been a contributing author to various DS6000™ and DS8000 related IBM Redbooks publications.



Ralf Wohlfarth is an IT Specialist in the IBM European Storage Competence Center in Mainz, working in technical sales support with focus on the IBM XIV Storage System. In 1998 he joined IBM and has been working in last level product support for IBM System Storage and Software since 2004. He had the lead for post sales education during a product launch of an IBM Storage Subsystem and resolved complex customer situations. During an assignment in the US, he acted as liaison into development and has been driving product improvements into hardware and software development. Ralf holds a masters degree in Electrical Engineering, with a major in Telecommunications from the University of Kaiserslautern, Germany.



Kip Wagner, Eugene Tsypin, Nils Nause, Carlo Saba, Markus Oscheka, Roger Eriksson, Bertrand Dufrasne

We extend special thanks to:

John Bynum
Worldwide Technical Support Management
IBM US, San Jose

For their technical advice, support, and other contributions to this project, many thanks to:

Rami Elron, Richard Heffel, Aviad Offer, Izhar Sharon, Omri Palmon, Iddo Jacobi, Orli Gan, Moshe Dahan, Dave Denny, Juan Yanes, John Cherbini, Alice Bird, Rosemary McCutchen, Brian Sherman, Bill Wiegand, Michael Hayut, Moriel Lechtman, Hank Sautter, Chip Jarvis, Avi Aharon, Shimon Ben-David, Chip Jarvis, Dave Adams, Eyal Abraham
IBM

Robby Jackard,
ATS Group, LLC

Thanks also to the authors of the previous editions:

Hank Sautter, Marc Kremkus, Christina Lara, Lisa Martinez, Alexander Safonov, Stephen Solewin, Aubrey Applewhaite, Ron Verbeek, Pete Wendler, Guenter Rebmann, Christopher Sansone, Attila Grosz, Giacomo Chiaparini.

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us.

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review IBM Redbooks publications form found at:

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks publications

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

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

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks publications 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>



IBM XIV Storage System overview

The IBM XIV Storage System is a revolutionary high-end fully scalable enterprise disk storage system that is based on a grid of standard, off-the-shelf hardware components. Its outstanding virtualized grid design allows massive parallelism with completely balanced paths, CPUs, caches, and disks, with virtually no possibility for hotspots. It has been designed with an easy to use and intuitive GUI that allows administrators to become productive in a very short time.

In this chapter we provide a high level overview of the IBM XIV Storage System.

1.1 Introduction

The IBM XIV Storage System architecture is designed to deliver performance, scalability, and ease of management while harnessing the high capacity and cost benefits of Serial Advanced Technology Attachment (SATA) drives. The system employs off-the-shelf components as opposed to traditional offerings that use proprietary designs, thus requiring more expensive components.

The IBM XIV Storage System incorporates fully automated recovery processes. The data redistribution takes place automatically after new hardware is added, removed, or has failed. Only the data necessary to keep the system fully redundant is redistributed, minimizing recovery time and data movement. Because of this innovative distribution algorithm, the performance is always consistent.

The virtualized grid architecture and algorithms used divide the host data into 1 MB partitions and distribute them pseudo-randomly across all disks, leading to a consistent load on all components and eliminating hotspots. During the distribution process, the data is always mirrored by ensuring that each 1 MB partition is kept in at least two separate locations within the system.

Customers receive exceptionally low total cost of ownership because the system software licensing includes snapshots, thin provisioning, data migration, asynchronous and synchronous mirroring, and an intuitive GUI, combined with dramatic efficiencies in capacity, power, and space. Due to the unique physical architectural design of the system, including off-the-shelf modules, network switches, and power components, new technologies can be adopted quite easily.

These design points provide the following benefits:

- ▶ Reduced Total Cost of Ownership (TCO)
- ▶ Flexibility to meet changing business needs
- ▶ Consolidate and virtualize data storage
- ▶ High availability and performance characteristics

1.2 Total Cost of Ownership (TCO)

Traditional components of TCO calculation for enterprise storage usually include the following costs:

- ▶ Cost of initial purchase
- ▶ Cost of subsequent upgrades
- ▶ Cost of hardware and software maintenance
- ▶ Cost of storage management
- ▶ Cost of electricity and cooling

The IBM XIV Storage System includes software licenses for all features and functionality. There is no need to purchase additional software licensing when you decide to add additional capacity to your system or use an advanced feature such as mirroring or data migration.

The IBM XIV Storage System Management Software suite, along with the virtualized grid architecture, greatly simplify the layout and management of data within the storage array, which reduces the cost of managing the system.

The IBM XIV Storage System uses a grid array of low cost, high capacity (1 TB or 2 TB) SATA drives for storing data, which provides performance similar to Fibre Channel (FC) drives in traditional storage arrays. The grid architecture used in the IBM XIV Storage Systems eliminates the need for “spare drives” in the event of a drive failure. As a result, all drives in the system are fully utilized, reducing the amount of idle components consuming resources.

These features and more reduce the overall TCO of the IBM XIV Storage System.

1.3 System models and components

The IBM XIV Storage System family consists of two machine types:

- ▶ Machine type 2812 Model A14
- ▶ Machine type 2810 Model A14

The 2812 Model A14 supports a 3-year warranty to complement the 1-year warranty offered by the existing and functionally equivalent 2810 Model A14.

Both machine types are available in the following configurations:

- ▶ 6 modules (including 3 Interface Modules)
- ▶ 9 -15 modules (including 6 Interface Modules)

Both machine types include the following components, which are visible in Figure 1-1:

- ▶ 3-6 Interface Modules, each with 12 SATA disk drives (1 TB or 2 TB, but no intermix)
- ▶ 3-9 Data Modules, each with 12 SATA disk drives (1 TB or 2 TB, but no intermix)
- ▶ An Uninterruptible Power Supply (UPS) module complex comprising three redundant UPS units
- ▶ Two Ethernet switches and an Ethernet Switch Redundant Power Supply (RPS)
- ▶ A Maintenance Module
- ▶ An Automatic Transfer Switch (ATS) for external power supply redundancy
- ▶ A modem, connected to the Maintenance Module for externally servicing the system; note that the modem (feature number 9101) is not available in all countries.

All of the modules in the system are linked through an internal redundant Gigabit Ethernet network, which enables maximum bandwidth utilization and is resilient to at least any single component failure.

The system and all of its components come pre-assembled and wired in a lockable rack.




Figure 1-1 IBM XIV Storage System components: Front and rear view

1.4 Key design features

This section describes the key design features of the IBM XIV Storage System architecture.

We discuss these key design points and underlying architectural concepts in detail in Chapter 2, “IBM XIV logical architecture and concepts” on page 13.

1.4.1 Massive parallelism

The system architecture ensures full exploitation of all system components. Any I/O activity involving a specific logical volume in the system is always inherently handled by all spindles. The system harnesses all storage capacity and all internal bandwidth, and it takes advantage of all available processing power, which is as true for host-initiated I/O activity as it is for system-initiated activity, such as rebuild processes and snapshot generation. All disks, CPUs, switches, and other components of the system contribute to the performance of the system at all times.

1.4.2 Workload balancing

The workload is evenly distributed over all hardware components at all times. All disks and modules are utilized equally, regardless of access patterns. Despite the fact that applications might access certain volumes more frequently than other volumes, or access certain parts of

a volume more frequently than other parts, the load on the disks and modules will be balanced perfectly.

Pseudo-random distribution ensures consistent load-balancing even after adding, deleting, or resizing volumes as well as adding or removing hardware. This balancing of all data on all system components eliminates the possibility of a hot-spot being created.

1.4.3 Self-healing

Protection against double disk failure is provided by an efficient rebuild process that brings the system back to full redundancy in minutes. In addition, the IBM XIV Storage System extends the self-healing concept, resuming redundancy even after failures in components other than disks.

1.4.4 True virtualization

Unlike other system architectures, storage virtualization is inherent to the basic principles of the IBM XIV Storage System design. Physical drives and their locations are completely hidden from the user, which dramatically simplifies storage configuration, letting the system lay out the user's volume in the optimal way. The automatic layout maximizes the system's performance by utilizing system resources for each volume, regardless of the user's access patterns. With XIV, you no longer have to decide which type of data protection to use for storing data or how many drives you need to dedicate for a specific application. Just allocate the needed storage space and the system does the rest. This allows the user to respond to growing storage needs in minutes instead of hours or days, as with traditional storage systems.

1.4.5 Thin provisioning

The system supports thin provisioning at the storage pool level, which is the capability to allocate additional storage space to applications on a just-in-time and as needed basis, allowing the most efficient use of available resources and as a result, significant cost savings compared to traditional provisioning techniques. This benefit is achieved by defining a logical capacity that is larger than the actual physical capacity used.

1.4.6 Processing power

The IBM XIV Storage System open architecture leverages the latest processor technologies and is more scalable than solutions that are based on a closed architecture. The IBM XIV Storage System avoids sacrificing the performance of one volume over another, and therefore requires little to no tuning.

1.5 The XIV Storage System software

The IBM XIV system software 10.1 (or later) provides the functions of the system, which include these:

- **Bundled Advanced Features:**

All the features of the IBM XIV including advanced features such as migration and mirroring are included free of charge and apply to the entire storage capacity.

► **Non-Disruptive Code Load (NDCL):**

System software code can be upgraded without requiring downtime. This enables “non-stop” production environments to remain running while new code is upgraded.

The code upgrade is run on all modules in parallel and the process is fast enough to minimize impact on hosts applications.

No data migration or rebuild process is allowed during the upgrade. Mirroring, if any, will be suspended during the upgrade and automatically reactivated upon completion.

Storage management operations are also not allowed during the upgrade, although the status of the system and upgrade progress can be queried. It is also possible to cancel the upgrade process up to a point of no return.

Firmware: NDCL does not apply to certain component firmware upgrades (for instance, module BIOS or similar firmware). These cases might require a phase in / phase out or a similar process of the impacted modules, which can be scheduled for a separate maintenance period to reduce impact to production environments.

► **Support for multiple snapshots:**

The snapshot capabilities within the IBM XIV Storage System Software utilize a metadata, redirect-on-write design that allows snapshots to occur in a subsecond time frame with little performance overhead. The system is designed to support multiple full or differential snapshots of a volume. Any of the snapshots can be made writable, and then snapshots can be taken of the newly writable snapshots (snapshots of snapshots). Volumes can even be restored from these writable snapshots.

► **Synchronous and asynchronous remote mirroring to another XIV Storage System:**

Synchronous or asynchronous (requires software level 10.2.1 or higher) remote mirroring can be performed over Fibre Channel (FC) or IP network Small Computer System Interface (iSCSI) connections. Synchronous remote mirroring is used when data at local and remote sites must remain synchronized at all times.

► **Support for thin provisioning:**

Thin provisioning allows administrators to over-provision allocated storage within a Storage Pool. This is done by defining volume sizes that are larger than the physical capacity of the storage pool. Unlike other approaches, the physical capacity of the storage pool only needs to be larger than the actual consumed size of the volumes within that pool. Physical capacity of the storage pool needs to be increased only when actual written data of the volumes increase.

► **Support for in-band data migration of heterogeneous storage:**

The IBM XIV Storage System is also capable of acting as a host, gaining access to volumes on an existing legacy storage system. The system is configured as a proxy to respond to requests between the current hosts and the legacy storage while migrating all existing data in the background. In addition, the IBM XIV Storage System supports thick-to-thin data migration, which allows the system to reclaim any allocated space that is not occupied by actual data.

► **Authentication using Lightweight Directory Access Protocol (LDAP):**

LDAP can be used to provide user logon authentication allowing the IBM XIV Storage System to integrate with Microsoft® Active Directory (AD) or Sun Java™ Systems Directory Server (formerly Sun ONE Directory). Multiple directory servers can be configured to provide redundancy in case one server becomes unavailable.

- ▶ Robust user auditing with access control lists:
The IBM XIV Storage System Software offers the capability for robust user auditing with Access Control Lists (ACLs) in order to provide more control and historical information.
- ▶ Support for Tivoli® Storage Productivity Center (TPC):
TPC can discover IBM XIV Storage Systems and all internal components, manage capacity for storage pools, including allocated, unallocated, and available capacity with historical trending on utilization. It can also receive events and define policy-based alerts based on user-defined triggers and thresholds.
- ▶ vStorage API Array Integration (VAAI) support:
(Available with XIV Storage System software Version 10.2.4.) Using VAAI moves the storage-related tasks that were previously performed by VMware onto the storage system. Transferring the processing burden dramatically reduces performance overhead; speeds processing; frees up VMware for more mission-critical tasks, such as adding applications; simplifies management; and positions the virtual environment for additional capacity and scalability. When hardware acceleration is enabled with XIV, operations such as VM provisioning, VM cloning, and VM migration complete dramatically faster, and with minimal impact on the ESX server, increasing scalability and storage performance.
- ▶ Host Rate Limiting - Quality of Service (QoS):
(Available with XIV Storage System software Version 10.2.4.) XIV system resources (such as storage and cache) constitute a virtualized environment that is shared by all hosts and applications. Such an approach lends itself exceptionally well to accommodate high performance requirements for multiple applications with similar performance objectives through fair resource allocation. In environments with applications that have various performance objectives, XIV's Host Rate Limiting feature enables the customer to restrict the system's processing power allocated to specified hosts/applications, thereby maximizing the power available for applications requiring the utmost performance.

1.5.1 IBM XIV Storage Management Tools

The IBM XIV Management Tools, including the IBM XIV Graphical User Interface (XIVGUI), IBM XIV online monitoring tool (XIVTop) and IBM XIV Command Line Interface (XCLI), allow the user to configure and monitor their IBM XIV Storage Systems.

The IBM XIV Management Tools (Version 2.4.3.b) are available for various operating system platforms, including these:

- ▶ Windows (32 and 64 bit versions)
- ▶ Linux (Red Hat 5, SLES 10, SLES 11)
- ▶ AIX 5.3, AIX 6.1
- ▶ Solaris v9, Solaris v10
- ▶ HP-UX 11i v2, HP-UX 11i v3

The tools can be downloaded at the following site:

<ftp://ftp.software.ibm.com/storage/XIV/GUI/>

Note that the GUI, XIVTop, and XCLI tools are bundled together in a single downloadable package for each supported operating system.

1.5.2 IBM XIV Storage Management GUI

The IBM XIV Storage Management GUI acts as the management console for the storage system. A simple and intuitive GUI enables storage administrators to manage and monitor all system aspects easily, with almost no learning curve.

Figure 1-2 shows one of the top level configuration panels.

The GUI also contains a demonstration mode.

To use the demonstration mode, login as user P10DemoMode. A password is not required.



Figure 1-2 The IBM XIV Storage Manager GUI

1.5.3 IBM XIV Storage Management online monitoring with XIVTop

The XIVTop application allows you to view and monitor performance information for defined volumes and hosts in real-time. See Figure 1-3 for an illustration.




Figure 1-3 XIVTop

The tool was designed to quickly display the volumes or hosts that are consuming the most system resources at any given time. Information can be sorted in a variety of ways and exported to a file for use in spreadsheet applications.

1.5.4 IBM XIV Storage Management XCLI

The IBM XIV Storage Management XCLI is a comprehensive command line interface to configure and monitor the system. All the functions available in the GUI are also available in the XCLI. The XCLI can be used in a shell environment to interactively configure the system or as part of a script to perform lengthy and/or complex tasks.

Figure 1-4 shows a command being run in a Windows DOS shell.

```
C:\>xcli -c XIV_LAB config_get
Name                                     Value
dns_primary                             192.168.1.2
dns_secondary                           192.168.1.3
system_name                             XIV LAB 01 EBC
snmp_location                           Unknown
snmp_contact                            Unknown
snmp_community                           XIV
snmp_trap_community                      XIV
system_id                               12345
machine_type                             2810
machine_model                            A14
machine_serial_number                    6012345
email_sender_address                     XIV_alert@customer.com
email_reply_to_address
email_subject_format                     {severity}: {description}
internal_email_subject_format             2810-A14: 6012345: {severity}:
{description}
iscsi_name                               iqn.2005-10.com.xivstorage:000105
maximal_snapshot_deletion_priority        4
timezone                                 -7200
fc_proof                                 yes
ntp_server                               192.168.1.100
ups_control                               yes
support_center_port_type                  Management
```

Figure 1-4 XCLI commands in a Windows DOS Shell

The XCLI is installed as part of the FULL GUI installation, or it can be installed alone.

1.6 Host support

The IBM XIV Storage System can be attached to a variety of host operating systems, including these:

- ▶ Microsoft Windows
- ▶ VMware ESX
- ▶ IBM AIX
- ▶ HPUX
- ▶ Linux
- ▶ Oracle Solaris

For details on each Operating System and versions supported, see the IBM XIV interoperability matrix or the System Storage Interoperability Center (SSIC) at:

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

There are also a variety of software “agents” available to help you manage environments that connect to an IBM XIV Storage System.

Note: Also see the IBM Redbooks publication, *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904.

► **Tivoli Storage Productivity Center Storage Resource Agent**

http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=%2Fcom.ibm.tpc_V42.doc%2Ffqz0_r_whats_new_release_4.2.html

This provides support for IBM XIV Storage System provisioning, data path explorer, and performance management reporting.

► **VMware vCenter Site Recovery Manager (SRM) Storage Replication Adapter**

http://downloads.vmware.com/d/details/srm100u1_patch5/dHBiZCp0QGJkZGhk

Storage Replication Adapters are software modules for VMware Site Recovery Manager that simplify the use of storage replication software with VMware vSphere.

► **IBM XIV Management Console for VMware vCenter version 1.0.1**

http://www-01.ibm.com/support/docview.wss?rs=1319&context=STJTAG&dc=D400&q1=ssg1*&uid=ssg1S4000884&loc=en_US&cs=utf-8&lang=en

IBM XIV Management Console for VMware vCenter integrates XIV storage with the VMWare VSphere management console to allow more efficient management of VMWare components residing on XIV storage.

► **IBM XIV MSCS Agent, version 1.0.1**

http://www-01.ibm.com/support/docview.wss?rs=1319&context=STJTAG&dc=D400&q1=ssg1*&uid=ssg1S4000869&loc=en_US&cs=utf-8&lang=en

The XIV MSCS Agent maintains volume mirroring between two IBM XIV Storage Systems and Microsoft Failover clusters.

► **IBM XIV VSS Provider (xProv) for Windows, v2.2.3 & v2.2.2**

http://www-01.ibm.com/support/docview.wss?rs=1319&context=STJTAG&dc=D400&q1=ssg1*&uid=ssg1S4000800&loc=en_US&cs=utf-8&lang=en

XIV VSS Provider integrates IBM XIV Storage System Snapshot™ capabilities with Microsoft Windows Shadow Copy Services.

► **IBM XIV MPIO Management Console version 1.0.0**

http://www-01.ibm.com/support/docview.wss?rs=1319&context=STJTAG&dc=D400&q1=ssg1*&uid=ssg1S4000835&loc=en_US&cs=utf-8&lang=en

The IBM XIV MPIO Console application provides monitoring, host storage connectivity, and multipath management capabilities through an intuitive and powerful GUI.



IBM XIV logical architecture and concepts

In this chapter we elaborate on various underlying design and architectural concepts of the IBM XIV Storage System that we introduced in Chapter 1, “IBM XIV Storage System overview” on page 1.

We cover the following topics:

- ▶ Architectural elements
- ▶ Parallelism
- ▶ Virtualization
- ▶ Data distribution
- ▶ Thin provisioning
- ▶ Self-healing and resiliency
- ▶ Rebuild redundancy

2.1 Architecture overview

In this section we present an overview of the IBM XIV Storage System architecture.

2.1.1 Features

The IBM XIV Storage System architecture incorporates a variety of features designed to uniformly distribute data across internal resources. This unique data distribution method fundamentally differentiates the IBM XIV Storage System from conventional storage subsystems, thereby offering numerous availability, performance, and management benefits across both physical and logical elements of the system.

2.1.2 Hardware elements

To convey the conceptual principles that comprise the IBM XIV Storage System architecture, it is useful to first provide a glimpse of the physical infrastructure. Further details are covered in Chapter 3, “IBM XIV physical architecture, components, and planning” on page 47.

Modules

The primary components of the IBM XIV Storage System are known as *modules*. Modules provide processing, cache, and host interfaces and are comprised of “off the shelf” Intel® based systems. The modules are redundantly connected to one another through an internal switched Ethernet network, as shown in Figure 2-1. All of the modules work together concurrently as elements of a grid architecture, and therefore, the system harnesses the powerful parallelism inherent in such a distributed computing environment. We discuss the grid architecture in 2.2, “Parallelism” on page 16.




Figure 2-1 IBM XIV Storage System major hardware elements

Data Modules

At a conceptual level, the Data Modules function as the elementary “building blocks” of the system, providing storage capacity, processing power, and caching, in addition to advanced system-managed services. The Data Module’s ability to share and manage system software and services are key elements of the physical architecture, as depicted in Figure 2-2.

Interface Modules

Interface Modules are equivalent to Data Modules in all aspects, with the following exceptions:

- ▶ In addition to disk, cache, and processing resources, Interface Modules are designed to include both Fibre Channel and iSCSI interfaces for host system connectivity, remote mirroring, and data migration activities. Figure 2-2 conceptually illustrates the placement of Interface Modules within the topology of the IBM XIV Storage System architecture.
- ▶ The system services and software functionality associated with managing external I/O reside exclusively on the Interface Modules.

Ethernet switches

The XIV Storage System contains a redundant switched Ethernet network that transmits both data and metadata traffic between the modules. Traffic can flow in any of the following ways:

- ▶ Between two Interface Modules
- ▶ Between two Data Modules
- ▶ Between an Interface Module and a Data Module




Figure 2-2 Architectural overview

Architecture: Figure 2-2 depicts the conceptual architecture only. Do not misinterpret the number of connections and such as a precise hardware layout.

2.2 Parallelism

The concept of *parallelism* pervades all aspects of the IBM XIV Storage System architecture by means of a balanced, redundant data distribution scheme in conjunction with a pool of distributed (or grid) computing resources. In order to explain the principle of parallelism further, it is helpful to consider the ramifications of both the hardware and software implementations independently. We subsequently examine virtualization principles in 2.3, “Full storage virtualization” on page 18.

Important: The IBM XIV Storage System exploits parallelism at *both* hardware and software levels.

2.2.1 Hardware parallelism and grid architecture

The XIV grid design (Figure 2-3) entails the following characteristics:

- ▶ Both interface and Data Modules work together in a distributed computing sense. However, the Interface Modules also have additional functions and features associated with host system connectivity.
- ▶ The modules communicate with each other through the internal, redundant Ethernet network.
- ▶ The software services and distributed computing algorithms running within the modules collectively manage all aspects of the operating environment.

Design principles

The IBM XIV Storage System grid architecture, by virtue of its distributed topology and “off the shelf” Intel components, ensures that the following design principles are possible:

- ▶ Performance:
 - The relative effect of the loss of a single component is minimized.
 - All modules are able to participate equally in handling the total workload.

This design principle is true regardless of access patterns. The system architecture enables excellent load balancing, even if certain applications access certain volumes, or certain parts within a volume, more frequently.
- ▶ Compatibility:
 - Modules consist of standard “off the shelf” components.

Because components are not specifically engineered for the system, the resources and time required for the development of newer hardware technologies are minimized. This benefit, coupled with the efficient integration of computing resources into the grid architecture, enables the system to realize the rapid adoption of the newest hardware technologies available without the need to deploy a whole new subsystem.
- ▶ Scalability:
 - Computing resources can be dynamically changed
 - “Scaled out” by adding new modules to accommodate both new capacity and new performance demands
 - “Scaled up” by upgrading modules




Figure 2-3 IBM XIV Storage System scalable conceptual grid architecture

Proportional scalability

Within the IBM XIV Storage System, each module contains all of the pertinent hardware elements that are necessary for a grid topology (processing, caching, and storage). All modules are connected through a scalable network. This aspect of the grid infrastructure enables the relative proportions of cache, processor, disk, and interconnect bandwidth to remain optimal even in the event that modules are added or removed:

- ▶ **Linear cache growth:** The total system cache size and cache bandwidth increase linearly with disk capacity, because every module is a self-contained computing resource that houses its own cache. Note that the cache bandwidth scales linearly in terms of both host-to-cache and cache-to-disk throughput, and the close proximity of cache, processor, and disk is maintained.
- ▶ **Proportional interface growth:** Interface Modules house Ethernet and fibre Channel host interfaces and are able to access not only the local resources within the module, but also the entire system. With every Interface Module added, the system proportionally scales both the number of host interfaces and the bandwidth to the internal resources.
- ▶ **Constant switching capacity:** The internal switching capacity is designed to scale proportionally as the system grows, preventing bottlenecks regardless of the number of modules. This capability ensures that internal throughput scales proportionally to capacity.
- ▶ **Embedded processing power:** Because each module incorporates its own processing power in conjunction with cache and disk components, the ability of the system to perform processor-intensive tasks, such as aggressive prefetch caching, sophisticated cache updates, snapshot management, and data distribution, is always maintained regardless of the system capacity.

2.2.2 Software parallelism

In addition to the hardware parallelism, the IBM XIV Storage System also employs sophisticated algorithms to achieve optimal software parallelism.

Modular software design

The IBM XIV Storage System internal operating environment consists of a set of software functions that are loosely coupled with the hardware modules. These software functions reside on one or more modules and can be redistributed among modules as required, thus ensuring resiliency under changing hardware conditions.

An example of this modular design resides specifically in the Interface Modules. All six Interface Modules actively manage system services and software functionality associated with managing external I/O. Also, three of the Interface Modules deliver the system's management interface service for use with the IBM XIV Storage System.

Data distribution algorithms

Data is distributed across all drives in a *pseudo-random* fashion. The patented algorithms provide a uniform yet random spreading of data across all available disks to maintain data resilience and redundancy. Figure 2-4 on page 21 provides a conceptual representation of the pseudo-random data distribution within the IBM XIV Storage System.

For more details about the topic of data distribution and storage virtualization, see 2.3.2, "Logical system concepts" on page 20.

2.3 Full storage virtualization

The data distribution algorithms employed by the IBM XIV Storage System are innovative and unique in that they are deeply integrated into the system architecture itself, instead of at the host or storage area network level.

To fully appreciate the value inherent to the virtualization design that is used by the IBM XIV Storage System, it is helpful to remember various aspects of the physical and logical relationships that comprise conventional storage subsystems. Specifically, traditional subsystems rely on storage administrators to carefully plan the relationship between logical structures, such as arrays and volumes, and physical resources, such as disk packs and drives, in order to strategically balance workloads, meet capacity demands, eliminate hot-spots, and provide adequate performance.

2.3.1 IBM XIV Storage System virtualization design

The implementation of full storage virtualization employed by the IBM XIV Storage System eliminates many of the potential operational drawbacks that can be present with conventional storage subsystems, while maximizing the overall usefulness of the subsystem.

The IBM XIV Storage System virtualization offers the following benefits:

- ▶ Easier volume management:
 - Logical volume placement is driven by the distribution algorithms, freeing the storage administrator from planning and maintaining volume layout. The data distribution algorithms manage all of the data in the system collectively without deference to specific logical volume definitions.
 - Any interaction, whether host or system driven, with a specific logical volume in the system is inherently handled by all resources; it harnesses all storage capacity, all internal bandwidth, and all processing power currently available in the system.
 - Logical volumes are not exclusively associated with a subset of physical resources.
 - Logical volumes can be dynamically resized.
 - Logical volumes can be thinly provisioned, as discussed in 2.3.5, "Capacity allocation and thin provisioning" on page 26.

- ▶ Consistent performance and scalability:
 - Hardware resources are always utilized equally, because all logical volumes always span all physical resources and are therefore able to reap the performance potential of the full system and maintain data integrity.
 - Virtualization algorithms automatically redistribute the logical volumes data and workload when new hardware is added, thereby maintaining the system balance while preserving transparency to the attached hosts.
 - In the event of a hardware failure, data is automatically, efficiently, and rapidly rebuilt across all the drives and modules in the system. This preserves host transparency, equilibrium, and data redundancy at all times while virtually eliminating any performance penalty associated with traditional RAID rebuilds.
 - When new hardware is added to the system, data is transparently redistributed across all resources to restore equilibrium to the system.
 - There are no “pockets” of capacity, “orphaned” disk space, or resources that are inaccessible due to array mapping constraints or data placement.
- ▶ Flexible snapshots:
 - Full storage virtualization incorporates snapshots that are differential in nature; only updated data consumes physical capacity:
 - Many concurrent snapshots

Multiple concurrent snapshots are possible because a snapshot uses physical space only after a change has occurred on the source.
 - Multiple snapshots of a single master volume can exist independently of each other.
 - Snapshots can be cascaded, in effect, creating snapshots of snapshots.
 - Creation and deletion of snapshots do not require data to be copied and hence occur immediately.
 - When updates occur to master volumes, the system's virtualized logical structure enables it to preserve the original point-in-time data associated with any and all dependent snapshots by redirecting the update to a new physical location on disk. This process, which is referred to as *redirect on write*, occurs transparently from the host perspective and uses the virtualized remapping of the updated data to minimize any performance impact associated with preserving snapshots, regardless of the number of snapshots defined for a given master volume.

Snapshot: The XIV snapshot process uses “redirect on write,” which is more efficient than the “copy on write” that is used by many other storage subsystems.

- ▶ Data migration efficiency:
 - XIV supports thin provisioning. When migrating from a system that only supports regular (or thick) provisioning, XIV allows thick-to-thin provisioning of capacity. Thin-provisioned capacity is discussed in 2.3.5, “Capacity allocation and thin provisioning” on page 26.
 - Due to the XIV pseudo-random distribution of data, the performance impact of data migration on production activity is minimized, because the load is spread evenly over all resources.

2.3.2 Logical system concepts

In this section, we elaborate on the logical system concepts, which form the basis for the system full storage virtualization.

Logical constructs

The IBM XIV Storage System logical architecture incorporates constructs that underlie the storage virtualization and distribution of data, which are integral to its design. The logical structure of the system ensures that there is optimum granularity in the mapping of logical elements to both modules and individual physical disks, thereby guaranteeing an equal distribution of data across all physical resources.

Partitions

The fundamental building block of a logical volume is known as a *partition*. Partitions have the following characteristics on the IBM XIV Storage System:

- ▶ All partitions are 1 MB (1024 KB) in size.
- ▶ A partition contains either a primary copy or secondary copy of data:
 - Each partition is mapped to a single physical disk:
 - This mapping is dynamically managed by the system through innovative data distribution algorithms in order to preserve data redundancy and equilibrium. For more information about the topic of data distribution, see “Logical volume layout on physical disks” on page 22.
 - The storage administrator has no control or knowledge of the specific mapping of partitions to drives.
 - Secondary copy partitions are always placed in a separate *Module* from the one containing the primary copy partition.

Important: In the context of the IBM XIV Storage System logical architecture, a partition consists of 1 MB (1024 KB) of data. Do not confuse this definition with other definitions of the term “partition.”

The diagram in Figure 2-4 illustrates that data is uniformly, yet randomly distributed over all disks. Each 1 MB of data is duplicated in a primary and secondary partition for redundancy. The system ensures that the primary partition and its corresponding secondary partition are never located within the same module in order to protect against a single point of failure.




Figure 2-4 Pseudo-random data distribution¹

Logical volumes

The IBM XIV Storage System presents logical volumes to hosts in the same manner as conventional subsystems; however, both the granularity of logical volumes and the mapping of logical volumes to physical disks differ:

- ▶ As discussed previously, every logical volume is comprised of 1 MB (1024 KB) constructs of data known as partitions.
- ▶ The physical capacity associated with a logical volume is *always a multiple of 17 GB (decimal)* because the volume spans all physical drives in the system due to the grid nature of the architecture.

Therefore, while it is possible to present a block-designated logical volume to a host that is *not* a multiple of 17 GB, the *actual physical space* that is allocated for the volume will always be the sum of the minimum number of 17 GB increments needed to meet the block-designated capacity.

Capacity: The initial physical capacity actually allocated by the system upon volume creation can be less than this amount, as discussed in “Actual and logical volume sizes” on page 27.

¹ Copyright 2005-2008 Mozilla. All Rights Reserved. All rights in the names, trademarks, and logos of the Mozilla Foundation, including without limitation, Mozilla, Firefox, as well as the Firefox logo, are owned exclusively by the Mozilla Foundation.

- A maximum total of 4000 volumes and snapshots can be concurrently created on the system.

Storage pools

Storage pools are administrative boundaries that enable storage administrators to manage relationships between volumes and snapshots and to define separate capacity provisioning and snapshot requirements for separate applications or departments. Storage pools are not tied in any way to physical resources, nor are they part of the data distribution scheme. We discuss storage pools and their associated concepts in 2.3.4, “Storage pool concepts” on page 24.

Snapshots

A *snapshot* represents a point-in-time copy of a volume. Snapshots are like volumes except snapshots incorporate dependent relationships with their source volumes, which can be either logical volumes or other snapshots. Because they are not independent entities, a given snapshot does not necessarily wholly consist of partitions that are unique to that snapshot. Conversely, a snapshot image will not share all of its partitions with its source volume if updates to the source occur after the snapshot was created.

Logical volume layout on physical disks

The IBM XIV Storage System manages the distribution of logical volumes over physical disks and modules by means of a dynamic relationship between primary data partitions, secondary data partitions, and physical disks. This virtualization of resources in the IBM XIV Storage System is governed by the data distribution algorithms.

Distribution table

The distribution table is created at system startup, and contains a mapping of every primary and secondary partition, as well as the module and physical disk they reside on. When hardware changes occur, a new distribution table is created and delivered to every module. Each module retains redundant copies of the distribution table.

Volume layout

At a conceptual level, the data distribution scheme can be thought of as an mixture of mirroring and striping. While it is tempting to think of this scheme in the context of RAID 1+0 (10) or 0+1, the low-level virtualization implementation precludes the usage of traditional RAID algorithms in the architecture.

As discussed previously, the IBM XIV Storage System architecture divides logical volumes into 1 MB partitions. This granularity and the mapping strategy are integral elements of the logical design that enable the system to realize the following features and benefits:

- Partitions that make up a volume are distributed on all disks using what is defined as a *pseudo-random distribution function*, which was introduced in 2.2.2, “Software parallelism” on page 17.
 - The distribution algorithms seek to preserve the equality of access among all physical disks under all conceivable conditions and volume access patterns. Essentially, while not truly random in nature, the distribution algorithms in combination with the system architecture preclude the occurrence of “*hot-spots*”:
 - A fully configured IBM XIV Storage System contains 180 disks, and each volume is allocated across at least 17 GB (decimal) of capacity that is distributed evenly across all disks.
 - Each logically adjacent partition on a volume is distributed across a separate disk; partitions are *not* combined into groups before they are spread across the disks.

- The pseudo-random distribution ensures that logically adjacent partitions are never striped sequentially across physically adjacent disks. See 2.2.2, “Software parallelism” on page 17 for a further overview of the partition mapping topology.
- Each disk has its data mirrored across all other disks, *excluding the disks in the same module*.
- Each disk holds approximately one percent of any other disk in other modules.
- Disks have an equal probability of being accessed regardless of aggregate workload access patterns.

Goal distribution: When the number of disks or modules changes, the system defines a new data layout that preserves redundancy and equilibrium. This target data distribution is called the *goal distribution* and is discussed in “Goal distribution” on page 39.

- ▶ As discussed previously in “IBM XIV Storage System virtualization design” on page 18:
 - The storage system administrator does *not* plan the layout of volumes on the modules.
 - Provided that there is space available, volumes can always be added or resized instantly with negligible impact on performance.
 - There are no unusable pockets of capacity known as “orphaned spaces.”
- ▶ When the system is scaled out through the addition of modules, a new *goal distribution* is created whereby just a minimum number of partitions are moved to the newly allocated capacity to arrive at the new distribution table.

The new capacity is fully utilized within a few hours and with no need for any administrative intervention. Thus, the system automatically returns to a state of equilibrium among all resources.

- ▶ Upon the failure or phase-out of a drive or a module, a new goal distribution is created whereby data in non-redundant partitions is copied and redistributed across the remaining modules and drives.

The system rapidly returns to a state in which all partitions are again redundant, because all disks and modules participate in achieving the new goal distribution.

2.3.3 System usable capacity

The XIV Storage System reserves physical disk capacity for:

- ▶ Global spare capacity
- ▶ Metadata, including statistics and traces
- ▶ Mirrored copies of data

Global spare capacity

The dynamically balanced distribution of data across all physical resources by definition eliminates the need of dedicated spare drives that are necessary with conventional RAID technologies. Instead, the IBM XIV Storage System reserves capacity on each disk in order to provide adequate space for the redistribution or rebuilding of redundant data in the event of a hardware failure.

This global spare capacity approach offers advantages over dedicated hot spare drives, which are used only upon failure and are not used otherwise, therefore reducing the number of spindles that the system can leverage for better performance. Also, those non-operating disks are not typically subjected to background scrubbing processes, whereas in XIV, all disks are operating all the time and are subject to examination, which helps detect potential reliability issues with drives.

The global reserved space includes sufficient capacity to withstand the failure of a full module plus three additional disk drives, and will still allow the system to execute a new *goal distribution* to return to full redundancy.

Tolerance of failures: A fully utilized system will tolerate multiple hardware failures, including up to an entire module with three subsequent drive failures *outside of the failed module*, provided that a new goal distribution is fully executed before a subsequent failure occurs. If the system is less than 100% full, it can sustain more subsequent failures based on the amount of unused disk space that will be allocated at the event of failure as a spare capacity. For a thorough discussion of how the system uses and manages reserve capacity under specific hardware failure scenarios, see 2.4, “Reliability, availability, and serviceability” on page 35.

Snapshots: The IBM XIV Storage System does not manage a separate global reserved space for snapshots. We explore this topic in the next section.

Metadata and system reserve

The system reserves roughly 4% of the physical capacity for statistics and traces, as well as the distribution table.

Net usable capacity

The calculation of the *net usable capacity* of the system consists of the total disk count, less disk space reserved for sparing (which is the equivalent of one module plus three more disks), multiplied by the amount of capacity on each disk that is dedicated to data (that is 96% because of metadata and system reserve), and finally reduced by a factor of 50% to account for data mirroring achieved by the secondary copy of data.

For example:

An XIV system with 15 modules populated with 1 TB drives will have a net usable capacity of approximately 79 TB based on the following formula: $((180 - 15) * 1 \text{ GB}) * .96 / 2 = 79.2$

2.3.4 Storage pool concepts

The concept of *storage pools* is purely administrative. While the hardware resources within the IBM XIV Storage System are virtualized in a global sense, the available capacity in the system can be administratively portioned into separate and independent storage pools. Essentially, storage pools function as a means to effectively manage a related group of similarly provisioned logical volumes and their snapshots.

Improved management of storage space

Storage pools form the basis for controlling the usage of storage space by imposing a capacity quota on specific applications, a group of applications, or departments, enabling isolated management of relationships within the associated group of logical volumes and snapshots.

A *logical volume* is defined within the context of one and only one storage pool. As storage pools are logical constructs, a volume and any snapshots associated with it can be moved to any other storage pool, as long as there is sufficient space within the target storage pool.

As a benefit of the system virtualization, there are no limitations on the associations between logical volumes and storage pools. In fact, manipulation of storage pools consists exclusively of metadata transactions and does not trigger any copying of data. Therefore, changes are completed instantly and without any system performance degradation.

Consistency groups

A *consistency group* is a group of volumes of which a snapshot can be made at the same point in time, thus ensuring a consistent image of all volumes within the group at that time. The concept of a consistency group is common among storage subsystems in which it is necessary to perform concurrent operations collectively across a set of volumes, so that the result of the operation preserves the consistency among volumes. For example, effective storage management activities for applications that span multiple volumes, or creating point-in-time backups, is not possible without first employing consistency groups.

This consistency between the volumes in the group is paramount to maintaining data integrity from the application perspective. By first grouping the application volumes into a consistency group, it is possible to later capture a consistent state of all volumes within that group at a given point-in-time using a special snapshot command for consistency groups.

Issuing this type of a command results in the following process:

1. Complete and destage writes across the constituent volumes.
2. Instantaneously suspend I/O activity simultaneously across all volumes in the consistency group.
3. Create the snapshots.
4. Finally, resume normal I/O activity across all volumes.

The IBM XIV Storage System manages these suspend and resume activities for all volumes within the consistency group.

Consistency: Additional mechanisms or techniques, such as those provided by the Microsoft Volume Shadow Copy Services (VSS) framework, might still be required to maintain full application consistency from a host system perspective. See the Redbooks publication, *IBM XIV Storage System: Copy Services and Migration*, SG24-7759.

Storage pool relationships

Storage pools facilitate the administration of relationships among logical volumes, snapshots, and consistency groups.

The following principles govern the relationships between logical entities within the storage pool:

- ▶ A logical volume can have multiple independent snapshots. This logical volume is also known as a *master volume*.
- ▶ A *master volume* and all of its associated snapshots always reside in the same storage pool.
- ▶ A volume can only be part of a single consistency group and a single storage pool.
- ▶ All volumes of a consistency group must reside in the same storage pool.

Storage pools have the following characteristics:

- ▶ The size of a storage pool can range from 17 GB (the minimum size that can be assigned to a logical volume) to 79 TB (entire system when equipped with 1 TB drives).
- ▶ *Snapshot reserve capacity* is defined within each Storage Pool and is effectively *maintained separately* from logical, or master, volume capacity. The same principles apply for thinly provisioned Storage Pools, which are discussed in “Thinly provisioned storage pools” on page 28, with the exception that space is not *guaranteed* to be available for snapshots due to the potential for hard space depletion, as discussed in “Depletion of hard capacity” on page 34.

The following principles apply to snapshots:

- Snapshots are structured in the same manner as logical, or master, volumes.

Snapshot reserve: The system preemptively deletes snapshots if the snapshots fully consume the allocated space of the storage pool. Therefore, you must ensure that adequate space is allocated for snapshot reserve when defining a storage pool.

The snapshot reserve needs to be a minimum of 34 GB.

- As mentioned before, snapshots will only be automatically deleted when there is inadequate physical capacity available within the context of each storage pool. This process is managed by a snapshot deletion priority scheme. Therefore, when the capacity of a storage pool is exhausted, only the snapshots that reside in the affected storage pool are deleted in order of the deletion priority.
- ▶ The space allocated for a storage pool can be dynamically changed by the storage administrator:
 - The storage pool can be increased in size.
 - The storage pool can be decreased in size. It is limited only by the space that is consumed by the volumes and snapshots that are defined within that storage pool.
- ▶ The designation of a storage pool as a regular pool or a thinly provisioned pool can be dynamically changed even for existing storage pools. Thin provisioning is discussed in-depth in 2.3.5, “Capacity allocation and thin provisioning” on page 26.
- ▶ The storage administrator can relocate logical volumes between storage pools without any limitations, provided that there is sufficient free space in the target storage pool:
 - If necessary, the target storage pool capacity can be dynamically increased prior to volume relocation, assuming sufficient unallocated capacity available in the system.
 - When a logical volume is relocated to a target storage pool, sufficient space must be available for all of its snapshots to reside in the target storage pool as well.

Notes:

- ▶ When moving a volume into a storage pool, the size of the storage pool is not automatically increased by the size of the volume. When removing a volume from a storage pool, the size of the storage pool does not decrease by the size of the volume.
- ▶ The system defines capacity using decimal metrics. 1 GB is 1 000 000 000 bytes using decimal metrics. By contrast, 1 GB is 1 073 741 824 bytes using binary metrics.

2.3.5 Capacity allocation and thin provisioning

Thin provisioning is a central theme of the virtualized design of the IBM XIV Storage System, because it uncouples the virtual, or apparent, allocation of a resource from the underlying hardware allocation.

The IBM XIV Storage System’s implementation of thin provisioning provides these benefits:

- ▶ Capacity associated with specific applications or departments can be dynamically increased or decreased per the demand imposed at a given point in time, without necessitating an accurate prediction of future needs. Physical capacity is only committed to the logical volume when the associated applications execute writes, as opposed to when the logical volume is initially allocated.

- Because the total system capacity is architected as a globally available pool, thinly provisioned resources share the same “buffer” of free space, which results in highly efficient aggregate capacity utilization without pockets of inaccessible unused space.

With the static, inflexible relationship between logical and physical resources commonly imposed by traditional storage subsystems, each application’s capacity must be managed and allocated independently. This situation often results in a large percentage of the total system capacity remaining unused, because the capacity is confined within each volume at a highly granular level.

- Capacity acquisition and deployment can be more effectively deferred until actual application and business needs demand additional space, in effect facilitating an on-demand infrastructure.

Actual and logical volume sizes

The physical capacity that is assigned to traditional volumes is equivalent to the logical capacity presented to hosts, which does not have to be the case with the IBM XIV Storage System. For a given logical volume, there are effectively two associated sizes. The physical capacity allocated for the volume is not static, but it increases as host writes fill the volume.

Logical volume size

The *logical volume size* is the size of the logical volume that is observed by the host, as defined upon volume creation or as a result of a resizing command. The storage administrator specifies the volume size in the same manner regardless of whether the storage pool the volume resides in will be a thin pool or a regular pool. The volume size is specified in one of two ways, depending on units:

- In terms of GB: The system will allocate the soft volume size as the minimum number of discrete 17 GB increments needed to meet the requested volume size.
- In terms of blocks: The capacity is indicated as a discrete number of 512 byte blocks. The system will still allocate the soft volume size consumed within the storage pool as the minimum number of discrete 17 GB increments needed to meet the requested size (specified in 512 byte blocks); however, the size that is reported to hosts is equivalent to the precise number of blocks defined.

Incidentally, the snapshot reserve capacity associated with each storage pool is a soft capacity limit, and it is specified by the storage administrator, though it effectively limits the hard capacity consumed collectively by snapshots as well.

Tip: Defining logical volumes in terms of blocks is useful when you must precisely match the size of an existing logical volume residing on another system.

Actual volume size

This value reflects the total size of volume areas that were written by hosts. The actual volume size is not controlled directly by the user and depends only on the application behavior. It starts from zero at volume creation or formatting and can reach the logical volume size when the entire volume has been written. Resizing of the volume affects the logical volume size, but does not affect the actual volume size.

The actual volume size reflects the physical space used in the volume, as a result of host writes. It is discretely and dynamically provisioned by the system, not the storage administrator. The discrete additions to actual volume size can be measured in two ways, by considering the allocated space or the consumed space. The allocated space reflects the physical space used by the volume in 17 GB increments. The consumed space reflects the physical space used by the volume in 1 MB partitions.

In both cases, the upper limit of this provisioning is determined by the logical size assigned to the volume:

- ▶ Capacity is allocated to volumes by the system in increments of 17 GB due to the underlying logical and physical architecture; there is no smaller degree of granularity than 17 GB. For more details, see 2.3.2, “Logical system concepts” on page 20.
- ▶ Application write access patterns determine the rate at which the allocated hard volume capacity is consumed and subsequently the rate at which the system allocates additional increments of 17 GB up to the limit defined by the logical volume size. As a result, the storage administrator has no direct control over the actual capacity allocated to the volume by the system at any given point in time.
- ▶ During volume creation, or when a volume has been formatted, there is *zero physical capacity assigned to the volume*. As application writes accumulate to new areas of the volume, the physical capacity allocated to the volume will grow in increments of 17 GB and can ultimately reach the full logical volume size.
- ▶ Increasing the logical volume size *does not affect* the actual volume size.

Thinly provisioned storage pools

Whereas volumes are effectively thinly provisioned automatically by the system, storage pools can be defined by the storage administrator (when using the GUI) as either *regular* or *thinly provisioned*. Note that when using the Extended Command Line Interface (XCLI), there is no specific parameter to indicate thin provisioning for a storage pool. You indirectly and implicitly create a storage pool as thinly provisioned by specifying a pool soft size greater than its hard size.

With a regular pool, the “host-apparent” capacity is *guaranteed* to be equal to the physical capacity reserved for the pool. The total physical capacity *allocated to* the constituent individual volumes and collective snapshots at any given time within a regular pool will reflect the current usage by hosts, because the capacity is dynamically consumed as required. However, the remaining unallocated space within the pool remains reserved for the pool and cannot be used by other storage pools.

In contrast, a thinly provisioned storage pool is not fully backed by hard capacity, meaning that the entirety of the logical space within the pool cannot be physically provisioned unless the pool is transformed first into a regular pool. However, benefits can be realized when physical space consumption is less than the logical space assigned, because the amount of logical capacity assigned to the pool that is not covered by physical capacity is available for use by other storage pools.

When a storage pool is created using thin provisioning, that pool is defined in terms of both a soft size and a hard size independently, as opposed to a regular storage pool in which these sizes are by definition equivalent. Hard pool size and soft pool size are defined and used in the following ways.

Hard pool size

Hard pool size is the maximum actual capacity that can be used by all the volumes and snapshots in the pool.

Thin provisioning of the storage pool maximizes capacity utilization in the context of a group of volumes, wherein the aggregate “host-apparent,” or soft, capacity assigned to all volumes surpasses the underlying physical, or hard, capacity allocated to them. This utilization requires that the aggregate space available to be allocated to hosts within a thinly provisioned storage pool must be defined independently of the physical, or hard, space allocated within the system for that pool.

Thus, the storage pool hard size that is defined by the storage administrator limits the physical capacity that is available collectively to volumes and snapshots within a thinly provisioned storage pool. The aggregate space that is assignable to host operating systems is specified by the storage pool soft size.

Regular storage pools effectively segregate the hard space reserved for volumes from the hard space consumed by snapshots by limiting the soft space allocated to volumes; however, thinly provisioned storage pools permit the totality of the hard space to be consumed by volumes with no guarantee of preserving any hard space for snapshots. Logical volumes take precedence over snapshots and can be allowed to overwrite snapshots if necessary as hard space is consumed. The hard space that is allocated to the storage pool that is unused (or in other words, the incremental difference between the aggregate logical and actual *volume* sizes) can, however, be used by snapshots in the same storage pool.

Careful management is critical to prevent hard space for both logical volumes and snapshots from being exhausted. Ideally, hard capacity utilization must be maintained under a certain threshold by increasing the pool hard size as needed in advance.

Storage pools and snapshots:

- ▶ As discussed in “Storage pool relationships” on page 25, storage pools control when and which snapshots are deleted when there is insufficient space assigned within the pool for snapshots.
- ▶ The soft snapshot reserve capacity and the hard space allocated to the storage pool are consumed only as changes occur to the master volumes or the snapshots themselves, not as snapshots are created.
- ▶ See the *IBM XIV Storage System: Copy Services and Migration*, SG24-7759 for a more detailed explanation of snapshot deletion priority.

Soft pool size

Soft pool size is the maximum logical capacity that can be assigned to all the volumes and snapshots in the pool.

Thin provisioning is managed for each storage pool independently of all other storage pools:

- ▶ Regardless of any unused capacity that might reside in other storage pools, snapshots within a given storage pool will be deleted by the system according to corresponding snapshot pre-set priority if the hard pool size contains insufficient space to create an additional volume or increase the size of an existing volume. (Note that snapshots will actually only be deleted when a write occurs under those conditions, and not when allocating more space).
- ▶ As discussed previously, the storage administrator defines both the soft size and the hard size of thinly provisioned storage pools and allocates resources to volumes within a given storage pool without any limitations imposed by other storage pools.

The designation of a storage pool as a regular pool or a thinly provisioned pool can be dynamically changed by the storage administrator.

- ▶ When a regular pool needs to be converted to a thinly provisioned pool, the soft pool size parameter must be explicitly set in addition to the hard pool size, which will remain unchanged unless updated.

- When a thinly provisioned pool needs to be converted to a regular pool, the soft pool size is automatically reduced to match the current hard pool size. If the combined allocation of soft capacity for existing volumes in the pool exceeds the pool hard size, the storage pool cannot be converted. This situation can be resolved if individual volumes are selectively resized or deleted to reduce the soft space consumed.

System-level thin provisioning

The definitions of hard size and soft size naturally apply at the subsystem level as well, because by extension, it is necessary to permit the full system to be defined in terms of thin provisioning in order to achieve the full potential benefit previously described: The ability to defer deployment of additional capacity on an as-needed basis.

The IBM XIV Storage System's architecture allows the global system capacity to be defined in terms of both a *hard system size* and a soft system size. When thin provisioning is not activated at the system level, these two sizes are equal to the system's physical capacity.

With thin provisioning, these concepts have the following meanings.

Hard system size

The hard system size represents the physical disk capacity that is available within the IBM XIV Storage System. Obviously, the system's hard capacity is the upper limit of the aggregate hard capacity of all the volumes and snapshots and can only be increased by installing new hardware components in the form of individual modules (and associated disks) or groups of modules.

There are conditions that can *temporarily* reduce the system's hard limit. For further details, see 2.4.2, "Rebuild and redistribution" on page 38.

Soft system size

The soft system size is the total, "global," logical space available for all storage pools in the system. When the soft system size exceeds the hard system size, it is possible to logically provision more space than is physically available, thereby allowing the aggregate benefits of thin provisioning of storage pools and volumes to be realized at the system level.

The soft system size obviously limits the soft size of all volumes in the system and has the following attributes:

- It is not related to any direct system attribute and can be defined to be larger than the hard system size if thin provisioning is implemented. Note that the storage administrator cannot set the soft system size.

Storage pools: If the storage pools within the system are thinly provisioned, but the soft system size does not exceed the hard system size, the total system hard capacity cannot be filled until all storage pools are regularly provisioned. Therefore, it is best to define all storage pools in a non-thinly provisioned system as regular storage pools.

- The soft system size is a purely logical limit; however, you must exercise care when the soft system size is set to a value greater than the maximum potential hard system size. Obviously, it must be possible to upgrade the system's hard size to be equal to the soft size, so defining an unreasonably high system soft size can result in full capacity depletion. It is for this reason that defining the soft system size is not within the scope of the storage administrator role.

There are conditions that might *temporarily* reduce the system's soft limit. For further details, see 2.4.2, "Rebuild and redistribution" on page 38.

Thin provisioning conceptual examples

In order to further explain the thin provisioning principles previously discussed, it is helpful to examine the following basic examples, because they incorporate all of the concepts inherent to the IBM XIV Storage System's implementation of thin provisioning.

System-level thin provisioning conceptual example

Figure 2-5 depicts the incremental allocation of capacity to both a regular storage pool and a thinly provisioned storage pool within the context of the global system soft and hard sizes. This example assumes that the soft system size has been defined to exceed its hard size. The unallocated capacity shown within the system's soft and hard space is represented by a discontinuity in order to convey the full scope of both the logical and physical view of the system's capacity. Each increment in the diagram represents 17 GB of soft or hard capacity.

When a regular storage pool is defined, only one capacity is specified, and this amount is allocated to the storage pool from *both the hard and soft global capacity within the system*.

When a thinly provisioned storage pool is defined, both the soft and hard capacity limits for the storage pool must be specified, and these amounts are deducted from the system's global available soft and hard capacity, respectively.

In the next example we focus on the regular storage pool introduced in Figure 2-5.




Figure 2-5 Thin provisioning at the system level

Regular storage pool conceptual example

Next, Figure 2-6 represents a focused view of the regular storage pool that is shown in Figure 2-5 and depicts the division of both soft and hard capacity among volumes within the pool. Note that the regular pool is the same size (102 GB) in both diagrams.

First, consider Volume 1. Although Volume 1 is defined as 19,737,900 blocks (10 GB), the soft capacity allocated will nevertheless be comprised of the minimum number of 17 GB increments needed to meet or exceed the requested size in blocks, which is in this case only a single 17 GB increment of capacity. The host will, however, see exactly 19,737,900 blocks. When Volume 1 is created, the *system does not initially allocate any hard capacity*. At the moment that a host writes to Volume 1, even if it is just to initialize the volume, the system will allocate 17 GB of hard capacity. The hard capacity allocation of 17 GB for Volume 1 is illustrated in Figure 2-6, although clearly this allocation will never be fully utilized as long as the host-defined capacity remains only 10 GB.

Unlike Volume 1, Volume 2 has been defined in terms of gigabytes and has a soft capacity allocation of 34 GB, which is the amount that is reported to any hosts that are mapped to the volume. In addition, the hard capacity consumed by host writes has not yet exceeded the 17 GB threshold, and hence, the system has thus far *only allocated one increment of 17 GB hard capacity*. However, because the hard capacity and the soft capacity allocated to a regular storage pool are equal by definition, the remaining 17 GB of soft capacity assigned to Volume 2 is effectively preserved and will remain available within the pool's hard space until it is needed by Volume 2. In other words, because the pool's soft capacity does not exceed its hard capacity, there is no way to allocate soft capacity to effectively "overcommit" the available hard capacity.

The final reserved space within the regular storage pool shown in Figure 2-6 is dedicated for the snapshot usage. The diagram illustrates that the specified snapshot reserve capacity of 34 GB is effectively deducted from both the hard and soft space defined for the regular storage pool, thus guaranteeing that this space will be available for consumption collectively by the snapshots associated with the pool. Although snapshots consume space granularly at the partition level, as discussed in "Storage pool relationships" on page 25, the snapshot reserve capacity is still defined in increments of 17 GB.

The remaining 17 GB within the regular storage pool have not been allocated to either volumes or snapshots. Note that all soft capacity remaining in the pool is "backed" by hard capacity; the remaining unused soft capacity will always be less than or equal to the remaining unused hard capacity.




Figure 2-6 Volumes and snapshot reserve space within a regular storage pool

Thinly provisioned storage pool conceptual example

The thinly provisioned storage pool that was introduced in Figure 2-5 on page 31 is explored in detail in Figure 2-7. Note that the hard capacity and the soft capacity allocated to this pool are the same in both diagrams: 136 GB of soft capacity and 85 GB of hard capacity are allocated. Because the available soft capacity exceeds the available hard capacity by 51 GB, it is possible to thinly provision the volumes collectively by *up to 66.7%*, assuming that the snapshots are preserved and the remaining capacity within the pool is allocated to volumes.

Consider Volume 3 in Figure 2-7. The size of the volume is defined as 34 GB; however, less than 17 GB has been consumed by host writes, so only 17 GB of hard capacity have been allocated by the system. In comparison, Volume 4 is defined as 51 GB, but Volume 4 has consumed between 17 GB and 34 GB of hard capacity and therefore has been allocated 34 GB of hard space by the system. It is possible for either of these two volumes to require up to an additional 17 GB of hard capacity to become fully provisioned, and therefore, at least 34 GB of additional hard capacity must be allocated to this pool in anticipation of this requirement.

Finally, consider the 34 GB of snapshot reserve space depicted in Figure 2-7. If a new volume is defined in the unused 17 GB of soft space in the pool, or if either Volume 3 or Volume 4 requires additional capacity, the system will sacrifice the snapshot reserve space in order to give priority to the volume requirements. Normally, this scenario does not occur, because additional hard space must be allocated to the storage pool as the hard capacity utilization crosses certain thresholds.




Figure 2-7 Volumes and snapshot reserve space within a thinly provisioned Storage Pool

Depletion of hard capacity

Using thin provisioning creates the inherent danger of exhausting the available physical capacity. If the soft system size exceeds the hard system size, the potential exists for applications to fully deplete the available physical capacity.

Important: Upgrading the system beyond the full 15 modules in a single frame is currently not supported.

Snapshot deletion

As mentioned previously, snapshots in regular storage pools can be automatically deleted by the system in order to provide space for newer snapshots, or in the case of thinly provisioned pools, to permit more physical space for volumes.

Volume locking

If more hard capacity is still required after all the snapshots in a thinly provisioned storage pool have been deleted, *all the volumes in the storage pool are locked, thereby preventing any additional consumption of hard capacity*. There are two possible behaviors for a locked volume: read only (the default behavior) or no I/O at all.

Important: Volume locking prevents writes to all volumes in the storage pool.

It is very important to note that thin provisioning implementation in the IBM XIV Storage System manages space allocation within each storage pool, so that hard capacity depletion in one storage pool will never affect the hard capacity available to another storage pool. There are both advantages and disadvantages:

- ▶ Because storage pools are independent, thin provisioning volume locking on one storage pool never cascades into another storage pool.
- ▶ Hard capacity cannot be reused across storage pools, even if a certain storage pool has free hard capacity available, which can lead to a situation where volumes are locked due to the depletion of hard capacity in one storage pool, while there is available capacity in another storage pool. Of course, it is still possible for the storage administrator to intervene in order to redistribute hard capacity.

2.4 Reliability, availability, and serviceability

The IBM XIV Storage System's unique modular design and logical topology fundamentally differentiate it from traditional monolithic systems, and this architectural divergence extends to the exceptional reliability, availability, and serviceability aspects of the system. In addition, the IBM XIV Storage System incorporates autonomic, proactive monitoring and self-healing features that are capable of not only transparently and automatically restoring the system to full redundancy within minutes of a hardware failure, but also taking preventive measures to preserve data redundancy even before a component malfunction actually occurs.

For further reading about the IBM XIV Storage System's parallel modular architecture, see 2.2, "Parallelism" on page 16.

2.4.1 Resilient architecture

As with any enterprise class system, redundancy pervades every aspect of the IBM XIV Storage System, including the hardware, internal operating environment, and the data itself. However, the design elements, including the distribution of volumes across the whole of the system, in combination with the loosely coupled relationship between the underlying hardware and software elements, empower the IBM XIV Storage System to realize unprecedented resiliency. The resiliency of the architecture encompasses not only high availability, but also excellent maintainability, serviceability, and performance under non-ideal conditions resulting from planned or unplanned changes to the internal hardware infrastructure, such as the loss of a module.

Availability

The IBM XIV Storage System maximizes operational availability and minimizes the degradation of performance associated with non disruptive planned and unplanned events, while providing for the capability to preserve the data to the fullest extent possible in the event of a disaster.

High reliability

The IBM XIV Storage System not only withstands individual component failures by quickly and efficiently reinstating full data redundancy, but also automatically monitors and phases out individual components before data redundancy is compromised. We discuss this topic in detail in "Proactive phase-out and self-healing mechanisms" on page 44. The collective high reliability provisions incorporated within the system constitute multiple layers of protection from unplanned outages and minimize the possibility of related service actions.

Maintenance freedom

While the potential for unplanned outages and associated corrective service actions are mitigated by the reliability attributes inherent to the system design, the IBM XIV Storage System's autonomic features also minimize the need for storage administrators to conduct non-preventative maintenance activities that are purely reactive in nature, by adapting to potential issues before they are manifested as a component failure. The continually restored redundancy in conjunction with the self-healing attributes of the system effectively enable maintenance activities to be decoupled from the instigating event (such as a component failure or malfunction) and safely carried out according to a predefined schedule. In addition to the system's diagnostic monitoring and autonomic maintenance, the proactive and systematic, rather than purely reactive, approach to maintenance is augmented, because the entirety of the logical topology is continually preserved, optimized, and balanced according to the physical state of the system.

The modular system design also expedites the installation of any replacement or upgraded components, while the automatic, transparent data redistribution across all resources eliminates the downtime, even in the context of individual volumes, associated with these critical activities.

High availability

The rapid restoration of redundant data across all available drives and modules in the system during hardware failures, and the equilibrium resulting from the automatic redistribution of data across all newly installed hardware, are fundamental characteristics of the IBM XIV Storage System architecture that minimize exposure to cascading failures and the associated loss of access to data.

Consistent performance

The IBM XIV Storage System is capable of adapting to the loss of an individual drive or module efficiently and with relatively minor impact compared to monolithic architectures. While traditional monolithic systems employ an N+1 hardware redundancy scheme, the IBM XIV Storage System harnesses the resiliency of the grid topology, not only in terms of the ability to sustain a component failure, but also by maximizing consistency and transparency from the perspective of attached hosts. The potential impact of a component failure is vastly reduced, because each module in the system is responsible for a relatively small percentage of the system's operation. Simply put, a controller failure in a typical N+1 system likely results in a dramatic (up to 50%) reduction of available cache, processing power, and internal bandwidth, whereas the loss of a module in the IBM XIV Storage System translates to only 1/15th of the system resources and does not compromise performance nearly as much as the same failure with a typical architecture.

Additionally, the IBM XIV Storage System incorporates innovative provisions to mitigate isolated disk-level performance anomalies through *redundancy-supported reaction*, which is discussed in "Redundancy-supported reaction" on page 45, and flexible handling of dirty data, which is discussed in "Flexible handling of dirty data" on page 45.

Disaster recovery

Enterprise class environments must account for the possibility of the loss of both the system and all of the data as a result of a disaster. The IBM XIV Storage System includes the provision for remote mirror functionality as a fundamental component of the overall disaster recovery strategy.

Write path redundancy

Data arriving from the hosts is temporarily placed in two separate caches before it is permanently written to disk drives located in separate modules. This design guarantees that the data is always protected against possible failure of individual modules, even before the data has been written to the disk drives.

Figure 2-8 illustrates the path taken by a write request as it travels through the system. The diagram is intended to be viewed as a conceptual topology, so do not interpret the specific numbers of connections and so forth as literal depictions. Also, for purposes of this discussion, the Interface Modules are depicted on a separate level from the Data Modules. However, in reality the Interface Modules also function as Data Modules. The following numbers correspond to the numbers in Figure 2-8:

1. A host sends a write request to the system. Any of the Interface Modules that are connected to the host can service the request, because the modules work in an active-active capacity. Note that the IBM XIV Storage System does not load balance the requests itself. Load balancing must be implemented by storage administrators to equally distribute the host requests among all Interface Modules.
2. The Interface Module uses the system configuration information to determine the location of the primary module that houses the referenced data, which can be either an Interface Module, including the Interface Module that received the write request, or a Data Module. The data is written only to the local cache of the primary module.
3. The primary module uses the system configuration information to determine the location of the secondary module that houses the copy of the referenced data. Again, this module can be either an Interface Module or a Data Module, but it will not be the same as the primary module. The data is redundantly written to the local cache of the secondary module.

After the data is written to cache in both the primary and secondary modules, the host receives an acknowledgement that the I/O is complete, which occurs independently of copies of either cached, or dirty, data being destaged to physical disk.




Figure 2-8 Write path

System quiesce and graceful shutdown

When an event occurs that compromises both sources of power to the IBM XIV Storage System's redundant uninterruptible power supplies, the system executes the graceful shutdown sequence. Full battery power is guaranteed during this event, because the system monitors available battery charge at all times and takes proactive measures to prevent the possibility of conducting write operations when battery conditions are non-optimal.

Due to the IBM XIV Storage System's grid topology, a system quiesce event essentially entails the graceful shutdown of all modules within the system. Each module can be thought of as an independent entity that is responsible for managing the destaging of "dirty" data, that is, written data that has not yet been destaged to physical disk. The dirty data within each module consists of equal parts primary and secondary copies of data, *but will never contain both primary and secondary copies of the same data.*

Write cache protection

Each module in the IBM XIV Storage System contains a local, independent space reserved for caching operations within its system memory.

Each module contains at least 8 GB of high speed volatile *memory* (a total of 120 GB), from which 5.5 GB (and 82.5 GB overall) is dedicated for caching data.

Destage: The system does *not* contain non-volatile memory space that is reserved for write operations. However, the close proximity of the cache and the drives, in conjunction with the enforcement of an upper limit for dirty, or non-destaged, data on a per-drive basis, ensures that the full destage will occur while operating under battery power.

Graceful shutdown sequence

The system executes the graceful shutdown sequence under either of these conditions:

- ▶ The battery charge remaining in two or more uninterruptible power supplies is below a certain threshold, which is conservatively predetermined in order to provide adequate time for the system to fully destage all dirty data from cache.
- ▶ The system detects the loss of external power for more than 30 seconds.

Power on sequence

Upon startup, the system will verify that the battery charge levels in all uninterruptible power supplies exceed the threshold necessary to guarantee that a graceful shutdown can occur. If the charge level is inadequate, the system will halt the startup process until the charge level has exceeded the minimum required threshold.

Battery: If the battery charge is inadequate, the system will remain fully locked until the battery charge has exceeded the necessary threshold to safely resume I/O activity.

2.4.2 Rebuild and redistribution

As discussed in "Data distribution algorithms" on page 18, the IBM XIV Storage System dynamically maintains the pseudo-random distribution of data across all modules and disks while ensuring that two copies of data exist at all times *when the system reports Full Redundancy*. Obviously, when there is a change to the hardware infrastructure as a result of a failed component, data must be restored to redundancy and distributed, or when a component is added, or *phased-in*, a new data distribution must accommodate the change.

Goal distribution

The process of achieving a new goal distribution while simultaneously restoring data redundancy due to the loss of a disk or module is known as a *rebuild*. Because a rebuild occurs as a result of a component failure that compromises full data redundancy, there is a period during which the *non-redundant data* is both restored to full redundancy and homogeneously redistributed over the remaining disks.

The process of achieving a new goal distribution (only occurring when redundancy exists) is known as a *redistribution*, during which all data in the system (including both primary and secondary copies) is redistributed, when it is a result of the following events:

- ▶ The replacement of a failed disk or module following a rebuild, also known as a “phase-in”
- ▶ When one or more modules are added to the system, known as a “scale out” upgrade

Following any of these occurrences, the IBM XIV Storage System immediately initiates the following sequence of events:

1. The IBM XIV Storage System distribution algorithms calculate which partitions must be relocated and copied based on the distribution table that is described in 2.2.2, “Software parallelism” on page 17. The resultant distribution table is known as the *goal distribution*.
2. The Data Modules and Interface Modules begin concurrently redistributing and copying (in the case of a rebuild) the partitions according to the goal distribution:
 - This process occurs in a parallel, any-to-any fashion concurrently among all modules and drives in the background, with complete host transparency.
 - The priority associated with achieving the new goal distribution is internally determined by the system. The priority cannot be adjusted by the storage administrator:
 - Rebuilds have the highest priority; however, the transactional load is homogeneously distributed over all the remaining disks in the system resulting in a very low density of system-generated transactions.
 - Phase-outs (caused by the XIV technician removing and replacing a failed module) have lower priority than rebuilds, because at least two copies of all data exist at all times during the phase-out.
 - Redistributions have the lowest priority, because there is neither a lack of data redundancy nor has the system detected the potential for an impending failure.
3. The system reports Full Redundancy after the goal distribution has been met.

Following the completion of goal distribution resulting from a rebuild or phase-out, a subsequent redistribution must occur when the system hardware is fully restored through a phase-in.

Goal distribution: The goal distribution is transparent to storage administrators and cannot be changed. In addition, the goal distribution has many determinants depending on the precise state of the system.

Important: Never perform a phase-in to replace a failed disk or module until after the rebuild process has completed. These operations must be performed by the IBM XIV technician.

Preserving data redundancy

Whereas conventional storage systems maintain a static relationship between RAID arrays and logical volumes by preserving data redundancy only across a subset of disks that are defined in the context of a particular RAID array, the IBM XIV Storage System dynamically and fluidly restores redundancy and equilibrium across all disks and modules in the system during the rebuild and phase-out operations. See “Logical volume layout on physical disks” on page 22 for a detailed discussion of the low-level virtualization of logical volumes within the IBM XIV Storage System.

Benefits

The proactive phase-out of non-optimal hardware through autonomic monitoring and the modules’ cognizance of the virtualization between the logical volumes and physical disks yield unprecedented efficiency, transparency, and reliability of data preservation actions, encompassing both rebuilds and phase-outs:

- ▶ The rebuild of data is many times faster than conventional RAID array rebuilds and can complete in a short period of time for a fully provisioned system, because the redistribution workload spans all drives in the system resulting in very low transactional density:
 - Statistically, the chance of exposure to data loss or a cascading hardware failure, which occurs when corrective actions in response to the original failure result in a subsequent failure, is minimized due to both the brevity of the rebuild action and the low density of access on any given disk.

Rebuilding conventional RAID arrays can take many hours to complete, depending on the type of the array, the number of drives, and the ongoing host-generated transactions to the array.

- The rebuild process can complete 25% to 50% more quickly for systems that are not fully provisioned, which equates to a rebuild completion in as little as 15 minutes.
- ▶ The system relocates only real data, as opposed to rebuilding the entire array, which consists of complete disk images that often include unused space, vastly reducing the potential number of transactions that must occur.

Conventional RAID array rebuilds can place many times the normal transactional load on the disks and substantially reduce effective host performance.

- ▶ The number of drives participating in the rebuild is about 20 times greater than in most average-sized conventional RAID arrays, and by comparison, the array rebuild workload is greatly dissipated, greatly reducing the relative impact on host performance.
- ▶ Whereas standard dedicated spare disks utilized during a conventional RAID array rebuild might not be globally accessible to all arrays in the system, the IBM XIV Storage System maintains universally accessible reserve space on all disks in the system, as discussed in “Global spare capacity” on page 23.
- ▶ Because the system maintains access density equilibrium, hot-spots are statistically eliminated, which reduces the chances of isolated workload-induced failures.
- ▶ The system-wide goal distribution alleviates localized drive stress and associated additional heat generation, which can significantly increase the probability of a double drive failure during the rebuild of a RAID array in conventional subsystems.
- ▶ Modules intelligently send information to each other directly. There is no need for a centralized supervising controller to read information from one disk module and write to another disk module.
- ▶ All disks are monitored for errors, poor performance, or other signs that might indicate that a full or partial failure is impending.

Dedicated spare disks in conventional RAID arrays are inactive, and therefore, unproven and potentially unmonitored, increasing the possibility for a second failure during an array rebuild.

Rebuild examples

When the full redundancy of data is compromised due to a module failure, as depicted in Figure 2-9, the system immediately identifies the non-redundant partitions and begins the rebuild process. Because none of the disks within a given module contain the secondary copies of data residing on any of the disks in the module, the secondary copies are read from the remaining modules in the system. Therefore, during a rebuild resulting from a module failure, there will be concurrently 168 disks (180 disks in the system minus 12 disks in a module) reading, and 168 disks writing, as is conceptually illustrated in Figure 2-9.




Figure 2-9 Non-redundant group of partitions following module failure

Figure 2-10 depicts a denser population of redundant partitions for both volumes A and B, thus representing the completion of a new goal distribution, as compared to Figure 2-9, which contains the same number of redundant partitions for both volumes distributed less densely over the original number of modules and drives.

Finally, consider the case of a single disk failure occurring in an otherwise healthy system (no existing phased-out or failed hardware). During the subsequent rebuild, there will be only 168 disks reading, because there is no non-redundant data residing on the other disks within the same module as the failed disk. Concurrently, there will be 179 disks writing in order to preserve full data distribution.

Rebuild: Figure 2-9 and Figure 2-10 conceptually illustrate the rebuild process resulting from a failed module. *The diagrams are not intended to depict in any way the specific placement of partitions within a real system, nor do they literally depict the number of modules in a real system.*




Figure 2-10 Performing a new goal distribution following module failure

Transient soft and hard system size

The capacity allocation that is consumed for purposes of either restoring non-redundant data during a rebuild, or creating a tertiary copy during a phase-out, will be sourced based on availability, with the following precedence:

1. *Unallocated system hard capacity*: The system might consume hard capacity that was not assigned to any storage pools at the time of the failure.
2. *Unallocated storage pool hard capacity*: The hard capacity of storage pools that is not assigned to any existing volumes or consumed by snapshots, as measured before the failure, is unallocated hard capacity. For details about the topic of storage pool sizes, see “Thinly provisioned storage pools” on page 28. Do not confuse this unallocated storage pool hard capacity with “unconsumed” capacity, which is unwritten hard space allocated to volumes.
3. *Reserve spare capacity*: As discussed previously, the system reserves enough capacity to sustain the consecutive, non-concurrent failure of three drives and an entire module before replacement hardware must be phased in to ensure that data redundancy can be restored during subsequent hardware failures.

In the event that sufficient unallocated hard capacity is available, the system will withhold allocating reserve spare space to complete the rebuild or phase-out process in order to provide additional protection. As a result, it is possible for the system to report a maximum soft size that is temporarily less than the allocated soft capacity. The soft and hard system sizes will not revert to the original values until a replacement disk or module is phased-in, and the resultant redistribution completes.

Important: Although it is possible to resize or create volumes, snapshots, or storage pools while a rebuild is underway, we strongly discourage these activities until the system has completed the rebuild process and restored full data redundancy.

Redistribution

The IBM XIV Storage System homogeneously redistributes all data across all disks whenever new disks or modules are introduced or phased in to the system. This redistribution process is not equivalent to the “striping volumes on all disks” employed in traditional systems:

- ▶ Both conventional RAID striping, as well as the data distribution, fully incorporate all spindles when the hardware configuration remains static; however, when new capacity is added and new volumes are allocated, ordinary RAID striping algorithms do not intelligently redistribute data to preserve equilibrium for all volumes through the pseudo-random distribution of data, which is described in 2.2.2, “Software parallelism” on page 17.
- ▶ Thus, the IBM XIV Storage System employs dynamic volume-level virtualization, obviating the need for ongoing manual volume layout planning.

The redistribution process is triggered by the “phase-in” of a new drive or module and differs from a rebuild or phase-out in that:

- ▶ The system does not need to create secondary copies of data to reinstate or preserve full data redundancy.
- ▶ The distribution density, or the concentration of data on each physical disk, decreases instead of increasing.
- ▶ The redistribution of data performs differently, because the concentration of write activity on the new hardware resource is the bottleneck:
 - When a replacement module is phased-in, there will be concurrently 168 disks reading and 12 disks writing, and thus the time to completion is limited by the throughput of the replacement module. Also, the read access density on the existing disks will be extremely low, guaranteeing extremely low impact on host performance during the process.
 - When a replacement disk is phased-in, there will be concurrently 179 disks reading and only one disk writing. In this case, the replacement drive obviously limits the achievable throughput of the redistribution. Again, the impact on host transactions is extremely small, or insignificant.

2.4.3 Minimized exposure

This section describes other features that contribute to the IBM XIV Storage System reliability and availability.

Disaster recovery

All high availability SAN implementations must account for the contingency of data recovery and business continuance following a disaster, as defined by the organization’s recovery point and recovery time objectives. The provision within the IBM XIV Storage System to efficiently and flexibly create nearly unlimited snapshots, coupled with the ability to define Consistency Groups of logical volumes, constitutes integral elements of the data preservation strategy. In addition, the IBM XIV Storage System’s data mirroring functionality facilitates excellent potential recovery point and recovery time objectives as a central element of the full disaster recovery plan.

Proactive phase-out and self-healing mechanisms

The IBM XIV Storage System can seamlessly restore data redundancy with minimal data migration and overhead. A further enhancement to the level of reliability standards attained by the IBM XIV Storage System entails self-diagnosis and early detection mechanisms that autonomically phase out components before the probability of a failure increases beyond a certain point.

In real systems, the failure rate is not constant with time, but rather increases with service life and duty cycle. By actively gathering component statistics to monitor this trend, the system ensures that components will not operate under conditions beyond an acceptable threshold of reliability and performance. Thus, the IBM XIV Storage System's self-healing mechanisms dramatically increase the already exceptional level of availability of the system, because they virtually preclude the possibility of data redundancy from ever being compromised along with the associated danger, however unlikely, of subsequent failures during the rebuild process.

The autonomic attributes of the IBM XIV Storage System cumulatively impart an enormous benefit to not only the reliability of the system, but also the overall availability, by augmenting the maintainability and serviceability aspects of the system. Both the monetary and time demands associated with maintenance activities, or in other words, the total cost of ownership (TCO), are effectively minimized by reducing reactive service actions and enhancing the potential scope of proactive maintenance policies.

Disk scrubbing

The IBM XIV Storage System maintains a series of scrubbing algorithms that run as background processes concurrently and independently scanning multiple media locations within the system in order to maintain the integrity of the redundantly stored data. This continuous checking enables the early detection of possible data corruption, alerting the system to take corrective action to restore the data integrity before errors can manifest themselves from the host perspective. Thus, redundancy is not only implemented as part of the basic architecture of the system, but it is also continually monitored and restored as required. In summary, the data scrubbing process has the following attributes:

- ▶ Verifies the integrity and redundancy of stored data
- ▶ Enables early detection of errors and early recovery of redundancy
- ▶ Runs as a set of background processes on all disks in parallel
- ▶ Checks whether data can be read from partitions and verifies data integrity by employing checksums
- ▶ Examines a partition approximately every two seconds

Enhanced monitoring and disk diagnostics

The IBM XIV Storage System continuously monitors the performance level and reliability standards of each disk drive within the system, using an enhanced implementation of Self-Monitoring, Analysis and Reporting Technology (SMART) tools. As typically implemented in the storage industry, SMART tools simply indicate whether certain thresholds have been exceeded, thereby alerting that a disk is at risk for failure and thus needs to be replaced.

However, as implemented in IBM XIV Storage System, the SMART diagnostic tools, coupled with intelligent analysis and low tolerance thresholds, provide an even greater level of refinement of the disk behavior diagnostics and the performance and reliability driven reaction. For instance, the IBM XIV Storage System measures the specific values of parameters including, but not limited to:

- ▶ *Reallocated sector count*: If the disk encounters a read or write verification error, it designates the affected sector as “reallocated” and relocates the data to a reserved area of spare space on the disk. Note that this spare space is a parameter of the drive itself and is *not* related in any way to the system reserve spare capacity that is described in “Global spare capacity” on page 23. The IBM XIV Storage System initiates phase-out at a much lower count than the manufacturer specifies.
- ▶ *Disk temperature*: The disk temperature is a critical factor that contributes to premature drive failure and is constantly monitored by the system.
- ▶ *Raw read error*: The raw read error count provides an indication of the condition of the magnetic surface of the disk platters and is carefully monitored by the system to ensure the integrity of the magnetic media itself.
- ▶ *Spin-up time*: The spin-up time is a measure of the average time that is required for a spindle to accelerate from zero to 7 200 rpm. The IBM XIV Storage System recognizes abnormal spin-up time as a potential indicator of an impending mechanical failure.

Likewise, for additional early warning signs, the IBM XIV Storage System continually monitors other aspects of disk-initiated behavior, such as spontaneous reset or unusually long latencies. The system intelligently analyzes this information in order to reach crucial decisions concerning disk deactivation and phase-out. The parameters involved in these decisions allow for a very sensitive analysis of the disk health and performance.

Redundancy-supported reaction

The IBM XIV Storage System incorporates *redundancy-supported reaction*, which is the provision to exploit the distributed redundant data scheme by intelligently redirecting reads to the secondary copies of data, thereby extending the system’s tolerance of above average disk service time when accessing primary data locations. The system will reinstate reads from the primary data copy when the transient degradation of the disk service time has subsided. Of course, a redundancy-supported reaction itself might be triggered by an underlying potential disk error that will ultimately be managed autonomically by the system according to the severity of the exposure, as determined by ongoing disk monitoring.

Flexible handling of dirty data

In a similar manner to the redundancy-supported reaction for read activity, the IBM XIV Storage System can also make convenient use of its redundant architecture in order to consistently maintain write performance. Because intensive write activity directed to any given volume is distributed across all modules and drives in the system, and the cache is independently managed within each module, the system is able to tolerate sustained write activity to an under-performing drive by effectively maintaining a considerable amount of “dirty,” or unwritten, data in cache, thus potentially circumventing any performance degradation resulting from the transient, anomalous service time of a given disk drive.

Non-disruptive code load

Non-disruptive code load (NDCL) enables upgrades to the IBM XIV Storage System software from a current version (starting with Version 10.1) to a later version without disrupting the application service.

The code upgrade is run on all modules in parallel and the process is fast enough to minimize impact on hosts applications.

No data migration or rebuild process is allowed during the upgrade. Mirroring, if any, will be suspended during the upgrade and automatically reactivated upon completion.

Storage management operations are also not allowed during the upgrade, although the status of the system and upgrade progress can be queried. It is also possible to cancel the upgrade process up to a point of no return.

Note that the NDCL does not apply to specific components firmware upgrades (for instance, module BIOS, HBA firmware). Those require a phase in / phase out process of the impacted modules.



IBM XIV physical architecture, components, and planning

In this chapter we describe the hardware architecture of the IBM XIV Storage System. We present the physical components that make up the IBM XIV Storage System, such as the system rack, Interface Modules, Data Modules, Management Module, disk drives, network switches, and power distribution devices.

We also include an overview of the planning aspects required before and after deployment of an IBM XIV Storage System.

3.1 IBM XIV Storage System models 2810-A14 and 2812-A14

The IBM XIV Storage System, seen in Figure 3-1, is designed to be a scalable enterprise storage system based upon a grid array of hardware components. The architecture offers the highest performance through maximized utilization of all disks, true distributed cache implementation, coupled with high bandwidth. It also offers superior reliability through distributed architecture, redundant components, self-monitoring, and self-healing attributes.



Figure 3-1 IBM XIV Storage System front and rear views

The IBM XIV Storage System family consists of two machine types, the 2810-A14 and the 2812-A14. The 2812 machine type comes standard with a 3 year manufacturer warranty. The 2810 machine type is delivered with a 1 year standard warranty.

All XIV hardware components come pre-installed in a standard 19" data center class rack. Data and Interface Modules provide the processing, caching, and storing of data. In addition, Interface Modules provide Fibre Channel and iSCSI host connectivity adapters. Each module contains twelve 1 TB or 2 TB Serial Advanced Technology Attachment (SATA) disk drives.

Important: 1 TB and 2 TB drives cannot be intermixed within the same XIV system. IBM XIV Storage System Software V10.2.1, or later, is required to support the 2 TB drives.

Installed near the middle of the rack, just above module 6, are a maintenance module used to service the machine, two Ethernet switches (1 Gbps, 48 ports each), and a Redundant Power Supply (RPS) for the switches.

The Ethernet switches form the basis of an internal Gigabit Ethernet network that provides redundant communications paths between each module. This grid network ensures communication between all modules even if one of the switches or a cable connection fails. Furthermore, this grid network provides the capabilities for parallelism and the execution of a data distribution algorithm that contributes to the excellent performance of the IBM XIV Storage System.

At the bottom of the rack, an Uninterruptible Power Supply (UPS) module complex, which is made up of three redundant UPS units, is installed and provides power to the various system components.

All cabling between the modules and switches, including the internal power connections, are fully redundant using two sets of cables. All cables use industry standard plugs.

Figure 3-2 summarizes the main hardware components of the IBM XIV Storage System 2810-A14 and 2812-A14.




Figure 3-2 Hardware overview - Machine type 2810-2812 - model A14

3.1.1 Fully populated configurations

A fully populated rack contains 9 Data Modules and 6 Interface Modules for a total of 15 modules. The total usable capacity is 79 TB (180 TB raw) for the complete system when equipped with 1 TB drives, or 161 TB usable capacity (369 TB raw capacity) when fully populated with the 2 TB drives.

3.1.2 Partially populated configurations

The IBM 2810-A14 and IBM 2812-A14 are also available in partially populated configurations of 6 or 9, 10, 11, 12, 13, 14, 15 modules, allowing for more granularity of capacity options. Additional modules can be added as storage needs increase to achieve configurations with a total of 9, 10, 11, 12, 13, 14, or 15 modules. Remember that you cannot intermix 1 TB and 2 TB drives within the same IBM XIV Storage System.

These partially configured racks are available in usable capacities ranging from 27 TB to 79 TB when using 1 TB disks or from 55 TB to 161 TB when using 2 TB disks.

Certain partial rack configurations do not utilize all host attachment interface ports even though they might be physically present. The interface ports will be activated automatically as additional modules are added to the system.

Details on these configuration options and the various capacities, drives, ports, and memory are provided in Figure 3-3.

Total Modules	6	9	10	11	12	13	14	15
Useable Capacity (TB)								
- Using 1 TB disks	27	43	50	54	61	66	73	79
- Using 2 TB disks	55	87	102	111	125	134	149	161
Memory (GB)								
- Using 1 TB disks	48	72	80	88	96	104	112	120
- Using 2 TB disks	96	144	160	176	192	208	224	240
Interface Modules (physically installed)	3	6	6	6	6	6	6	6
Interface Modules (with active interfaces)	3	4	4	5	5	6	6	6
Data Modules	3	3	4	5	6	7	8	9
Disk Drives	72	108	120	132	144	156	168	180
Fibre Channel Ports	8	16	16	20	20	24	24	24
iSCSI Ports	0	4	4	6	6	6	6	6

Figure 3-3 XIV partial configurations using 1 TB or 2 TB drives

For more information about *usable capacity*, see 2.3, “Full storage virtualization” on page 18.

3.2 IBM XIV hardware components

The system architecture of the IBM XIV Storage System is designed specifically upon off-the-shelf components¹ that are not dependent upon specifically designed hardware or proprietary technology. This design in architecture is optimized for ease of use such that, as newer and higher performing components are made available in the marketplace, development is able to incorporate this newer technology into the base system design at a faster pace than was traditionally possible. In the sections that follow, we explore these base components in further detail.

¹ With the exception of the Automatic Transfer Switch (ATS)

At a minimum, for all configurations, the following components are supplied with the system:

- ▶ 3 UPS Units
- ▶ 2 Ethernet Switches
- ▶ 1 Ethernet Switch Redundant Power Supply
- ▶ 1 Maintenance Module
- ▶ 1 Automatic Transfer Switch (ATS)
- ▶ 1 Modem
- ▶ 8-24 Fibre Channel Ports
- ▶ 0-6 iSCSI Ports (iSCSI ports are not supplied in configurations with only 6 Modules)
- ▶ Complete set of internal cabling

Hardware: For the same reason that the system is not dependent on specially developed parts, there might be differences in the actual hardware components that are used in your particular system compared with those components described next.

3.2.1 Redundancy

In this section we discuss features that contribute to redundancy.

Hardware redundancy

The IBM XIV Storage System hardware is redundant to prevent machine outage when any single hardware component is failing. The combination of hardware redundancy with the logical architecture that is described in Chapter 2, “IBM XIV logical architecture and concepts” on page 13 makes the IBM XIV Storage System extremely resilient to outages.

Power redundancy

To prevent the complete rack or single components from failing due to power problems, all power components in the IBM XIV Storage System are redundant:

- ▶ To ensure redundant power availability at the rack level, a device must be present to enable switching from one power source to another available power source, which is performed by an Automatic Transfer Switch (ATS). In the case of a failing external power source, this switch transfers the load to the remaining external power source without interrupting the system.
- ▶ Each module has two independent power supplies. During normal operation, both power supplies operate on half of the maximal load. If one power supply fails, the operational power supply can take over, and the module continues its operation without any noticeable impact. After the failing power supply is replaced, the power load balancing is restored.
- ▶ The two Gigabit-Ethernet switches are powered by the internal UPSes as well as the PowerConnect RPS-600 Redundant Power Bank to eliminate the A/C power supply as a single point of failure.

Switch and interconnect redundancy

The IBM XIV Storage System internal network is built around two Ethernet switches which are interconnected for redundancy. Each module (Data and Interface) has connections to both switches to eliminate any failing hardware component within the network from becoming a single point of failure.

3.2.2 Rack and UPS modules

This section describes the hardware rack and UPS modules.

The IBM XIV Storage System hardware components are installed in a 24U rack. Adequate space is provided to house all components and to properly route all cables. The rack door and/or side panels can be locked with a key to prevent unauthorized access to the installed components.

An illustration is shown in Figure 3-4.



Figure 3-4 IBM XIV Model 2810 & 2812

For detailed dimensions, clearances, and the weight of each rack and its components, see 3.3.2, “Physical site planning” on page 67.

UPS module complex

The Uninterruptible Power Supply (UPS) module complex consists of three UPS units. The UPS units work together to provide a “cache battery” function to sustain the IBM XIV Storage System in the event of power fluctuations or loss. The UPS complex will sustain the system for up to 30 seconds in the event of a temporary loss of external power. In the event of an extended external power failure or outage which is greater than 30 seconds, the UPS module complex maintains battery power long enough to allow a safe and orderly shutdown of the IBM XIV Storage System. The complex can sustain the failure of one UPS unit and still provide adequate protection against external power outages.

The UPS complex is dedicated to supporting the components inside the XIV rack. No external equipment can be plugged into the UPS units, because doing so might jeopardize the complex’s ability to provide adequate power to allow the system to complete the shutdown process during a power loss event. The UPS complex is not a substitute for traditional power

conditioning or sustaining equipment typically found in a data center environment and is intended to be used in conjunction with existing equipment.

Figure 3-5 shows an illustration of one UPS module.




Figure 3-5 UPS module

The three UPS modules are located at the bottom of the rack. Each UPS is 3U in height. Each UPS has an output of 6 kVA to supply power to all other components in the system. The UPS module complex design allows proactive detection of temporary power problems and can compensate for them without affecting the system. In the case of a complete power outage, integrated batteries continue to supply power to the entire system. The batteries are designed to last long enough for a safe and orderly shutdown of the IBM XIV Storage System.

Warning: Do *not* power off the XIV using the UPS power button because this can result in the loss of data and system configuration information. Powering the system off must be done solely from either the XIV GUI or the XCLI.

Automatic Transfer Switch (ATS)

The Automatic Transfer Switch (ATS) seen in Figure 3-7 supplies power to all three Uninterruptible Power Supplies (UPS) and to the Maintenance Module. In case of a power problem on one line, the ATS reorganizes the power and switches to the other line. The operational components take over the load from the failing power source or power supply. This rearrangement is performed by the ATS in a seamless manner such that the system operation continues without any application impact.

The ATS is capable of being monitored by the system if the ATS monitoring cables are installed and the system is running software Version 10.2.1 or higher. The ATS will generate system event messages in case of problems. The status of the ATS can also be viewed with the XCLI command `ats_list` as shown in Figure 3-6.

```

XCLI Session - XIV 6000037 PFE
connecting...
XIV 6000037 PFE>>ats_list
Component ID  Status  Currently Functioning  Input Line 1  Input Line 2  Output 30A #1  Output 30A #2  Output 30A #3  Output 10A  Firmware Version  3-Phase
1-ATS:1       OK       yes                  yes          yes          no           yes          no           yes       4              no
XIV 6000037 PFE>>

```

Figure 3-6 Output of XCLI command ats_list

The ATS is available as a single-phase power or three-phase power option. Multiple line cord configurations are available depending on the ATS option and your geography. For the appropriate line cord selection, see the *IBM XIV Storage System (Types 2810 and 2812) Model A14 (Gen2) Introduction and Planning Guide for Customer Configuration*, GA52-1327-07.

Internal UPS: Although the system is protected by an uninterruptible power supply (UPS) for internal usage, you can reduce the risk of a power outage if you connect the system to an external UPS, a backup generator, or both.




Figure 3-7 Automatic Transfer Switch ATS

Two separate external main power sources supply power to the ATS. The following power options are available:

- ▶ Single-phase power ATS:
 - Two 60 A, 200-240 V ac, single-phase, two-pole, line-line ground female receptacles, each connected to a separate power source
- ▶ Three-phase power ATS:
 - Two 60 A, 200-240 V ac, three-phase receptacles, each connected to a separate power source
 - Two 30 A, 200-240 V ac, three-phase receptacles, each connected to a separate power source

3.2.3 Data Modules and Interface Modules

The hardware of the Interface Modules and Data Modules is based on an Intel server platform optimized for data storage services. A module is 87.9 mm (3.46 inches) (2U) tall, 483 mm (19 inches) wide, and 707 mm (27.8 inches) deep. The weight depends on configuration and type (Data Module or Interface Module) and is a maximum of 30 kg (66.14 lbs).

Figure 3-8 shows a representation of a module in perspective.

Interface Modules contain dual-low voltage CPU and are a compatible replacement for previous single-CPU Interface Modules. The dual-low voltage CPUs provide additional computing power to the interface functions of the IBM XIV Storage System while reducing overall power consumption.

Data Modules contain a single-low voltage CPU and are a compatible replacement for previous Data Modules.

Each module can be ordered equipped with either 1 TB or 2 TB drives (no intermix).




Figure 3-8 Data Module/Interface Module

Data Module

The fully populated rack hosts 9 Data Modules (Module 1-3 and Module 10-15). The only difference between Data Modules and Interface Modules (see “Interface Module” on page 59) are the additional host adapters and GigE adapters in the Interface Modules, along with a single-low voltage CPU as opposed to the dual-low voltage CPU in the Interface Module. The main components of the modules, in addition to the 12 disk drives (1 TB or 2 TB) are:

- ▶ System Planar (motherboard)
- ▶ Processor
- ▶ Memory/cache

- ▶ Enclosure Management Card
- ▶ Cooling devices (fans)
- ▶ Memory Flash Card
- ▶ Redundant power supplies
- ▶ Gigabit Ethernet ports

Each Data Module contains two redundant Gigabit Ethernet ports. These ports together with the two switches form the internal network, which is the communication path for data and metadata between all modules. The Dual GigE adapter is integrated in the System Planar (port 1 and 2) as seen in Figure 3-9.




Figure 3-9 Data Module connections that are used

System planar

The system planar used in the Data Modules and the Interface Modules is an ATX motherboard from Intel. This high-performance server board with a built-in Serial-Attached SCSI (SAS) adapter supports:

- ▶ Single or Dual (for Interface Modules only) 64-bit quad-core Intel Xeon® processors:
 - The dual-CPU Interface Modules ship with two low-voltage Central Processing Units (CPUs). These modules show an improvement in performance on sequential read and write operations over the single-CPU predecessor. The single CPU Interface Module uses the same low voltage CPU.
 - Dual-CPU Interface Modules (feature number 1101) can be used to complete the Interface Module portion of a six-module configuration that already has three single-CPU Interface Modules (feature number 1100).
 - A feature conversion from single-CPU Interface Modules (1100) to dual-CPU Interface Modules (1101) is *not* allowed.

- ▶ 8 x 1 GB or 4 x 2 GB on systems with 1 TB drives and 4 x 4 GB on systems with 2 TB drives, fully buffered 533/667 MHz Dual Inline Memory Module (DIMMs) to increase capacity and performance
- ▶ Dual Gigabit Ethernet with Intel I/O Acceleration Technology to improve application and network responsiveness by moving data to and from applications faster
- ▶ Four PCI Express slots to provide the I/O bandwidth needed by servers

Processor

The processor consists of either one or two Intel Xeon Quad Core Processors. This 64-bit processor has the following characteristics:

- ▶ 2.33 GHz clock
- ▶ 12 MB cache
- ▶ 1.33 GHz Front Serial Bus
- ▶ Low voltage power profile

For systems already deployed, there are no options to replace the current CPU in the Data and Interface Modules with the newer low voltage processor. For partially populated configurations, it is possible to expand the system with modules utilizing the newer low voltage CPU design.

Memory and cache

Every module on 1 TB drive systems has 8 GB of memory installed (either 4 x 2 GB or 8 x 1 GB); on systems with 2 TB drives, 16 GB of memory (4 x 4 GB) are installed. Fully Buffered DIMM (FBDIMM). FBDIMM memory technology increases reliability, speed, and density of memory for use with Xeon Quad Core Processor platforms. This processor memory configuration can provide three times higher memory throughput, enable increased capacity and speed to balance capabilities of quad core processors, perform reads and writes simultaneously, and eliminate the previous read to write blocking latency.

Part of the memory is used as module system memory, while the rest is used as cache memory for caching data previously read, pre-fetching of data from disk, and for delayed destaging of previously written data. For a description of the cache algorithm, see “Write cache protection” on page 38.

Cooling fans

To provide enough cooling for the disks, processor, and board, the system includes 10 fans located between the disk drives and the board. The cool air is aspirated from the front of the module through the disk drives. An air duct leads the air around the processor before it leaves the module through the back. The air flow and the alignment of the fans assure proper cooling of the entire module, even if a fan is failing.

Enclosure management card

The enclosure management card is located between the disk drives and the system planar. In addition to the internal module connectivity between the drive backplane and the system planar, this card is the backplane for the 10 fans. Furthermore, it includes fan control and the logic to generate hardware alarms in the case of problems in the module.

Compact Flash Card

Each module contains a Compact Flash Card (1 GB) in the right-most rear slot. See Figure 3-10.




Figure 3-10 Compact Flash Card

This card is the boot device of the module and contains the software and module configuration files.

Important: Due to the configuration files, the Compact Flash Card is not interchangeable between modules.

Power supplies

Figure 3-11 shows the redundant power supply units.




Figure 3-11 Redundant module power supply units

The modules are powered by a redundant Power Supply Unit (PSU) cage with a dual 850W PSU assembly as seen in Figure 3-11. These power supplies are redundant and can be individually replaced with no need to stop the system. The power supply is a Field-Replaceable Unit (FRU).

Interface Module

Figure 3-12 shows an Interface Module with iSCSI ports.

The Interface Module is similar to the Data Module. The only differences are as follows:

- ▶ Interface Modules contain iSCSI and Fibre Channel ports, through which hosts can attach to the IBM XIV Storage System. These ports can also be used to establish Remote Mirror links and data migration paths with another remote IBM XIV Storage System.
- ▶ There are two 4-port GigE PCIe adapters installed for additional internal network connections as well as for iSCSI host connections.

All Fibre Channel ports, iSCSI ports, and Ethernet ports used for external connections are internally connected to a patch panel where the external cable connections are made. See 3.2.5, “Patch panel” on page 63.




Figure 3-12 Interface Module with iSCSI ports

Fibre Channel connectivity

There are four Fibre Channel ports available in each Interface Module for a total of 24 Fibre Channel ports. Certain partial rack configurations do not utilize all ports even though they might be physically present. See Figure 3-3 on page 50 for more information.

The Fibre Channel ports support 2 Gbps and 4 Gbps full-duplex data transfer over short wave fibre links, using 50 micron multi-mode cable and support new end-to-end error detection through a Cyclic Redundancy Check (CRC) for improved data integrity during reads and writes.

The allocation and use of specific ports in each module depend on your environment, your specific requirements in terms of resiliency, the nature of your host I/O traffic, and whether you will use mirroring or not.

Most illustrations in this book show ports 1 and 3 allocated for host connectivity, while ports 2 and 4 are reserved for additional host connectivity or remote mirror and data migration connectivity. This is generally the choice for customers who want more resiliency (ports 1 and 3 are on separate adapters), or availability (in case of adapter firmware upgrade, one connection remains available through the other adapter), and also if the workload needs more performance (each adapter has its own PCI-e bus). If you will not use mirroring, you can also change port 4 to a target port.

Important: Customers are encouraged to discuss with their IBM support representatives to determine what port allocation will be most desirable in their environment.

Fibre Channel ports: Utilizing more than 12 Fibre Channel ports for host connectivity will not necessarily provide more bandwidth. Best practice is to utilize enough ports to support multipathing, without overburdening the host with too many paths to manage. See the *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904 for more information.

iSCSI connectivity

There are 6 iSCSI ports (two ports per Interface Modules 7 through 9) available for iSCSI over IP/Ethernet services. These ports support 1 Gbps Ethernet network connection. These ports connect to the user's IP network through the Patch Panel and provide connectivity to the iSCSI hosts. See Figure 3-14 on page 63 for additional details on the cabling of these ports.

iSCSI ports: Configurations with only six modules (three Data Modules, three Interface Modules) do not have available iSCSI ports. iSCSI ports only become available on systems with nine or more modules installed.

You can operate iSCSI connections for various uses:

- ▶ As an iSCSI target that the server hosts through the iSCSI protocol
- ▶ As an iSCSI initiator for Remote Mirroring when connected to another iSCSI port
- ▶ As an iSCSI initiator for data migration when connected to a third-party iSCSI storage system

For each iSCSI IP interface, you can define these configuration options:

- ▶ IP address (mandatory)
- ▶ Network mask (mandatory)
- ▶ Default gateway (optional)
- ▶ MTU; Default and Maximum: 4500 MTU

3.2.4 SATA disk drives

The SATA disk drives, which are shown in Figure 3-13 and used in the IBM XIV Storage System, are 1 TB or 2 TB, 7200 rpm hard drives designed for high-capacity storage in enterprise environments. These drives are manufactured to a higher set of standards than typical off the shelf SATA disk drives to ensure a longer life and increased reliability.




Figure 3-13 SATA disks

The IBM XIV Storage System was engineered with substantial protection against data corruption and data loss. Various features and functions implemented in the disk drive also increase reliability and performance. We describe the highlights next.

Performance features and benefits

The following performance features and benefits are included:

- ▶ 32 MB cache buffer:

The internal 32 MB cache buffer enhances the data transfer performance.

- ▶ Rotation Vibration Safeguard (RVS):

In multi-drive environments, rotational vibration, which results from the vibration of neighboring drives in a system, can degrade hard drive performance. To aid in maintaining high performance, the disk drive incorporates the enhanced Rotation Vibration Safeguard (RVS) technology, providing up to a 50% improvement over the previous generation against performance degradation, and therefore, leading the industry.

Reliability features and benefits

Reliability features and benefits include:

- ▶ Advanced magnetic recording heads and media:
There is an excellent soft error rate for improved reliability and performance.
- ▶ Self-Protection Throttling (SPT):
SPT monitors and manages I/O to maximize reliability and performance.
- ▶ Thermal Fly-height Control (TFC):
TFC provides a better soft error rate for improved reliability and performance.
- ▶ Fluid Dynamic Bearing (FDB) Motor:
The FDB Motor improves acoustics and positional accuracy.
- ▶ Load/unload ramp:
The R/W heads are placed outside the data area to protect user data when the power is removed.
- ▶ Self-Monitoring, Analysis and Reporting Technology (SMART) tools:
SMART tools simply indicate whether certain thresholds have been exceeded, thereby alerting that a disk is at risk for failure and thus needs to be replaced.

All IBM XIV Storage System disks are installed in the front of the modules, twelve disks per module. Each single SATA disk is installed in a disk tray, which connects the disk to the backplane and includes the disk indicators on the front. If a disk is failing, it can be replaced easily from the front of the rack. The complete disk tray is one FRU, which is latched in its position by a mechanical handle.

Important: SATA disks in the IBM XIV Storage System must never be swapped within a module or placed in another module because of internal tracing and logging data that they maintain.

3.2.5 Patch panel

The patch panel is located at the rear of the rack. Interface Modules are connected to the patch panel using 50 micron cables. All external connections must be made through the patch panel. In addition to the host connections and to the network connections, more ports are available on the patch panel for service connections. Figure 3-14 shows the details for the patch panel and the ports.

The patch panel has had various re-designs with respect to the labelling based on production date and is much easier to read in the latest configurations.




Figure 3-14 Connection details of Patch Panel ports

3.2.6 Interconnection and switches

The internal network is based on two redundant 48-port Gigabit Ethernet switches. Each of the modules (Data or Interface) is directly attached to each of the switches (see Figure 3-9 on page 56 and Figure 3-12 on page 59), and the switches are also linked to each other. This network topology enables maximum bandwidth utilization, as the switches are used in an active-active configuration, while being tolerant to any failure of the following individual network components:

- ▶ Ports
- ▶ Links
- ▶ Switches

Figure 3-15 shows the two Ethernet switches and the cabling to them.




Figure 3-15 48 Port Gigabit Ethernet Switch

The Gigabit Ethernet Layer 3-Switch contains 48 copper and 4 fiber ports (small form-factor pluggable (SFP) capable of one of 3 speeds, 10/100/10000 Mbps), robust stacking, and 10 Gigabit-Ethernet uplink capability. The switches are powered by redundant power supplies to eliminate any single point of failure.

3.2.7 Support hardware

This section covers important features of the IBM XIV Storage System used by internal functions and/or IBM maintenance, if a problem arises with the system.

Module USB to Serial connections

The Module USB to Serial connections are used by internal system processes to keep the communication to the modules alive in the event that the network connection is not operational. Modules are linked together with USB to Serial cables in groups of three modules. This emergency link is needed to communicate between the modules for internal processes and are used by IBM Maintenance in the event of repair to the internal network.

The USB to Serial connection always connects a group of three modules:

- ▶ USB Module 1 is connected to Serial Module 3.
- ▶ USB Module 3 is connected to Serial Module 2.
- ▶ USB Module 2 is connected to Serial Module 1.

This connection sequence is repeated for USB Modules 4-6, 7-9, 10-12, and 13-15.

For partially configured systems, such as 10 or 11 modules, the USB to Serial connections follow the same pattern as applicable.

In the case of a system with 10 modules, Modules 1-3 will be connected together, as will Modules 4-6 and 7-9, while Module 10 will be unconnected by this method.

This connection sequence for a fully configured system (15 modules) is depicted in the diagram shown in Figure 3-16.




Figure 3-16 Module: USB to serial

Maintenance module

The 1U maintenance module and the modem, which are installed in the middle of the rack, are used for IBM XIV Support and the IBM SSR to maintain and repair the machine. This device has fairly generic requirements because it is only used to gain remote access to the IBM XIV Storage System through the Secure Remote Support connection or modem for support personnel. When there is a software or hardware problem that needs the attention of the IBM XIV Support Center, a remote connection will be required and used to analyze and possibly repair the faulty system. This connection can be initiated by the customer and is done either through the XIV Remote Support Center (XRSC) facility or through a phone line and modem. For further information about remote connections, see “IBM XIV Remote Support Center (XRSC)” on page 77.

Modem

The modem installed in the rack is optionally used for remote support if the preferred choice of XIV Remote Support Center is not selected. It enables the IBM XIV Support Center specialists and, if necessary, a higher level of support to connect to the IBM XIV Storage System. Problem analysis and repair actions without a remote connection can be complicated and time-consuming.

Modem: The modem is not available in all countries.

3.3 Hardware planning overview

This section provides an overview of planning considerations for the IBM XIV Storage System, including a reference listing of the information required for the setup. The information in this section includes requirements for:

- ▶ Physical installation
- ▶ Delivery requirements
- ▶ Site requirements
- ▶ Cabling requirements

For more detailed planning information, see the *IBM XIV Storage System Planning Guide* v2.3, GA32-0770.

Additional documentation is available from the XIV InfoCenter at:

<http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp>

For a smooth and efficient installation of the IBM XIV Storage System, planning and preparation tasks must take place before the system is scheduled for delivery and installation in the data center. A marketing or sales representative will arrange a Technical and Delivery Assessment (TDA) meeting to go over site specific details and to ensure that proper information is gathered prior to the delivery of the system.

There are four major areas involved in installation planning:

- ▶ Ordering the appropriate IBM XIV hardware configuration:
 - Current requirements
 - Future requirements
- ▶ Physical site planning:
 - Space, dimensions, and weight
 - Raised floor
 - Power, cooling, cabling, and additional equipment
- ▶ Configuration planning:
 - Basic configurations
 - Network connections
 - Management connections
- ▶ Physical installation

3.3.1 Ordering IBM XIV Storage System hardware

This part of the planning describes ordering the appropriate XIV hardware configuration. At this point, consider actual requirements, but also consider potential future requirements.

Feature codes and hardware configuration

The IBM XIV Storage System hardware is mostly pre-configured and consists of two machine types, namely the 2810-A14 and the 2812-A14, which as previously mentioned, only differ in their initial warranty coverage and are now equipped with dual CPU Interface Modules. All features are identical on each system.

Optional features

There are only a few optional or specific hardware features that you can select as part of the initial order, as for instance the choice between 1 TB or 2 TB disks, or a partially populated configuration.

3.3.2 Physical site planning

Physical planning considers the size, weight, and the environment on which you will install the IBM XIV Storage System.

Site requirements

The physical requirements for the room where the IBM XIV Storage System is going to be installed must be checked well ahead of the arrival of the machine:

- ▶ The floor must be able to withstand the weight of the IBM XIV Storage System to be installed.

Consider also possible future machine upgrades. The IBM XIV Storage System can be installed on a non-raised floor, but it is certainly best to use a raised floor for increased air circulation and cooling.

- ▶ Enough clearance around the system must be left for cooling and service.

The airflow enters the system rack on the front and is expelled to the rear. You must ensure that there is enough service clearance to fully open the front and back doors.

- ▶ Building particularities, such as any ramps, elevators, and floor characteristics according to the height and weight of the machine must also be considered.

Figure 3-17 gives a general overview of the clearance needed for airflow and service around the Standard XIV rack.




Figure 3-17 Dimensions of the Standard XIV Rack

For detailed information and further requirements, see the *IBM XIV Storage System Planning Guide*, GA32-0770.

Weight and raised floor requirements

The system is normally shipped as one unit. There is a weight reduction feature available for the Standard XIV rack in case the access to the site cannot accommodate the weight of the fully populated rack during its movement to the final install destination.

The following measurements are provided for your convenience. For the latest and most accurate information, see the *IBM XIV Storage System Planning Guide*, GA32-0770.

- ▶ Weight is 884 kg (1949 lb.) using the Standard XIV Rack.
- ▶ Raised floor requirements:
 - Provide reinforcement to support the system weight in an area of 60 cm x 109 cm.
 - Provide enough ventilation tiles in front of the rack.
 - Provide a cutout (opening) for the cables according to the template in the *IBM XIV Storage System Planning Guide*, GA32-0770.

See the system planning guide for these additional requirements:

- ▶ Power requirements
- ▶ Cooling requirements
- ▶ Delivery requirements

3.3.3 Basic configuration planning

You must complete the configuration planning worksheets and checklists as found in the Planning section of the IBM XIV Storage System Information Center to allow the IBM SSR to physically install and configure the system. The configuration planning worksheets can also be found in the *IBM XIV Storage System Planning Guide v2.3, GA32-0770*.

Additional documentation is available from the XIV InfoCenter at:

<http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp>

You must provide the IBM SSR with the information required to attach the system to your network for operations and management, as well as enabling remote connectivity for IBM support and maintenance.

Fill out each document with all required information to prevent further inquiry and delays during the installation (see 3.3.4, “IBM XIV physical installation” on page 78):

- ▶ **Interface Module:**

Interface Modules (4,5 and 6) need an IP address, Netmask, and Gateway. This address is necessary to manage and monitor the IBM XIV Storage System with either the GUI or Extended Command Line Interface (XCLI). Each Interface Module needs a separate IP address in case a module is failing.

- ▶ **DNS server:**

If Domain Name System (DNS) is used in your environment, the IBM XIV Storage System needs to have the IP address, Netmask, and Gateway from the primary DNS server and, if available, also from the secondary server.

- ▶ **SMTP Gateway:**

The Simple Mail Transfer Protocol (SMTP) Gateway is needed for event notification through email. The IBM XIV Storage System can initiate an email notification, which will be sent out through the configured SMTP Gateway (IP Address or server name, Netmask, and Gateway).

- ▶ **NTP (Time server):**

IBM XIV can be used with a Network Time Protocol (NTP) time server to synchronize the system time with other systems. To use this time server, IP Address, or server name, Netmask and Gateway must be configured.

- ▶ **Time zone:**

Usually the time zone depends on the location where the system is installed. However, exceptions can occur for remote locations where the time zone equals the time of the host system location.

- ▶ **Email sender address:**

This is the email address that is shown in the email notification as the sender.

- ▶ **Remote access:**

The modem number or a client side IP Address needs to be configured for remote support. This network connection must have outbound connectivity to the Internet.

This basic configuration data will be entered in the system by the IBM SSR following the physical installation. See “Basic configuration” on page 79.

Other configuration tasks, such defining storage pools, volumes, and hosts, are the responsibility of the storage administrator and are described in Chapter 4, “Configuration” on page 83.

Network connection considerations

Network connection planning is also essential to prepare to install the IBM XIV Storage System. To deploy and operate the system in your environment, a number of network connections are required:

- ▶ Fibre Channel connections for host I/O over Fibre Channel
- ▶ Gigabit Ethernet connections for host I/O over iSCSI
- ▶ Gigabit Ethernet connections for management
- ▶ Gigabit Ethernet connections for IBM XIV Storage System remote support

All external IBM XIV Storage System connections are hooked up through the patch panel as explained in 3.2.5, “Patch panel” on page 63.

For details about the host connections, see the IBM Redbooks publication, *XIV Storage System: Host Attachment and Interoperability*, SG24-7904.

Fibre Channel connections

When shipped, the IBM XIV Storage System is by default equipped with 24 Fibre Channel ports (assuming a fully populated 15 Module rack). The IBM XIV Storage System supports 50 micron fiber cabling. If you have other requirements or special considerations, contact your IBM Representative.

The 24 FC ports are available from the six Interface Modules, four in each module, and they are internally connected to the patch panel. Of the 24 ports, 12 are provided for connectivity to the switch network for host access, and the remaining 12 are for use in remote mirroring or data migration scenarios (however, they can be reconfigured for host connectivity). Adhere to this guidance on Fibre Channel connectivity. The external (client-provided) cables are plugged into the patch panel. For planning purposes, Table 3-1 highlights the maximum values for various Fibre Channel parameters for your consideration.

Table 3-1 Maximum values in context of FC connectivity

FC Parameters	Maximum values
Maximum number of Interface Modules	6
Maximum number of 4 GB FC ports per Interface Module	4
Maximum queue depth per FC Adapter Port (Same Type Target/LUN)	1400
Maximum queue depth per FC Adapter Port (separate Type Target/LUN)	955/1845
Maximum queue depth per mapped volume per (Initiator, target, LUN)	256
Maximum FC ports for host connections (default configuration)	12
Maximum FC ports for migration/mirroring (default config)	12
Maximum volumes mapped per host	512
Maximum number of clusters	100
Maximum number of hosts (defined WWPNs/iSCSI qualified names)	1000
Maximum number of mirroring coupling (number of mirrors)	512
Maximum number of mirrors on remote machine	512
Maximum number of remote targets	8

Values in Table 3-1 apply to Release 10.2.2 of the IBM XIV Storage System software.

For details about Fibre Channel configuration and connections, see the IBM Redbooks publication, *XIV Storage System: Host Attachment and Interoperability*, SG24-7904.

iSCSI connections

When shipped, the IBM XIV Storage System is by default equipped with 6 iSCSI ports (assuming a fully populated 15 Module rack with Interface Modules 7-9 providing the 6 ports).

The external (client-provided) Ethernet cables are plugged into the patch panel. For planning purposes, Table 3-2 highlights the maximum values for various iSCSI parameters for your consideration. As with Fibre Channel, it is important to plan your connectivity based on these maximums.

Table 3-2 Maximum values in context of iSCSI connectivity

iSCSI parameters	Maximum values
Maximum number of Interface Modules with iSCSI ports	3
Maximum number of 1GB iSCSI ports per Interface Module	2
Maximum queue depth per iSCSI host port	1400
Maximum queue depth per mapped volume per host port/ target port/volume ordered set	256
Maximum iSCSI ports for any connection (host or XDRP)	6
Maximum number of hosts (iSCSI qualified names (IQNs))	1000
Maximum number of mirroring coupling (number of mirrors)	512
Maximum number of mirrors on remote machine	512
Maximum number of remote targets	8

These values are correct at the time of writing this book for Release 10.2.2 of the IBM XIV Storage System software.

Redundant configuration

To configure the Fibre Channel connections (SAN) for high availability, see the configuration illustrated in Figure 3-18. This configuration is best for all production systems to maintain system access and operations following a single hardware element or SAN component failure.

Our illustrations show ports 1 and 3 allocated for host connectivity while ports 2 and 4 are reserved for additional host connectivity or remote mirror and data migration connectivity. If you will not use mirroring, you can also change port 4 to a target port.

Important: Customers are encouraged to discuss with their IBM support representatives to determine what port allocation will be most desirable in their environment.




Figure 3-18 Redundant configuration

For a redundant Fibre Channel configuration, use the following guidelines:

- ▶ Each IBM XIV Storage System Interface Module is connected to two Fibre Channel switches, using two ports of the module. One patch panel connection to each Fibre Channel switch.
- ▶ Each host is connected to two switches using two host bus adapters or a host bus adapter (HBA) with two ports.

This configuration assures full connectivity and no single point of failure:

- ▶ Switch failure: Each host remains connected to all modules through the second switch.
- ▶ Module failure: Each host remains connected to the other modules.
- ▶ Cable failure: Each host remains connected through the second physical link.
- ▶ Host HBA failure: Each host remains connected through the second physical link.

Paths: Although up to 12 paths per host can be defined for most servers, a maximum of 6 paths per host provides sufficient redundancy and performance.

Ports: Utilizing more than 12 Fibre Channel ports for host connectivity will not necessarily provide more bandwidth. Best practice is to utilize enough ports to support multipathing, without overburdening the host with too many paths to manage. See the *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904 for more information.

A redundant configuration accessing still all Interface Modules, but using only a minimum number of 6 paths per LUN on the host, is depicted in Figure 3-19.




Figure 3-19 Fibre Channel redundant configuration

In this redundant configuration:

- ▶ Each host is equipped with dual HBAs. Each HBA (or HBA port) is connected to one of two FC switches.
- ▶ Each of the FC switches has a connection to a separate FC port of each of the six Interface Modules.
- ▶ One host is using the first 3 paths per fabric and the next host is using the 3 other paths per fabric.
- ▶ If a fabric fails, all Interface Modules will still be used.
- ▶ Each LUN has 6 paths.

Single switch solution

This configuration is resilient to the failures of a single Interface Module, host bus adapter, and cables, however in this configuration the switch represents a single point of failure. If the switch goes down due to a hardware failure or simply because of a software update, the connected hosts will lose all data access. Figure 3-20 depicts this configuration option.




Figure 3-20 Non-redundant configuration

Only use a single switch solution when no second switch is available and/or for test environments only.

Single HBA host connectivity

Hosts that are equipped with a single Fibre Channel port can only access one switch. So, this configuration is resilient to the failure of an individual Interface Module, but there are various possible points of failure (switch, cable, and HBA), which can cause access loss from the host to the IBM XIV Storage System. This configuration, which is not desirable for any production system, must be used if there is no way of adding a second Fibre Channel port to the host.

Restriction: Direct host to XIV connectivity is *not* permitted. Implementation must make use of a SAN fabric (either single or dual SAN switches), and dual fabric configurations are best.

Fibre Channel cabling and configuration

Fibre Channel cabling must be prepared based on the required fibre length and depending on the selected configuration.

When installing an IBM XIV Storage System, perform the following Fibre Channel configuration procedures:

- ▶ You must configure Fibre Channel switches zoned correctly allowing access between the hosts and the IBM XIV Storage System. The specific configuration to follow depends on the specific Fibre Channel switch. It is best to have a separate zone for each initiator.
- ▶ Hosts need to be set up and configured with the appropriate multipathing software to balance the load over various paths. For multipathing software and setup, see the specific operating system section in the IBM Redbooks publication, *XIV Storage System: Host Attachment and Interoperability*, SG24-7904.

iSCSI network configurations

Logical network configurations for iSCSI are equivalent to the logical configurations that are suggested for Fibre Channel networks. Four options are available:

- ▶ Redundant Configuration: Each module connects through two ports to two Ethernet switches, and each host is connected to the two switches. This design provides a network architecture resilient to a failure of any individual network switch or module.
- ▶ Single switch configuration: A single switch interconnects all modules and hosts.
- ▶ Single port host solution: Each host connects to a single switch, and a switch is connected to two modules.

IP configuration

The configuration of the IBM XIV Storage System iSCSI connection is highly dependent on your network. In the high availability configuration, the two client-provided Ethernet switches used for redundancy can be configured as either two IP subnets or as part of the same subnet. The IBM XIV Storage System iSCSI configuration must match the client's network. You must provide the following configuration information for each Ethernet port:

- ▶ IP address
- ▶ Net mask
- ▶ MTU (optional):

Maximum Transmission Unit (MTU) configuration is required if your network supports an MTU, which differs from the standard one. The largest possible MTU must be specified. From the XIV side the largest possible MTU size is 4500. If the iSCSI hosts reside on another subnet than the IBM XIV Storage System, a default IP gateway per port must be specified.

- ▶ Default gateway (optional):

Because IBM XIV Storage System always acts as a TCP server for iSCSI connections, packets are always routed through the Ethernet port from which the iSCSI connection was initiated. The default gateways are required only if the hosts do not reside on the same layer-2 subnet as the IBM XIV Storage System.

The IP network configuration must be ready to ensure connectivity between the IBM XIV Storage System and the host prior to the physical system installation:

- ▶ Ethernet Virtual Local Area Networks (VLANs), if required, must be configured correctly to enable access between hosts and the IBM XIV Storage System.
- ▶ IP routers (if present) must be configured correctly to enable access between hosts and the IBM XIV Storage System.

Mixed iSCSI and Fibre Channel host access

IBM XIV Storage System does not support mixed concurrent access from the same host to the same volume through FC and iSCSI simultaneously. You can, however, access the same volume from separate hosts using various connection methods.

Management connectivity

IBM XIV Storage System is managed through three IP addresses over Ethernet interfaces on the patch panel in order to be resilient to two hardware failures. Thus, you must have three Ethernet ports available for management. If you require management to be resilient to a single network failure, be sure to connect these ports to two switches.

Make sure as well that the networking equipment providing the management communication is protected by an Uninterruptible Power Supply (UPS).

Management IP configurations

For each of the three management ports, you must provide the following configuration information to the IBM SSR upon installation (refer also to 3.3.3, “Basic configuration planning” on page 69):

- ▶ IP address of the port (all three ports must belong to the same subnet)
- ▶ Subnet mask
- ▶ Default IP gateway (if required)

Provide the following system-level IP information (not port-specific):

- ▶ IP address of the primary and secondary DNS servers
- ▶ IP address or DNS names of the SNMP manager, if required
- ▶ IP address or DNS names of the Simple Mail Transfer Protocol (SMTP) servers

Protocols

The IBM XIV Storage System is managed through dedicated management ports running TCP/IP over Ethernet. Management is carried out through the following protocols (consider this design when configuring firewalls, other security protocols, and SMTP relaying):

- ▶ Proprietary XIV protocols are used to manage the IBM XIV Storage System from the GUI and the XCLI. This management communication is performed over TCP port 7778, where the GUI/XCLI, as the client, always initiates the connection, and the IBM XIV Storage System performs as the server.
- ▶ IBM XIV Storage System sends and responds to SNMP management packets.
- ▶ IBM XIV Storage System initiates SNMP packets when sending traps to SNMP managers.
- ▶ IBM XIV Storage System initiates SMTP traffic when sending emails (for either event notification through email or for email-to-SMS gateways).
- ▶ IBM XIV Storage System communicates with remote SSH connections over standard TCP port 22.

SMTP server

For proper operation of the XIV Call Home function, the SMTP server must function as follows:

- ▶ Be reachable on port 25 for the XIV customer specified management IP addresses.
- ▶ Permit relaying from the XIV customer specified management IP addresses.
- ▶ Permit the IBM XIV Storage System to send emails. The default sender address is fromxiv@il.ibm.com, but it can be customized.
- ▶ Permit recipient addresses of xiv-callhome-western-hemisphere@vnet.ibm.com and xiv-callhome-eastern-hemisphere@vnet.ibm.com. Additionally, you can choose the shorter alias recipient address xiv-west@vnet.ibm.com and xiv-east@vnet.ibm.com

Laptop ports

The IBM XIV Storage System has two Ethernet laptop ports. The system will serve as a Dynamic Host Configuration Protocol (DHCP) server and will automatically configure the mobile computer. These ports are only for IBM Maintenance personnel.

Restriction: Do not connect these ports to the user (client) network.

IBM XIV Remote Support Center (XRSC)

To facilitate remote support by IBM XIV Storage System personnel, you configure a dedicated Ethernet port for remote access. This port must be connected through the organizational firewall so that IBM XIV Storage System personnel can access the IBM XIV Storage System, if required.

The IBM XIV Remote Support Center comprises XIV internal functionality together with a set of globally deployed supporting servers to provide secure IBM support access to the IBM XIV Storage System when necessary and when authorized by the customer personnel. Figure 3-21 provides a representation of the data flow of the XIV to IBM Support.




Figure 3-21 XIV Remote Support Center

An optional Remote Support Proxy can be used when one or more XIV systems do not have direct access to the Internet (for example, due to firewall restrictions), you can use the XIV Remote Support Proxy to facilitate the connection to the XIV Remote Support Center. For more information about the Remote Support Proxy, see the Installation and User's Guide, GA32-0795.

To initiate the remote connection process, the following procedure is performed:

1. Customer initiates an Internet based SSH connection to XRSC either by the GUI or XCLI.
2. XRSC identifies the IBM XIV Storage System and marks it as "connected".
3. IBM Support person connects to XRSC using SSH over the IBM Intranet.
4. XRSC authenticates the IBM Support person against the IBM Intranet.

5. XRSC then displays the connected customer system available to the IBM Support person.
6. The IBM Support person then chooses which system to support and connect to:
 - Only permitted XIV systems are shown.
 - IBM Support person logs the intended activity.
7. A fully recorded support session starts.
8. When complete, the IBM Support person terminates the session and the XRSC disconnects the XIV array from the remote support system.

The XRSC Internet servers are hard coded in XIV Software and no further configuration is required by the customer to enable this function, aside from turning this feature on in the GUI or XCLI. This provides an expedient manner for IBM Support to gather required information from the system in the most timely fashion and with the least impact to the customer.

Modem connectivity

The modem installed in the rack is optionally used for remote support if the preferred choice of XIV Remote Support Center is not used. It enables the IBM XIV Support Center specialists and, if necessary, a higher level of support to connect to the IBM XIV Storage System. Problem analysis and repair actions without a remote connection can be complicated and time-consuming.

Remote mirroring connectivity

Planning the physical connections also includes considerations when the IBM XIV Storage System is installed in a Remote Copy environment. Contact advanced IBM XIV support for assistance for planning Remote Mirroring connectivity to assure the maximum resilience to hardware failures and connection failures.

Remote Copy links, which connect the direct primary system and secondary system, must also be planned for prior to the physical installation. The physical Remote Copy links can be Fibre Channel links, direct or through a SAN, or iSCSI port connections using Ethernet; however, iSCSI is not desirable for this use.

Planning for growth

Ensure consideration for growth and for the future IO demands of your business. Most applications and databases grow quickly and the need for greater storage capacity increases rapidly. Planning for growth prior to the implementation of the first IBM XIV Storage System in the environment can save time and re-configuring effort in the future.

3.3.4 IBM XIV physical installation

After all previous planning steps are completed and the system is delivered to its final location, the physical installation can begin. A IBM SSR will perform all the necessary tasks and perform the first logical configuration steps up to the point where you can connect the IBM XIV Storage System through the GUI and the XCLI. Configuring Storage Pools, logical unit numbers (LUNs), and attaching the IBM XIV Storage System to the host are storage administrator responsibilities. See Chapter 4, “Configuration” on page 83.

Physical installation

It is the responsibility of the customer or moving contractor to unpack and move the IBM XIV Storage System as close as possible to its final destination before an IBM SSR can start the physical installation. Carefully check and inspect the delivered crate and hardware for any visible damage. If there is no visible damage, and the tilt and shock indicators show no problem, sign for the delivery.

Also arrange that with the start, or during the physical installation, an electrician is available who is able to handle the power requirements in the environment up to the IBM XIV Storage System power connectors.

The physical installation steps are as follows:

1. Place and adjust the rack in its final position in the data center.
2. Check the IBM XIV Storage System hardware. When the machine is delivered with the weight reduction feature (FC 0200), the IBM SSR will install the removed modules and components into the rack.
3. Connect the IBM XIV Storage System line cords to the client-provided power source and advise an electrician to switch on the power connections.
4. Perform the initial power-on of the machine and perform necessary checks according to the given power-on procedure.
5. To complete the physical steps of the installation, the IBM SSR will perform various final checks of the hardware before continuing with the basic configuration.

Basic configuration

After the completion of the physical installation steps, the IBM SSR establishes a connection to the IBM XIV Storage System through the patch panel (see 3.2.5, “Patch panel” on page 63) and completes the initial setup. You must provide the required completed information sheet that is referenced in 3.3.3, “Basic configuration planning” on page 69.

The basic configuration steps are as follows:

1. Set the Management IP Addresses, (Client Network) Gateway, and Netmask.
2. Set the system name.
3. Set the email sender address and SMTP server address.
4. Set the primary DNS and the secondary DNS.
5. Set the SNMP management server address.
6. Set the time zone.
7. Set the NTP server address.
8. Configure the system to send events to IBM (Call Home).
9. Configure and test Remote Support.

Completing the physical installation

After the IBM SSR completes the physical installation and initial setup, the IBM SSR performs the final checks for the IBM XIV Storage System:

- ▶ Power off and power on the machine by using the XIV Storage Management (GUI) and XCLI.
- ▶ Check the Events carefully for problems.
- ▶ Verify that all settings are correct and persistent.

At this point, the installation is complete, and the IBM XIV Storage System is ready to be handed over to the customer to configure and use. See Chapter 4, “Configuration” on page 83.

3.3.5 System power-on and power-off

Strictly follow these procedures to power on and power off your XIV system.

Power-on

To power on the system, follow these steps:

1. On each UPS, look for the Test button located on the Control Panel (front of the UPS) as illustrated in Figure 3-22.

Important: Do not confuse the Test button with the Power Off button, which is normally protected by a cover. The Test button is the one circled in Figure 3-22.



Figure 3-22 Locate Test button

2. Use both hands as shown in Figure 3-23 to press each of the three Test buttons simultaneously.



Figure 3-23 Use both hands to hit the three Test buttons simultaneously

This will start applying power to the components in the rack, and initiate the boot process for the Interface Modules and Data Modules.

Power-off

Powering the system off must be done solely from either the XIV GUI or the XCLI. You must be logged on as Storage Administrator (*storageadmin* role).

Warning: Do *not* power off the IBM XIV Storage System using the UPS power button because this can result in the loss of data and system configuration information.

Using the GUI

From the XIV GUI, follow these steps:

1. Click **File** in the system main window toolbar and choose the option **Shutdown System**, as seen in Figure 3-24.




Figure 3-24 System shutdown

2. You have to confirm, first by clicking the **OK** button, and second by entering your login password, as shown in Figure 3-25.




Figure 3-25 Confirm system shutdown

The shutdown takes less than 5 minutes. When done, all fans and front lights on modules and all UPS lights are off.

Tip: Using the GUI is the most convenient and best way to power off the system.

Using the XCLI

From the command prompt, issue the command:

```
xcli -c "XIV_01" -u admin -p adminadmin shutdown -y
```

You receive the response:

Command executed successfully

Where XIV_01 is the system name

Or, if using the XCLI session, see the shutdown procedure in Example 3-1.

Example 3-1 Executing a shutdown from the XCLI session

```
User Name: itso
Password: *****
Machine IP/Hostname: 9.155.53.250
connecting.
XIV_01>>shutdown
Password: *****

Warning:  ARE_YOU_SURE_YOU_WANT_TO_SHUT_DOWN y/n:
Command executed successfully
```

The shutdown takes less than 5 minutes. When done, all fans and front lights on modules and all UPS lights are off.

Emergency Power Off (EPO)

The IBM XIV Storage System Machine Types 2810 and 2812 do not have an emergency power off (EPO) switch; however, you can connect the IBM XIV Storage System to a room EPO.

EPO precaution: Powering off the IBM XIV Storage System using a room EPO switch results in data loss and possible loss of configuration. An IBM service representative is required to recover an XIV system that was turned off using a room EPO switch. If the IBM XIV Storage System loses ac power but is not powered off using an EPO circuit, data and configuration are preserved.

Local laws: National or local building, electrical, fire prevention, safety and/or other laws or regulations can address or control the manner in which information technology equipment is installed within certain facilities and environments. The application of those laws or regulations can depend on considerations of factors beyond the nature or design of the equipment to be installed. It is a customer responsibility to interpret and identify any laws or regulations applicable to the installation of information technology in its environment and to inform IBM, IBM Business Partners or their designated installers of any actions not identified in this document that are necessary to install information technology equipment in the customer's facilities in accordance with such applicable laws or regulations.

Contact your IBM service representative for more information about connecting to a room EPO switch.



Configuration

In this chapter we discuss the tasks to be performed by the storage administrator to configure the IBM XIV Storage System using the XIV Management Software.

We provide step-by-step instructions covering the following topics, in this order:

1. Installing and customizing the XIV Management Software
2. Connecting to and managing XIV using graphical and command line interfaces
3. Organizing system capacity by storage pools
4. Creating and managing volumes in the system
5. Host definitions, mappings and connectivity
6. XCLI scripts

Important: Illustrations in this chapter mostly apply to an IBM XIV Storage System fully configured with 1 TB drives.

4.1 IBM XIV Storage Management software

The IBM XIV Storage System software supports the functions of the IBM XIV Storage System. The software provides the functional capabilities of the system. It is preloaded on each module (Data and Interface Modules) within the IBM XIV Storage System. The functions and nature of this software are equivalent to what is usually referred to as microcode or firmware on other storage systems.

The XIV Storage Management software is used to communicate with the IBM XIV Storage System Software, which in turn interacts with the IBM XIV Storage System hardware.

The XIV Storage Manager can be installed on a Windows, Linux, AIX, HP-UX, or Solaris workstation that will then act as the management console for the IBM XIV Storage System. The Storage Manager software is provided at the time of installation, or is optionally downloadable from the following website:

<http://www.ibm.com/systems/support/storage/XIV>

For detailed information about the XIV Storage Management software compatibility, see the XIV interoperability matrix or the System Storage Interoperability Center (SSIC) at:

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

4.1.1 XIV Storage Management user interfaces

The IBM XIV Storage Manager includes a user-friendly and intuitive Graphical User Interface (GUI) application, an XIV Command Line Interface (XCLI) component offering a comprehensive set of commands to configure and monitor the system, as well as XIV Top, a GUI application to monitor system statistics and performance.

Graphical User Interface

A simple and intuitive Graphical User Interface (GUI) lets you perform most administrative and technical operations (depending upon the user role) for multiple XIV systems, quickly and easily, with minimal training and knowledge.

The main motivation behind the XIV management and GUI design is the desire to eliminate the complexities of system management. The most important operational challenges, such as overall configuration changes, volume creation or deletion, snapshot definitions, and many more, are achieved with a few clicks.

This chapter contains descriptions and illustrations of tasks performed by a storage administrator when using the XIV graphical user interface (GUI) to interact with the system.

XIV Command Line Interface

The XIV Command Line Interface (XCLI) is a powerful text-based, command line-based tool that enables an administrator to issue simple commands to configure, manage, or maintain the system, including the definitions required to connect to hosts and applications. The XCLI tool can be used in an XCLI Session environment to interactively configure the system or as part of a script to perform lengthy and more complex tasks.

Tip: Any operation that can be performed by the XIV GUI is also supported by the XIV Command Line Interface (XCLI).

This chapter presents the most common XCLI commands and tasks normally used by the administrator to interact with the system.

XIV Top GUI

The XIV Top GUI provides the user a simple tool for monitoring configured system statistics and performance on a host and volume level. See Chapter 7, “Monitoring” on page 203 for additional information.

4.1.2 XIV Storage Management software installation

This section illustrates the step-by-step installation of the XIV Storage Management software under Microsoft Windows XP.

The GUI can be downloaded at:

<ftp://ftp.software.ibm.com/storage/XIV/GUI/>

It also contains a demo mode. To use the demo mode, log on as user P10DemoMode and no password.

Important: Minimum requirements for installation in Windows XP are as follows:

CPU:	Dual Core or equivalent
Memory:	512 MB RAM (1 GB RAM and above are best)
Graphic Memory:	128 MB
Disk capacity:	100 MB Free
Screen resolution:	1024x768 (Best:1024x768 - 1920x1200)
Graphics:	24/32 True Color: Best

At the time of writing, the IBM XIV Storage Manager Version 2.4.3.b was available. Later GUI releases might differ slightly in appearance.

Perform the following steps to install the XIV Storage Management software:

1. Locate the XIV Storage Manager installation file (either on the installation CD or a copy that you downloaded from the Internet). Running the installation file first shows the Welcome window displayed in Figure 4-1. Click **Next**.



Figure 4-1 Installation: Welcome window

2. A Setup dialog window is displayed (Figure 4-2) where you can specify the installation directory. Keep the default installation folder or change it according to your needs. When done, click **Next**.




Figure 4-2 Specify the installation directory

3. The next installation dialog is displayed. You can choose between a FULL installation method or just a command line interface installation method. Choose **FULL installation** as shown in Figure 4-3. In this case, the Graphical User Interface and the Command Line Interface as well will be installed. Click **Next**.




Figure 4-3 Choose the installation type

4. The next step is to specify the Start Menu Folder as shown in Figure 4-4. When done,click **Next**.




Figure 4-4 Select Start Menu Folder

5. The dialog shown in Figure 4-5 is displayed. Select the desktop icon placement and click **Next**.




Figure 4-5 Select additional tasks

6. The dialog window shown in Figure 4-6 is displayed. Review your selections and click **Install**.




Figure 4-6 Review selections

7. The dialog window shown in Figure 4-7 is displayed. The XIV Storage Manager requires the Java Runtime Environment Version 6, which will be installed during the setup if needed. Click **Finish**.




Figure 4-7 Completing setup

4.2 Managing the IBM XIV Storage System

The basic storage system configuration sequence followed in this chapter includes the initial installation steps, followed by disk space definition and management.

Additional configuration or advanced management tasks are cross-referenced to specific chapters where they are discussed in more detail. For example, allocating and mapping volumes to hosts is covered in the IBM Redbooks publication, *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904.

Figure 4-8 presents an overview of the configuration flow. Note that the XIV GUI is very intuitive, and you can easily and quickly achieve most configuration tasks.




Figure 4-8 Basic configuration sequence and advanced tasks

After the installation and customization of the XIV Management Software on a Windows, Linux, AIX, HP-UX, or Solaris workstation, a physical Ethernet connection must be established to the IBM XIV Storage System itself.

The Management workstation is used to:

- ▶ Execute commands through the XCLI interface (see 4.2.2, “Logging on to the system with XCLI” on page 99).
- ▶ Control the IBM XIV Storage System through the GUI.
- ▶ Configure email notification messages and Simple Network Management Protocol (SNMP) traps upon occurrence of specific events or alerts. See Chapter 7, “Monitoring” on page 203.

To ensure management redundancy in case of Interface Module failure, the IBM XIV Storage System management functionality is accessible from three IP addresses. Each of the three IP addresses is linked to another (hardware) Interface Module. The various IP addresses are transparent to the user, and management functions can be performed through any of the IP addresses. These addresses can also be used simultaneously for access by multiple management clients. Users only need to configure the GUI or XCLI for the set of IP addresses that are defined for the specific system.

Notes: All management IP interfaces must be on the same subnet and use the same components:

- ▶ Network mask
- ▶ Gateway
- ▶ Maximum Transmission Unit (MTU)

Both XCLI and GUI management run over TCP port 7778 with all traffic encrypted through the Secure Sockets Layer (SSL).

4.2.1 The XIV Storage Management GUI

This section reviews the XIV Storage Management GUI.

Launching the XIV Storage Management GUI

Upon launching the XIV GUI application, a login window prompts you for a user name and its corresponding password before granting access to the IBM XIV Storage System. The default user is “*admin*” and the default corresponding password is “*adminadmin*”, as shown in Figure 4-9.

Important: Remember to change the default passwords to properly secure your system.

The default admin user comes with storage administrator (storageadmin) rights. The IBM XIV Storage System offers role-based user access management.

For more information about user security and roles, see Chapter 5, “Security” on page 131.




Figure 4-9 Login window with default access

To connect to an IBM XIV Storage System, you must initially add the system to make it visible in the GUI by specifying its IP addresses.

To add the system:

1. Make sure that the management workstation is set up to have access to the LAN subnet where the IBM XIV Storage System resides. Verify the connection by pinging the IP address of the IBM XIV Storage System.

If this is the first time you start the GUI on this management workstation and no IBM XIV Storage System had been previously defined to the GUI, the Add System Management dialog window is automatically displayed:

- If the default IP address of the IBM XIV Storage System was not changed, check **Connect Directly**, which populates the IP/DNS Address1 field with the default IP address. Click **Add** to effectively add the system to the GUI.
- If the default IP address had already been changed to a client-specified IP address (or set of IP addresses, for redundancy), you must enter those addresses in the IP/DNS Address fields. Click **Add** to effectively add the system to the GUI. See Figure 4-10.




Figure 4-10 Add Managed System

2. You are now returned to the main XIV Management window. Wait until the system is displayed and shows as enabled. Under normal circumstances, the system will show a status of Full Redundancy displayed in a green label box. Left-click the label box to display additional status indicators including Systems IOPS and Utilization levels.
3. Move the mouse cursor over the image of the IBM XIV Storage System and click to open the IBM XIV Storage System Management main window as shown in Figure 4-11.

XIV Storage Management GUI main menu

The XIV Storage Management GUI is mostly self-explanatory with a well-organized structure and simple navigation.




Figure 4-11 XIV Storage Manager main window: System view

The main window is divided into the following areas:

- **Function icons:** Located on the left side of the main window, you find a set of vertically stacked icons that are used to navigate between the functions of the GUI, according to the icon selected.

Moving the mouse cursor over an icon brings up a corresponding option menu. The various menu options available from the function icons are presented in Figure 4-12 on page 94.

- **Main display:** This area occupies the major part of the window and provides graphical representation of the IBM XIV Storage System. Moving the mouse cursor over the graphical representation of a specific hardware component—module, disk, and Uninterruptible Power Supply (UPS) unit—brings up a status callout.

When a specific function is selected, the main display shows a tabular representation of that function.

- ▶ **Menu bar:** This area is used for configuring the system and as an alternative to the Function icons for accessing the various functions of the IBM XIV Storage System.
- ▶ **Toolbar:** It is used to access a range of specific actions linked to the individual functions of the system.
- ▶ **Status bar:** These indicators are located at the bottom of the window. This area indicates the overall operational status of the IBM XIV Storage System:
 - The first indicator on the left shows the amount of soft or hard storage capacity currently allocated to storage pools and provides alerts when certain capacity thresholds are reached. As the physical, or hard, capacity consumed by volumes within a storage pool passes certain thresholds, the color of this meter indicates that additional hard capacity might need to be added to one or more storage pools.
 - The second indicator (in the middle) displays the number of I/O operations per second (IOPS).
 - The third indicator on the far right shows the general system status and will, for example, indicate when a redistribution is underway.

Additionally, an Uncleared Event indicator is visible when events occur for which a repetitive notification was defined that has not yet been cleared in the GUI (these notifications are called *Alerting Events*).




Figure 4-12 Menu items in XIV Storage Management software

Tip: The configuration information regarding the connected systems and the GUI itself is stored in various files under the user's home directory.

As a useful and convenient feature, all the actions performed from the GUI are saved in a log in the format of XCLI syntax. The syntax includes quoted strings; however, the quotes are only needed if the value specified contains blanks.

The default location is in the "Documents and Settings" folder of the Microsoft Windows current user, for example:

```
%HOMEDRIVE%%HOMEPATH%\Application Data\XIV\GUI10\logs\guiCommands*.log
```

Recently issued commands during the current GUI session can also be accessed through the Tools menu drop-down under "Commands Log."

XIV Storage Management GUI features

There are various features that enhance the usability of the GUI menus. These features enable the user to quickly and easily execute tasks.

As shown in Figure 4-13, commands can be executed against multiple selected objects.




Figure 4-13 Multiple Objects commands

As shown in Figure 4-14, menu tips are now displayed when placing the mouse cursor over greyed out menu items, explaining why they are not selectable in a given context.




Figure 4-14 Menu Tool Tips

Figure 4-15 illustrates keyboard navigation shortcuts.




Figure 4-15 Keyboard Navigation

Additional keyboard shortcuts are listed under 'GUI Keyboard Shortcuts' in the Help menu bar drop-down.

Figure 4-16 illustrates the use of - and + icons to collapse or expand tree views.




Figure 4-16 Collapse/Expand entire tree

In multiple system environments, the main GUI can now register and display up to 64 XIV systems using a group management system (capability for 8 user-defined groups containing up to 8 systems). The GUI also offers automatic grouping based on system status or version. Systems are displayed in a circular arrangement as seen in Figure 4-17.




Figure 4-17 Systems window

As shown in Figure 4-18, mirroring targets are managed in a star view topology that offers the user a convenient way to select which system targets to configure and monitor connectivity status.




Figure 4-18 Mirroring Connectivity 'star' view

Another useful feature is the ability of the GUI to autodetect target connectivity and let the user select which individual system connections to view, as illustrated Figure 4-19.




Figure 4-19 Improved target connectivity

4.2.2 Logging on to the system with XCLI

After the installation of the XIV Storage Management software, you will find the XCLI Session Link on the desktop. Usage of the XCLI Session environment is simple and has an interactive help for command syntax. XCLI command examples are given in this section.

There are various methods of invoking the XCLI functions:

- ▶ **XCLI Session:** Click the desktop link or use the drop-down menu from the Systems Menu in the *GUI* (as shown in Figure 4-20). Starting from the GUI will automatically provide the current userid and password and connect to the system selected. Otherwise you will be prompted for user information and the IP address of the system.

Tip: The XCLI Session is the easiest way to issue XCLI commands against *XIV systems*, and is the best method to use.




Figure 4-20 Launching XCLI from GUI

- ▶ **Invoking the XCLI in order to define configurations:** A configuration is a mapping between a user-defined name and a list of three IP addresses. This configuration can be referenced later in order to execute a command without having to specify the system IP addresses (see the following command execution method in this list).

These various configurations are stored on the local host running the XCLI utility and must be defined for each host. This system name can also be used for XCLI Sessions.

- ▶ **Invoking the XCLI to execute a command:** This method is the most basic and is used in scripts that contain XCLI commands. When invoking an XCLI command directly or in a script, you must also provide either the system's IP addresses or a configuration name.
- ▶ **Invoking the XCLI for general purpose functions:** These invocations can be used to get the XCLI's software version or to print the XCLI's help text.

The command to execute is generally specified along with parameters and their values.

A *script* can be defined to specify the name and path to the commands file (lists of commands will be executed in User Mode only).

For complete and detailed documentation of the IBM XIV Storage Manager Software, see the *Product Overview*, GA32-0791 and the *IBM XIV Storage System XCLI Reference Guide*, GC27-2213, which can be found in the IBM XIV Storage System Information Center:

<http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp>

XCLI Session features

XCLI Session offers command and argument completions, along with possible values for the arguments. Within a session, there is no need to enter user information or IP addresses for each command:

- ▶ Executing a command: Simply type the command.
- ▶ Command completion: Type part of a command and press Tab to see possible valid commands.
- ▶ Command argument completion: Type a command and press Tab to see a list of values for the command argument.

```
>> user_ <TAB>
user_define      user_delete      user_group_add_user  user_group_create
user_group_delete user_group_list  user_group_remove_user user_group_rename
user_group_update user_list        user_rename          user_update

>> user_list <TAB>
show_users=      user=

>> user_list user= <TAB>
xiv_development  xiv_maintenance  admin    technician
ITS0

>> user_list user=admin
Name  Category      Group Active Email Address Area Code Phone Number Access
All
admin storageadmin      yes
```

Figure 4-21 XCLI Session example

Customizing the XCLI environment

For convenience and more efficiency in using the XCLI, use the XCLI Session environment and invoke XCLI Session from the GUI menu. However, if you want to write scripts to execute XCLI commands, it is possible to customize your management workstation environment.

As part of XIV's high-availability features, each system is assigned three IP addresses. When executing a command, the XCLI utility is provided with these three IP addresses and tries each of them sequentially until communication with one of the IP addresses is successful. You must pass at least one of the IP addresses (IP1, IP2, and IP3) with each command.

In order to avoid too much typing and having to remember IP addresses, you can use a predefined configuration name. By default, XCLI uses the system configurations defined when adding systems to the XIV GUI. To list the current configurations, use the command shown in Example 4-1.

Example 4-1 List Configurations XCLI command

```
C:\> xcli -L
```

System	Managements IPs
XIV LAB 01 EBC	192.168.0.1, 192.168.0.2, 192.168.0.3
XIV LAB 3 1300203	192.168.0.4, 192.168.0.5, 192.168.0.6

Commands: When executing a command, you must specify either a configuration or IP addresses, but not both.

To issue a command against a specific IBM XIV Storage System, you also need to supply the username and the password for it. The default user is `admin` and the default password is `adminadmin`, which can be used with the following parameters:

`-u user` or `-user`

- This sets the user name that will be used to execute the command.

`-p password` or `-password`

- This is the XCLI password that must be specified in order to execute a command in the system.

`-m IP1 [-m IP2 [-m IP3]]`

- This defines the IP addresses of the IBM XIV Storage System.

Example 4-2 illustrates a common command execution syntax on a given IBM XIV Storage System.

Example 4-2 Simple XCLI command

```
xcli -u admin -p adminadmin -m 192.168.0.1 user_list
```

Managing the IBM XIV Storage System by using the XCLI always requires that you specify these same parameters. To avoid repetitive typing, you can instead define and use specific environment variables. Create a batch file in which you set the value for these specific environment variables, as shown in Example 4-3.

Example 4-3 Script file setup commands

```
@echo off
set XIV_XCLIUSER=admin
set XIV_XCLIPASSWORD=adminadmin
REM add the following command only to change the default xiv-systems.xml file
REM set XCLI_CONFIG_FILE=%HOMEDRIVE%%HOMEPATH%\My Documents\xcli\xiv-systems.xml
REM List the current configuration
xcli -L
```

The XCLI utility requires user and password options. If user and passwords are not specified, the default environment variables `XIV_XCLIUSER` and `XIV_XCLIPASSWORD` are utilized.

The configurations are stored in a file under the user's home directory. A separate file can be specified by `-f` or `--file` switch (applicable to configuration creation, configuration deletion, listing configurations, and command execution). Alternatively, the environment variable `XCLI_CONFIG_FILE`, if defined, determines the file's name and path. The default file is in `%HOMEDRIVE%%HOMEPATH%\Application Data\XIV\GUI10\properties\xiv-systems.xml`.

After executing the setup commands, the shortened command syntax works as shown in Example 4-4.

Example 4-4 Short command syntax

```
REM S1 can be used in all commands to save typing and script editing
set S1="XIV LAB 01 EBC"
xccli -c %S1% user_list
```

Getting help with XCLI commands

To get help about the usage and commands (assuming that you have set and logged in as a valid user), proceed as shown in Example 4-5.

Example 4-5 XCLI help commands

```
xccli
xccli -c "XIV LAB 01 EBC" help
xccli -c "XIV LAB 01 EBC" help command=user_list format=full
```

The first command prints out the usage of **xccli**. The second one prints all the commands that can be used by the user in that particular system. The third one shows the usage of the **user_list** command with all the parameters.

There are various parameters to get the result of a command in a predefined format. The default is the user readable format. Specify the **-s** parameter to get it in a comma-separated format or specify the **-x** parameter to obtain an XML format.

Fields: The XML format contains all the fields of a particular command. The user and the comma-separated formats provide just the default fields as a result.

To list the field names for a specific **xccli** command output, use the **-t** parameter as shown in Example 4-6.

Example 4-6 XCLI field names

```
xccli -c "XIV LAB 01 EBC" -t name,fields help command=user_list
```

4.3 Storage pools

We have introduced the concept of XIV storage pools in 2.3.4, "Storage pool concepts" on page 24.

4.3.1 Function of storage pools

Storage pools function as a means to effectively manage a related group of logical volumes and their snapshots. Storage pools offer the following key benefits:

- ▶ Improved management of storage space: Specific volumes can be grouped within a storage pool, giving you the flexibility to control the usage of storage space by specific applications, a group of applications, or departments.
- ▶ Improved regulation of storage space: Automatic snapshot deletion occurs when the storage capacity limit is reached for each storage pool independently. Therefore, when a storage pool's size is exhausted, only the snapshots that reside in the affected storage pool are deleted.

Size of storage pools

The size of storage pools and the associations between volumes and storage pools are constrained by the following circumstances:

- ▶ The size of a storage pool can range from as small as possible (17.1 GB) to as large as 79 TB.
- ▶ The size of a storage pool can always be increased up to the 79 TB limit.
- ▶ The size of a storage pool can always be decreased, limited only by the space already consumed by the volumes and snapshots in that storage pool.
- ▶ Volumes can be moved between storage pools without any limitations, as long as they are not part of a consistency group and there is enough free space in the target storage pool. Volumes that are part of a consistency group can be moved together as a group.

Important: All of these operations are handled by the system at the metadata level, and they do not cause any data movement (copying) from one disk drive to another. Hence, they are completed almost instantly and can be done at any time without impacting the applications.

Thin provisioned pools

Thin provisioning is the practice of allocating storage on a “just-in-time” and “as needed” basis by defining a logical, or soft, capacity that is larger than the physical, or hard, capacity. Thin provisioning enables IBM XIV Storage System administrators to manage capacity based on the total space actually consumed rather than just the space allocated.

Thin provisioning can be specified at the storage pool level. Each thinly provisioned pool has its own hard capacity (which limits the actual disk space that can be effectively consumed) and soft capacity (which limits the total logical size of volumes defined).

The difference is in the pool size:

- ▶ *Hard pool size:* The hard pool size represents the physical storage capacity allocated to volumes and snapshots in the storage pool. The hard size of the storage pool limits the total of the hard volume sizes of all volumes in the storage pool and the total of all storage consumed by snapshots.
- ▶ *Soft pool size:* This size is the limit on the total soft sizes of all the volumes in the storage pool. The soft pool size has no effect on snapshots.

For more detailed information about the concept of XIV thin provisioning and a detailed discussion of hard and soft size for storage pools and volumes, see 2.3.5, “Capacity allocation and thin provisioning” on page 26.

When using the GUI, you specify what type of pool is desired (*Regular Pool* or a *Thin Provisioning Pool*) when creating the pool. See “Creating storage pools” on page 105. When using the XCLI, you create a thinly provisioned pool by setting the soft size to a greater value than its hard size.

In case of changing requirements, the pool’s type can be changed (non-disruptively) later.

Tip: The thin provisioning management is performed individually for each storage pool, and running out of space in one pool does not impact other pools.

4.3.2 Managing storage pools with the XIV GUI

In this section we describe how to manage storage pools using the XIV GUI.

Important: Illustrations in this chapter mostly apply to an IBM XIV Storage System fully configured with 1 TB drives.

Managing pools with the GUI is fairly simple and intuitive. As always, the related tasks can be reached by either the menu bar or the corresponding function icon on the left (called Pools), as shown in Figure 4-22.



Figure 4-22 Opening the Pools menu

To view overall information about the storage pools, select **Pools** from the Pools menu shown in Figure 4-22 to display the Storage Pools window seen in Figure 4-23.




Figure 4-23 Storage Pools view

The Storage Pools GUI window displays a table of all the pools in the system combined with a series of gauges for each pool. This view gives the administrator a quick grasp and general overview of essential information about the system pools.

The capacity consumption by volumes and snapshots within a given storage pool is indicated by various colors. The default threshold values are described next:

- ▶ Green is the indicator for consumed capacity below 80%.
- ▶ Yellow represents a capacity consumption above 80%.
- ▶ Orange is the color for a capacity consumption of over 90%.
- ▶ Any storage pool with depleted hard capacity displays in red within this view.

The name, the size, and the separated segments are labeled appropriately. Figure 4-24 shows the meaning of the various numbers.




Figure 4-24 Storage pool and size numbers

Creating storage pools

The creation and resizing of storage pools is relatively straightforward, and you only need to take care with the size allocation and re-allocation. The name of a storage pool must be unique in the system.

Pool size: The size of the storage pool is specified as an integer multiple of 10^9 bytes, but the actual size of the created storage pool is rounded up to the nearest integer multiple of 16×2^{30} bytes. According to this rule, the smallest pool size is 17.1 GB.

When creating a storage pool, the system initially provides a default snapshot size of zero. However, you need to specify a value between 34 GB or very conservatively up to 80% of the volume size, depending on how snapshots will be utilized, your workload characteristics, and your requirements for how long snapshots need to be kept. This value can be set at the time of creation or changed later depending on your needs. In any case, your IBM representative can help you determine the appropriate ratio for your environment.

Snapshot size: The snapshot size (default or specified) is included in the specified pool size. It is not an additional space.

Sizing must take into consideration volumes that are to be added to (or already exist in) the specific storage pool, the current allocation of storage in the total system capacity, and future activity within the storage pool, especially with respect to snapshot propagation resulting from creating too many snapshots. Upon depletion of space in a pool, the system will progressively delete snapshots in the pool to free up space for additional write requests. For more information, see Chapter 1, "Snapshots" in the Redbooks publication, *IBM XIV Storage System: Copy Services and Migration*, SG24-7759.

The system enables the assignment of the entire available capacity to user-created storage pools. The storage pool is initially empty and does not contain volumes. However, you cannot create a storage pool with zero capacity.

To create a storage pool:

1. Click **Add Pool** in the Storage Pools view or simply right-click in the body of the window. An Add Pool window displays as shown in Figure 4-25.




Figure 4-25 Add Pool

2. In the Select Type drop-down list box, choose **Regular Pool** or **Thin Provisioning Pool** according to your needs. For thinly provisioned pools, new fields appear:
 - Pool Hard Size: Here, you specify the upper limit of hard capacity.
 - Pool Soft Size: Here, you specify the upper limit of soft capacity.
 - Lock Behavior: Here, you specify the behavior in case of depleted capacity.
This value specifies whether the storage pool is locked for write or whether it is disabled for both read and write when running out of storage space. The default value is *Read only*.
3. In the Pool Size field, specify the required size of the storage pool.
4. In the Snapshots Size field, enter the required size of the reserved snapshot area.

Tip: Although it is possible to create a pool with identical snapshot and pool size, you cannot create a new volume in this type of a pool afterward without resizing first.

5. In the Pool Name field, enter the desired name (it must be unique across the Storage System) for the storage pool.
6. Click **Add** to add this storage pool.

Resizing storage pools

This action can be used to both increase or decrease a storage pool size. Capacity calculation is performed in respect to the total system net capacity. All reductions and increases are reflected in the remaining free storage capacity.

Storage pool size:

- ▶ When increasing a storage pool size, you must ensure that the total system capacity holds enough free space to enable the increase in storage pool size.
- ▶ When decreasing a storage pool size, you must ensure that the storage pool itself holds enough free capacity to enable a reduction in size.

This operation is also used to shrink or increase the snapshot capacity inside the storage pool. This alteration only affects the space within the storage pool. In other words, increasing snapshot size will consume the free capacity only from the corresponding pool.

To change the size of one storage pool in the system, simply right-click in the Storage Pools view (Figure 4-23 on page 104) on the desired pool and choose **Resize**.

The window shown in Figure 4-26 is displayed. Change the pool hard size, soft size, or the snapshot size to match your new requirements. The green bar at the top of the window represents the system's actual hard capacity that is used by storage pools. The vertical red line indicates the current size, and the yellow part is the desired new size of the particular pool. The remaining space in the bar without color indicates the consumable free capacity in the system.




Figure 4-26 Resizing pool

The resize operation can also be used to change the type of storage pool from thin provisioned to regular or from regular to thin provisioned (see Figure 4-27). Just change the type of pool in the Resize Pool window Select Type list box:

- ▶ When a regular pool is converted to a thin provisioned pool, you have to specify an additional soft pool size parameter besides the existing hard pool size. The soft pool size must be greater than the hard pool size.
- ▶ When a thin provisioned pool is changed to a regular pool, the soft pool size parameter will disappear from the window; in fact, its value will be equal to the hard pool size.

If the space consumed by existing volumes exceeds the pool's actual hard size, the pool cannot be changed to a regular type pool. In this case, you have to specify a minimum hard pool size equal to the total capacity consumed by all the volumes within this pool.




Figure 4-27 Resizing and changing the type of a pool

The remaining soft capacity is displayed in red characters and calculated by the system in the following manner:

Remaining Soft Capacity = [Current Storage Pool Soft Size + Remaining System Soft Size] - Current Storage Pool Hard Size

Deleting storage pools

To delete a storage pool, right-click the storage pool and select **Delete**. The system will ask for a confirmation before deleting this storage pool.

The capacity of the deleted storage pool is reassigned to the system's free capacity, which means that the free hard capacity increases by the size of the deleted storage pool.

Restriction: You cannot delete a storage pool if it still contains volumes.

Moving volumes between storage pools

In order for a volume to be moved to a specific storage pool, there must be enough room for the volume to reside there. If there is not enough free capacity (meaning that adequate capacity has not been allocated), the storage pool must be resized, or other volumes must be moved out first to make room for the new volume.

When moving a master volume from one storage pool to another, all of its snapshots are moved along with it to the destination storage pool. You cannot move a snapshot alone, independently of its master volume.

The destination storage pool must have enough free storage capacity to accommodate the volume and its snapshots. The exact amount of storage capacity allocated from the destination storage pool is released at the source storage pool.

A volume that belongs to a consistency group cannot be moved without moving the entire consistency group.

As shown in Figure 4-28, in the Volume by Pools report, just select the appropriate volume with a right-click and initiate a Move to Pool operation to change the location of a volume.



Figure 4-28 Volumes by Pools

In the pop-up window, select the appropriate storage pool as shown in Figure 4-29 and click **OK** to move the volume into it.




Figure 4-29 Move volume to another pool

Pool alert thresholds

You can use the XIV GUI to configure pool usage thresholds to trigger color-coded visual alerts at various severity levels. These alerts are system based across all pools and viewable in the GUI by all users configured with the particular system. System events are also produced based on specified values. The default pool usage threshold values for warning, minor, and major are 80%, 90%, and 95% respectively.

From the main GUI management window, select **Tools** → **Configure** → **Pool Alerts Thresholds** or select the **Configure Pool Thresholds** from the **Storage Pools** view to get the menu shown in Figure 4-30.




Figure 4-30 Set pool alert thresholds

Events: The XIV System has a number of predefined events that are triggered when system pool usage is depleting. Rules can be set to receive email alerts for these events. See Chapter 7, “Monitoring” on page 203 for more information about setting up email event notifications.

4.3.3 Managing storage pools with XCLI

All of the operations described in 4.3.2, “Managing storage pools with the XIV GUI” on page 104, can also be done through the command line interface.

Listing pool-related commands

To get a list of all the storage pool-related XCLI commands, type the following command from the XCLI command shell:

```
help category=storage-pool
```

Important: Note that the commands shown in this section assume that you have started an XCLI Session to the system selected; see “XCLI Session features” on page 100.

The output shown in Example 4-7 is displayed.

Example 4-7 All the storage pool-related commands

Category	Name	Description
storage-pool	cg_move	Moves a Consistency Group, all its volumes and all their snapshots and Snapshot Sets from one Storage Pool to another.
storage-pool	pool_change_config	Changes the lock behavior of the Storage Pool.
storage-pool	pool_config_snapshots	Changes the Storage Pool snapshot settings.
storage-pool	pool_create	Creates a Storage Pool.
storage-pool	pool_delete	Deletes a Storage Pool.
storage-pool	pool_list	Lists all Storage Pools or the specified one.
storage-pool	pool_rename	Renames a specified Storage Pool.
storage-pool	pool_resize	Resizes a Storage Pool.
storage-pool	vol_move	Moves a volume and all its snapshot from one Storage Pool to another.

To list the existing storage pools in a system, use the following command:

```
pool_list
```

A sample result of this command is illustrated in Figure 4-31.

.	Name	Size(GB)	Hard Size(GB)	Snapshot Size(GB)	Empty Space(GB)	Used by Volumes(GB)	Used by Snapshots(GB)	Locked
	test_pool	76948	71829	29068	7198	15633	0	no
	testing_cim	154	154	17	103	0	0	no
	ITS0 Pool 1	1013	807	206	807	0	0	no
	...							

Figure 4-31 Result of the pool_list command

For the purpose of new pool creation, enter the following command:

```
pool_create pool="ITS0 Pool 2" size=515 snapshot_size=103
```

The size of the storage pool is specified as an integer multiple of 10^9 bytes, but the actual size of the created storage pool is rounded up to the nearest integer multiple of 16×2^{30} bytes.

The snapshot_size parameter specifies the size of the snapshot area within the pool. It is a mandatory parameter, and you must specify a positive integer value for it. Allocate a minimum of 34 GB of snapshot space. A value of 80% of the volume space is best.

The following command shows how to resize one of the existing pools:

```
pool_resize pool="ITS0 Pool 2" size=704 snapshot_size=154
```

With this command, you can increase or decrease the pool size. The **pool_create** and the **pool_resize** commands are also used to manage the size of the snapshot area within a storage pool.

To rename an existing pool, issue this command:

```
pool_rename pool="ITS0 Pool 1" new_name="ITS0 Pool "
```

To delete a pool, type:

```
pool_delete pool="ITS0 Pool "
```

Approve or deny deletion by responding y/n when prompted or use the -y parameter with the **pool_delete** command to approve deletion.

Use the following command to move the volume named **cim_test_02** to the storage pool **ITS0 Pool 2**:

```
vol_move pool="ITS0 Pool 2" vol="cim_test_02"
```

The command only succeeds if the destination storage pool has enough free storage capacity to accommodate the volume and its snapshots.

The aforementioned command will move a particular volume and its snapshots from one storage pool to another, but if the volume is part of a consistency group, the entire group must be moved. In this case, the **cg_move** command is the correct solution:

```
cg_move cg="ITS0 CG" pool="ITS0 Pool 2"
```

All volumes, volume snapshots, and snapshot groups of the consistency group are moved.

Thin provisioned pools

To create thinly provisioned pools, specify the **hard_size** and the **soft_size** parameters. For thin provisioning concepts, see 2.3.5, “Capacity allocation and thin provisioning” on page 26.

A typical storage pool creation command with thin provisioning parameters can be issued as shown in the following example:

```
pool_create pool="ITS0 Pool" hard_size=807 soft_size=1013 lock_behavior=read_only  
snapshot_size=206
```

The **soft_size** is the maximal storage capacity seen by the host and cannot be smaller than the **hard_size**, which is the hard physical capacity of the storage pool.

If a storage pool runs out of hard capacity, all of its volumes are locked to all write commands. Even though write commands that overwrite existing data can be technically serviced, they are blocked as well in order to ensure consistency.

To specify the behavior in case of depleted capacity reserves in a thin provisioned pool, use the following command:

```
pool_change_config pool="ITS0 Pool" lock_behavior=no_io
```

This command specifies whether the storage pool is locked for write or whether it disables both read and write when running out of storage space.

Lock: The **lock_behavior** parameter can be specified for non-thin provisioning pools, but it has no effect.

4.4 Volumes

After defining storage pools, the next step in the IBM XIV Storage System configuration is volume management.

The IBM XIV Storage System offers logical volumes as the basic data storage element for allocating usable storage space to attached hosts. This logical unit concept is well known and is widely used by other storage subsystems and vendors. However, neither the volume segmentation nor its distribution over the physical disks is conventional in the IBM XIV Storage System.

Traditionally, logical volumes are defined within various RAID arrays, where their segmentation and distribution are manually specified. The result is often a suboptimal distribution within and across modules (expansion units) and is significantly dependent upon the administrator's knowledge and expertise.

As explained in 2.3, "Full storage virtualization" on page 18, the IBM XIV Storage System uses true virtualization as one of the basic principles for its unique design. With XIV, each volume is divided into tiny 1 MB partitions, and these partitions are distributed randomly and evenly, and duplicated for protection. The result is optimal distribution in and across all modules, which means that for any volume, the physical drive location and data placement are invisible to the user. This method dramatically simplifies storage provisioning, letting the system lay out the user's volume in an optimal way.

This method offers complete virtualization, without requiring preliminary volume layout planning or detailed and accurate stripe or block size pre-calculation by the administrator. All disks are equally used to maximize the I/O performance and exploit all the processing power and all the bandwidth available in the storage system.

IBM XIV Storage System virtualization incorporates an advanced snapshot mechanism with unique capabilities, which enables creating a virtually unlimited number of point-in-time copies of any volume, without incurring any performance penalties. The concept of snapshots is discussed in detail in the *Product overview*, GA32-0791, available for download from the IBM XIV Storage System Information Center:

<http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp>

Volumes can also be grouped into larger sets called *consistency groups* and *storage pools*. See 4.3, "Storage pools" on page 102.

Important: As shown in Figure 4-32, the basic hierarchy is as follows:

- ▶ A volume can have multiple snapshots.
- ▶ A volume can be part of one and only one consistency group.
- ▶ A volume and its corresponding snapshots are always part of one and only one storage pool.
- ▶ All volumes of a consistency group must belong to the same storage pool.




Figure 4-32 Basic storage hierarchy

4.4.1 Managing volumes with the XIV GUI

To start a volume management function from the XIV GUI, you can either select **View** → **Volumes** → **Volumes and Snapshots** from the menu bar or click the **Volumes** icon and then select the appropriate menu item. See Figure 4-33.



Figure 4-33 Opening the Volumes menu

The Volumes & Snapshots menu item is used to list all the volumes and snapshots that have been defined in this particular IBM XIV Storage System. An example of the resulting window can be seen in Figure 4-34.

Name	Size	Used (GB)	Master	Consistency	Pool	Created	Serial
vmesx35_test_1	2010	2005			vmesx_35_...		1025
vmesx35_test_10	2010	2009			vmesx_35_...		1026
vmesx35_test_2	2010	1457			vmesx_35_...		1027
sofs_prf0	1992	104			sofs_perf		812
p570_sapdata_1vol	1013	351			SAP_on_XL...		813
hursley1000	1013	10			Hursley_KK...		1024
Conrad	1013	6			Conrad		1275
Conrad.snapshot_00001	1013		Conrad		Conrad	2010-05-28 10:49	10163
Conrad.snapshot_00002	1013		Conrad		Conrad	2010-06-07 12:54	5277
Conrad.snapshot_00003	1013		Conrad		Conrad	2010-06-07 12:54	5278
Conrad.snapshot_00004	1013		Conrad		Conrad	2010-06-07 12:55	5280
sapISR	704	511			HGH_SAP_...		12534
sapISR.snap_25june_db2	704		sapISR		HGH_SAP_...	2010-06-25 13:48	8805
sapISRora	704	543			HGH_SAP_...		5367
sapISRdb2_asynch	704	512			HGH_SAP_...		8807
last-replicated-sapISRdb2...	704		sapISR...		HGH_SAP_...	2010-07-28 14:31	14939
sapISRora_asynch	704	547			HGH_SAP_...		8808
last-replicated-sapISRora...	704		sapISR...		HGH_SAP_...	2010-09-24 14:07	6153
hursley600	601	66			Hursley_KK...		261
test	601	0			at12677_js...		3465
test.snapshot_00001	601		test		at12677_js...	2010-09-22 20:49	6940
hursley500	515	69			Hursley_KK...		175
hursley_BlueArc500	515	133			Hursley_KK...		1180
GTS_Large_Linux	515	83			1_Sales_thin		1347
SNGPL_Demo_1	515	0			SNGPL_1		2990

Figure 4-34 Volumes and Snapshots view

Volumes are listed in a tabular format. If the volume has snapshots, then a + or a - icon displays on the left. Snapshots are listed under their master volumes, and the list can be expanded or collapsed at the volume level by clicking the + or - icon respectively.

Snapshots are listed as a sub-branch of the volume of which they are a replica, and their row is indented and highlighted in off-white.

The Master column of a snapshot shows the name of the volume of which it is a replica. If this column is empty, the volume is the master.

Tip: To customize the columns in the lists, just right-click one of the column headings and make the required selection of attributes. The default column set does not contain the Master column. You can also resize the columns to allow for longer names or to make more columns visible.

Table 4-1 shows columns of this view with their descriptions.

Table 4-1 Columns in the Volumes and Snapshots view

Column	Description	Default
Name	Name of a volume or snapshot	Y
Size (GB)	Volume or snapshot size. (value is zero if the volume is specified in blocks)	Y
Used (GB)	Used capacity in a volume	Y
Size (Blocks)	Volume size in blocks	N
Used Capacity (MiB)	Consumed capacity	N
Master	Snapshot master's name	Y
Consistency Group	Consistency group name	Y
Pool	Storage pool name	Y
Locked Status	Indicates the locking status of a volume or snapshot. Lock icon.	Y
Snapshot Modification Status	Shows if the snapshot was unlocked or modified	Y
Deletion Priority	Indicates the priority of deletion by numbers for snapshots	Y
Created	Shows the creation time of a snapshot	N
Creator	Volume or snapshot creator name	N
Serial Number	Volume or snapshot serial number	N
Coupling Status	Shows the mirroring type status	Y

Most of the volume-related and snapshot-related actions can be selected by right-clicking any row in the table to display a drop-down menu of options. The options in the menu differ slightly for volumes and snapshots.

Menu option actions

The following actions can be performed through these menu options:

- ▶ Adding or creating volumes; see “Creating volumes” on page 117
- ▶ Resizing a volume; see “Resizing volumes” on page 121
- ▶ Deleting a volume or snapshot; see “Deleting volumes” on page 122
- ▶ Formatting a volume
- ▶ Renaming a volume or snapshot
- ▶ Creating a consistency group with these volumes
- ▶ Adding to a consistency group
- ▶ Removing from a consistency group
- ▶ Moving volumes between storage pools; see “Moving volumes between storage pools” on page 109

- ▶ Creating a snapshot
- ▶ Creating a snapshot (advanced)
- ▶ Overwriting a snapshot
- ▶ Copying a volume or snapshot
- ▶ Locking/unlocking a volume or snapshot
- ▶ Mappings
- ▶ Displaying properties of a volume or snapshot
- ▶ Changing a snapshot's deletion priority
- ▶ Duplicating a snapshot or a snapshot (advanced)
- ▶ Restoring from a snapshot

Creating volumes

When you create a volume in a *traditional or regular storage pool*, the entire volume storage capacity is reserved (static allocation). In other words, you cannot define more space for volumes in a regular storage pool than the actual hard capacity of the pool, which guarantees the functionality and integrity of the volume.

When you create a volume in a *thin provisioned pool*, the capacity of the volume will not be reserved immediately to the volumes, but a basic 17.1 GB piece, taken out of the storage pool hard capacity, will be allocated at the first I/O operation. In a thin provisioned pool, you are able to define more space for volumes than the actual hard capacity of the pool, up to the soft size of the pool.

The volume size is the actual “net” storage space, as seen by the host applications, not including any mirroring or other data protection overhead. The free space consumed by the volume will be the smallest multiple of 17 GB that is greater than the specified size. For example, if we request an 18 GB volume to be created, the system will round this volume size to 34 GB. In case of a 16 GB volume size request, it will be rounded to 17 GB.

Figure 4-35 gives you various basic examples of volume definition and planning in a thinly provisioned pool. It depicts the volumes with the minimum amount of capacity, but the principle can be used for larger volumes as well.

As shown in this picture, plan carefully the number of volumes or the hard size of the thinly provisioned pool because of the minimum hard capacity that is consumed by one volume.

If you create more volumes in a thinly provisioned pool than the hard capacity can cover, the I/O operations against the volumes will fail at the first I/O attempt.

Volumes: Plan the volumes in a thin provisioned pool in accordance with this formula:
 $\text{Pool Hard Size} \geq 17 \text{ GB} \times (\text{number of volumes in the pool})$




Figure 4-35 Planning the number of volumes in a thin provisioning pool

The size of a volume can be specified either in gigabytes (GB) or in blocks (where each block is 512 bytes). If the size is specified in blocks, volumes are created in the exact size specified, and the size will be not rounded up. It means that the volume will show the exact block size and capacity to the hosts but will nevertheless consume a 17 GB size in the IBM XIV Storage System. This capability is relevant and useful in migration scenarios.

If the size is specified in gigabytes, the actual volume size is rounded up to the nearest 17.1 GB multiple (making the actual size identical to the free space consumed by the volume, as just described). This rounding up prevents a situation where storage space is not fully utilized because of a gap between the free space used and the space available to the application.

The volume is logically formatted at creation time, which means that any read operation results in returning all zeros as a response.

To create volumes with the XIV Storage Management GUI:

1. Click the “add volumes” icon in the Volume and Snapshots view (Figure 4-34 on page 115) or right-click in the body of the window (not on a volume or snapshot) and select Add Volumes. The window shown in Figure 4-36 is displayed.
2. From the Select Pool field, select the Pool in which this volume is stored. See 4.3, “Storage pools” on page 102 for a description of how to define storage pools. The storage size and allocation of the selected storage pool is shown textually and graphically in a color-coded bar:
 - Green indicates the space already allocated in this storage pool.
 - Yellow indicates the space that will be allocated to this volume (or volumes) after it is created.
 - Gray indicates the space that remains free after this volume (or volumes) is allocated.

Figure 4-36 Create Volumes

3. In the Number of Volumes field, specify the required number of volumes.
4. In the Volume Size field, specify the size of each volume to define. The size can also be modified by dragging the yellow part of the size indicator.

Volume size: When multiple volumes are created, they all have the same size as specified in the Volume Size field.

- 5. In the Volume Name field, specify the name of the volume to define. The name of the volume must be unique in the system. If you specified that more than one volume be defined, they are successively named by appending an incrementing number to end of the specified name. You can also add an initial sequence number.
- 6. Click **Create** to effectively create and add the volumes to the storage pool (Figure 4-37).




Figure 4-37 Volume creation progress indicator

After a volume is successfully added, its state is unlocked, meaning that write, format, and resize operations are permitted. The creation time of the volume is set to the current time and is never changed. Notice the volume name sequence in Figure 4-38.

All Systems (View By My Groups) > ITSO XIV Redbook > XIV LAB 01 EBC > Volumes by Pools					
Used	Name	Size (GB)	Used (GB)	Consiste	
0 %	ITSO Pool	807			
	vol_201	51	0		
	vol_202	51	0		
	vol_203	51	0		
	vol_204	51	0		
	vol_205	51	0		
	vol_206	51	0		
	vol_207	51	0		
	vol_208	51	0		
	vol_209	51	0		
	vol_210	51	0		
0 %	itso_test_pool	51			

Figure 4-38 Volumes created

Resizing volumes

Resizing volumes is an operation very similar to their creation. Only an unlocked volume can be resized. When you resize a volume, its size is specified as an integer multiple of 10^9 bytes, but the actual new size of the volume is rounded up to the nearest valid size, which is an integer multiple of 17 GB.

Attention: The size of the volume can be decreased with the GUI. However, to avoid possible data loss, you must contact your IBM XIV support personnel if you need to decrease a volume size. (Mapped volume size cannot be decreased.)

The volume address space is extended (at the end of the existing volume) to reflect the increased size, and the additional capacity is logically formatted (that is, zeroes are returned for all read commands).

When resizing a regular volume (not a writable snapshot), all storage space that is required to support the additional volume capacity is reserved (static allocation), which guarantees the functionality and integrity of the volume, regardless of the resource levels of the storage pool containing that volume.

Resizing a master volume does not change the size of its associated snapshots. These snapshots can still be used to restore their individual master volumes at their initial sizes.




Figure 4-39 Resize an existing volume

To resize volumes with XIV Storage Management GUI:

1. Right-click the row of the volume to be resized and select **Resize**.

The total amount of storage is presented both textually and graphically. The amount that is already allocated by the other existing volumes is shown in green. The amount that is free is shown in gray. The current size of the volume is displayed in yellow, to the left of a red vertical bar. This red bar provides a constant indication of the original size of the volume as you resize it. Place the mouse cursor over the red bar to display the volume's initial size.

2. In the New Size field, use the arrows to set the new size or type the new value.
3. Click Update to resize the volume.

Deleting volumes

With the GUI, the deletion of a volume is as easy as creating one.

Important: After you delete a volume or a snapshot, all data stored on the volume is lost and cannot be restored.

All the storage space that was allocated (or reserved) for the volume or snapshot is freed and returned to its storage pool. The volume or snapshot is then removed from all the logical unit number (LUN) maps that contain mapping of this volume.

Deleting a volume deletes all the snapshots associated with this volume, even snapshots that are part of snapshot groups. This deletion can only happen when the volume was in a consistency group and was removed from it.

You can delete a volume regardless of the volume's lock state, but you cannot delete a volume that is part of a consistency group.

To delete a volume or a snapshot:

1. Right-click the row of the volume to be deleted and select **Delete**.
2. Click to delete the volume.

Maintaining volumes

There are various other operations that can be issued on a volume. See “Menu option actions” on page 116.

The usage of these operations is obvious, and you can initiate an operation with a right-mouse click. These operations are:

- ▶ *Format a volume:* A formatted volume returns zeros as a response to any read command. The formatting of the volume is done logically, and no data is actually written to the physical storage space allocated for the volume. Consequently, the formatting action is performed instantly.
- ▶ *Rename a volume:* A volume can be renamed to a unique name in the system. A locked volume can also be renamed.
- ▶ *Lock/Unlock a volume:* You can lock a volume so that hosts cannot write to it. A volume that is locked is write-protected, so that hosts can read the data stored on it, but they cannot change it. The volume displays then as a lock icon. In addition, a locked volume cannot be formatted or resized. In general, locking a volume prevents any operation (other than deletion) that changes the volume's image.

Lock/unlock: Master volumes are set to *unlocked* when they are created. Snapshots are set to *locked* when they are created.

- ▶ *Consistency groups:* IBM XIV Storage System enables a higher level of volume management provided by grouping volumes and snapshots into sets called *consistency groups*. This kind of grouping is especially useful for cluster-specific volumes.
- ▶ *Copy a volume:* You can copy a source volume onto a target volume. All the data that was previously stored on the target volume is lost and cannot be restored.
- ▶ *Snapshot functions:* The IBM XIV Storage System's advanced snapshot feature has unique capabilities that enable the creation of a virtually unlimited number of copies of any volume, with no performance penalties.

- *Map a volume:* While the storage system sees volumes and snapshots at the time of their creation, the volumes and snapshots are visible to the hosts only after the mapping procedure. To get more information about mapping, see 4.5, “Host definition and mappings” on page 124.

4.4.2 Managing volumes with XCLI

All of the operations that are explained in 4.4.1, “Managing volumes with the XIV GUI” on page 114 can also be performed through the command line interface. To get a list of all the volume-related commands, enter the following command in the XCLI Session:

```
help category=volume
```

Important: Note that the commands shown in this section assume that you have started an XCLI Session to the system selected; see “XCLI Session features” on page 100. Simply replace the example name values to perform the commands.

Example 4-8 shows the output of the command.

Example 4-8 All the volume-related commands

Category	Name	Description
volume	reservation_clear	Clear reservations of a volume.
volume	reservation_key_list	Lists reservation keys.
volume	reservation_list	Lists volume reservations.
volume	vol_by_id	Prints the volume name according to its specified SCSI serial number.
volume	vol_copy	Copies a source volume onto a target volume.
volume	vol_create	Creates a new volume.
volume	vol_delete	Deletes a volume.
volume	vol_format	Formats a volume.
volume	vol_list	Lists all volumes or a specific one.
volume	vol_lock	Locks a volume so that it is read-only.
volume	vol_rename	Renames a volume.
volume	vol_resize	Resizes a volume.
volume	vol_unlock	Unlocks a volume, so that it is no longer read-only and can be written to.

To list the existing volumes in a system, use the following command:

```
vol_list pool="ITS0 Pool"
```

The result of this command is similar to the illustration given in Figure 4-40.

Name	Size (GB)	Master Name	Consistency Group	Pool	Creator	Used Capacity(GB)
vol_201	51			ITS0 Pool	ITS0	0
vol_202	103			ITS0 Pool	ITS0	0
vol_203	51			ITS0 Pool	ITS0	0
vol_204	51			ITS0 Pool	ITS0	0
vol_205	51			ITS0 Pool	ITS0	0
vol_206	51			ITS0 Pool	ITS0	0
vol_207	51			ITS0 Pool	ITS0	0
vol_208	51			ITS0 Pool	ITS0	0
vol_209	51			ITS0 Pool	ITS0	0
vol_210	51			ITS0 Pool	ITS0	0
>>						

Figure 4-40 vol_list command output

To find and list a specific volume by its SCSI ID, issue the following command:

```
vol_by_id id=6716
```

To create a new volume, enter the following command:

```
vol_create size=51 pool="ITS0 Pool" vol="vol_211"
```

The size can be specified either in gigabytes or in blocks (where each block is 512 bytes). If the size is specified in blocks, volumes are created in the exact size specified. If the size is specified in gigabytes, the actual volume size is rounded up to the nearest 17 GB multiple (making the actual size identical to the free space consumed by the volume, as described before). This rounding up prevents a situation where storage space is not fully utilized because of a gap between the free space used and the space available to the application.

The volume is logically formatted at creation time, which means that any read operation results in returning all zeros as a response. To format a volume, use the following command:

```
vol_format vol="vol_201"
```

Note that all data stored on the volume will be lost and unrecoverable. If you want to bypass the warning message, just put -y right after the XCLI command.

The following example shows how to resize one of the existing volumes:

```
vol_resize vol="vol_201" size=103
```

Cannot shrink: If you attempt to decrease volume size with the XCLI, you will receive a "CAN_NOT_SHRINK_VOLUME" message. Create a new volume and migrate the data to the new volume.

To rename an existing volume, issue this command:

```
vol_rename new_name="vol_200" vol="vol_210"
```

To delete an existing created volume, enter:

```
vol_delete vol="vol_200"
```

4.5 Host definition and mappings

Because the IBM XIV Storage System can be attached to multiple, heterogeneous hosts, it is necessary to specify which particular host can access which specific logical drives in the IBM XIV Storage System. In other words, mappings must be defined between hosts and volumes in the IBM XIV Storage System.

The IBM XIV Storage System is able to manage single hosts or hosts grouped together in clusters.

4.5.1 Assigning LUNs to a host using the GUI

There are a number of steps required in order to define a new host and assign LUNs to it. Prerequisites are that volumes have been created in a storage pool.

Defining a host

To define a host, follow these steps:

1. In the IBM XIV Storage System main GUI window, move the mouse cursor over the **Hosts and Clusters** icon and select **Hosts and Clusters** (see Figure 4-41).



Figure 4-41 Hosts and Clusters menu

2. The Hosts window is displayed showing a list of hosts (if any) that are already defined. To add a new host or cluster, click either the **Add Host** or **Add Cluster** in the menu bar (see Figure 4-42). In our example, we select **Add Host**. The difference between the two is that *Add Host* is for a single host that will be assigned a LUN or multiple LUNs, whereas *Add Cluster* is for a group of hosts that will share a LUN or multiple LUNs.



Figure 4-42 Add new host

3. The **Add Host** dialog is displayed as shown in Figure 4-43. Enter a name for the host. If a cluster definition was created in the previous step, it is available in the cluster drop-down list box. To add a server to a cluster, select a cluster name. Because we do not create a cluster in our example, we select **None**. Type, CHAP Name, and CHAP Secret can be specified here as necessary. CHAP is an authentication process of an iSCSI initiator. Because the example host is neither hpux or zvm, Type remains at **default**.




Figure 4-43 Add host details

4. Repeat steps 2 and 3 to create additional hosts. In our scenario we add another host called itso_win2008_iscsi.

- Host access to LUNs is granted depending on the host adapter ID. For an FC connection, the host adapter ID is the FC HBA WWPN; for an iSCSI connection, the host adapter ID is the host IQN. To add a WWPN or IQN to a host definition, right-click the host and select **Add Port** from the context menu (see Figure 4-44).




Figure 4-44 GUI example: Add port to host definition

- The Add Port dialog is displayed as shown in Figure 4-45. Select port type FC or iSCSI.
In this example, an FC host is defined. Add the WWPN for HBA1. If the host is correctly connected and has done a port login to the SAN switch at least once, the WWPN is shown in the drop-down list box. Otherwise, you can manually enter the WWPN. Adding ports from the drop-down list is less prone to error and is the best method. However, if hosts have not yet been connected to the SAN or zoned, then manually adding the WWPNs is the only option.




Figure 4-45 GUI example: Add FC port WWPN

Repeat steps 5 and 6 to add the second HBA WWPN; ports can be added in any order.

7. To add an iSCSI host, in the **Add Port** dialog, specify the port type as iSCSI and enter the IQN of the HBA as the iSCSI Name. See Figure 4-46.




Figure 4-46 GUI example: Add iSCSI port

8. The host will appear with its ports in the Hosts dialog box as shown in Figure 4-47.

[-] itso_win2008	default
10000000C97D295C	FC
10000000C97D295D	FC
[-] itso_win2008_iscsi	default
iqn.1991-05.com.microsoft:sand.storage.tucson.ibm.com	iSCSI

Figure 4-47 List of hosts and ports

In this example, the hosts *itso_win2008* and *itso_win2008_iscsi* are in fact the same physical host, however, they have been entered as separate entities so that when mapping LUNs, the FC and iSCSI protocols do not access the same LUNs.

Mapping LUNs to a host

The final configuration step is to map LUNs to the host. To do this, follow these steps:

1. While still in the **Hosts and Clusters** configuration pane, right-click the host to which the volume is to be mapped and select **Modify LUN Mappings** from the context menu (see Figure 4-48).




Figure 4-48 Map LUN to host

The **LUN Mapping for Host** window opens as shown in Figure 4-49.

2. Select an available volume from the left pane.

The GUI will suggest a LUN ID to which to map the volume, however, this can be changed to meet your requirements.

3. Click **Map** and the volume is assigned immediately.




Figure 4-49 Map FC volume to FC host

There is no difference in mapping a volume to an FC or iSCSI host in the XIV GUI **LUN Mapping for Host** view.

4. To complete this example, power up the host server and check connectivity. The IBM XIV Storage System has a real-time connectivity status overview.
5. Select **Hosts Connectivity** from the **Hosts and Clusters** menu to access the connectivity status. See Figure 4-50.



Figure 4-50 Hosts Connectivity

The host connectivity window is displayed. In our example, the *ExampleFCHost* was expected to have dual path connectivity to every module. However, only two modules (5 and 6) show as connected (see Figure 4-51), and the iSCSI host has no connection to module 9.

Name	1:4	1:5	1:6	1:7	1:8	1:9
Bladecenter4_etc_1						
Bladecenter4_etc_2						
Bladecenter4_etc_3						
Bladecenter4_etc_4						
Bladecenter4v9						
ExampleFHost						
210000E08BFA29E						
210100E08BFA29E						
ExampleISCSiHost						
iqn.2000-04.com.qlogic...						

Figure 4-51 GUI example: Host connectivity matrix

The setup of the new FC and/or iSCSI hosts on the IBM XIV Storage System is complete. At this stage there might be operating system dependent steps that need to be performed; these are described in the operating system chapters.

4.5.2 Assigning LUNs to a host using the XCLI

There are a number of steps required in order to define a new host and assign LUNs to it. Prerequisites are that volumes have been created in a storage pool.

Defining a new host

Follow these steps to use the XCLI to prepare for a new host:

1. Create a host definition for your FC and iSCSI hosts, using the **host_define** command. See Example 4-9.

Example 4-9 XCLI session example: Create host definition

```
>>host_define host=itso_win2008
Command executed successfully.
```

```
>>host_define host=itso_win2008_iscsi
Command executed successfully.
```

Host access to LUNs is granted depending on the host adapter ID. For an FC connection, the host adapter ID is the FC HBA WWPN. For an iSCSI connection, the host adapter ID is the IQN of the host.

2. Add the WWPN of the FC host for HBA1 and HBA2 with the **host_add_port** command by specifying an **fcaddress** as shown in Example 4-10.

Example 4-10 Create FC port and add to host definition

```
>> host_add_port host=itso_win2008 fcaddress=10000000c97d295c
Command executed successfully.
>> host_add_port host=itso_win2008 fcaddress=10000000c97d295d
Command executed successfully.
```

3. Add the IQN of the iSCSI host as shown in Example 4-11. Notice that this is the same **host_add_port** command, but with the **iscsi_name** parameter.

Example 4-11 Create iSCSI port and add to the host definition

```
>> host_add_port host=itso_win2008_iscsi
iscsi_name=iqn.1991-05.com.microsoft:sand.storage.tucson.ibm.com
Command executed successfully
```

Mapping LUNs to a host

The final configuration step is to map LUNs to the host definition. To do so, follow these steps:

1. Map LUNs to the host definition using the commands shown in Example 4-12. Note that for a cluster, the volumes are mapped to the cluster host definition. There is no difference for FC or iSCSI mapping to a host.

Example 4-12 XCLI example: Map volumes to hosts

```
>> map_vol host=itso_win2008 vol=itso_win2008_vol1 lun=1
Command executed successfully.

>> map_vol host=itso_win2008 vol=itso_win2008_vol2 lun=2
Command executed successfully.

>> map_vol host=itso_win2008_iscsi vol=itso_win2008_vol3 lun=1
Command executed successfully.
```

2. To complete the example, power up the server and check the host connectivity status from the IBM XIV Storage System point of view. Example 4-13 shows the output for both hosts.

Example 4-13 XCLI example: Check host connectivity

```
>> host_connectivity_list host=itso_win2008
Host          Host Port          Module          Local FC port  Type
itso_win2008  10000000C97D295C  1:Module:6      1:FC_Port:6:1  FC
itso_win2008  10000000C97D295C  1:Module:4      1:FC_Port:4:1  FC
itso_win2008  10000000C97D295D  1:Module:5      1:FC_Port:5:3  FC
itso_win2008  10000000C97D295D  1:Module:7      1:FC_Port:7:3  FC

>> host_connectivity_list host=itso_win2008_iscsi
Host          Host Port          Module          Local FC port  Type
itso_win2008_iscsi  iqn.1991-05.com.microsoft:sand.storage.tucson.ibm.com  1:Module:8      iSCSI
itso_win2008_iscsi  iqn.1991-05.com.microsoft:sand.storage.tucson.ibm.com  1:Module:7      iSCSI
```

In Example 4-13, there are two paths per host FC HBA and two paths for the single Ethernet port that was configured.

The setup of the new FC and/or iSCSI hosts on the IBM XIV Storage System is now complete.

At this stage there might be operating system dependent steps that need to be performed; these steps are described in the operating system chapters.

See the *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904 for details related to host definitions and volume mapping.

4.6 Scripts

IBM XIV Storage Manager software XCLI commands can be used in scripts or batch programs in case you need to use repetitive or complex operations. The XCLI can be used in a shell environment to interactively configure the system or as part of a script to perform specific tasks; see Figure 4-3 on page 101. In general, the XIV GUI or the XCLI Session environment will virtually eliminate the need for scripts.



Security

In this chapter we discuss the XIV Storage System security features from various perspectives.

We cover the following topics:

- ▶ System physical access security
- ▶ Native user authentication
- ▶ LDAP managed user authentication
- ▶ Securing LDAP communication with Security Socket Layer
- ▶ Audit event logging

5.1 Physical access security

When installing an XIV Storage System, you need to apply the same security best practices that you apply to any other business critical IT system. A good reference on storage security can be found at the Storage Networking Industry Association (SNIA) website:

http://www.snia.org/forums/ssif/programs/best_practices

A common risk with storage systems is the retention of volatile caches. The XIV Storage System is perfectly safe in regard to external operations and a loss of external power. In the case of a power failure, the internal Uninterruptible Power Supply (UPS) units provide power to the system. The UPS enables the XIV Storage System to gracefully shut down.

However, if someone were to gain physical access to the equipment, that person might manually shut off components by bypassing the preferred process. In this case, the storage system is likely to lose the contents of its volatile caches, resulting in a data loss and system unavailability. To eliminate or greatly reduce this risk, the XIV rack is equipped with lockable doors. You can prevent unauthorized people from accessing the rack by simply locking the doors, which will also protect against unintentional as well as malicious changes *inside* the system rack.

Important: Protect your XIV Storage System by locking the rack doors and monitoring physical access to the equipment.

5.2 Native user authentication

To prevent unauthorized access to the configuration of the storage system and ultimately to the information stored on its volumes, the XIV Storage System uses password based user authentication. Password based authentication is a form of challenge-response authentication protocol where the authenticity of a user is established by presenting that user with a question “challenge” and comparing the answer “response” with information stored in a credential repository.

5.2.1 XIV local credential repository

By default, the XIV Storage System is configured to use native (XIV managed) user authentication. Native user authentication makes use of the credential repository stored locally on the IBM XIV Storage System. The XIV local credential repository maintains the following types of objects:

- ▶ User name
- ▶ User password
- ▶ User role
- ▶ User group
- ▶ Optional account attributes

User name

A user name is a string of 1-64 characters that can only contain ‘a-z’, ‘A-Z’, ‘0-9’, ‘.-_~’ and ‘space’ symbols. User names are case sensitive. The XIV Storage System is configured with a set of predefined user names. Predefined user names and corresponding default passwords exist to provide initial access to XIV at the time of installation, for system maintenance, and for integration with application such as the Tivoli Storage Productivity Center.

The following user accounts are predefined on the IBM XIV Storage System:

- *technician*: This account is used by the IBM support representative to install the XIV Storage system
- *admin*: This account provides the highest level of customer access to the system. It can be used for creating new users and change passwords for existing users in native authentication mode.

Important: Use of the *admin* account must be limited to the initial configuration when no other user accounts are available. Access to the *admin* account must be restricted and securely protected.

- *xiv_development* and *xiv_maintenance* user: These IDs are special case pre-defined internal IDs that can only be accessed by qualified IBM development and service support representatives (SSRs).

Predefined user accounts cannot be deleted from the system and are always authenticated natively by the XIV Storage System even if the system operates under LDAP authentication mode.

New user accounts can initially be created by the *admin* user only. After the *admin* user creates a new user account and assigns it to the *storageadmin* (Storage Administrator) role, then other user accounts can be created by this newly created *storageadmin* user.

In native authentication mode, the system is limited to creating up to 32 user accounts. This number includes the predefined users.

User password

The *user password* is a secret word or phrase used by the account owner to gain access to the system. The user password is used at that time of authentication to establish the identity of that user. User passwords can be 6 to 12 characters long, using the characters [a-z][A-Z]~!@#\$%^&*()_+~{}|!;<>?,./[], and must not include any space between characters. In native authentication mode, the XIV Storage System verifies the validity of a password at the time the password is assigned.

Predefined users have the following default passwords assigned at the time of XIV Storage System installation (see Table 5-1).

Table 5-1 Default passwords

Predefined user	Default password
admin	adminadmin
technician	predefined and is used only by the IBM XIV Storage System technicians
xiv_development	predefined and is used only by the IBM XIV development team
xiv_maintenance	predefined and is used only by the IBM XIV maintenance team

Important: As a best practice, the default admin password must be changed at installation time to prevent unauthorized access to the system.

The following restrictions apply when working with passwords in native authentication mode:

- ▶ For security purposes, passwords are not shown in user lists.
- ▶ Passwords are user changeable. Users can change only their own passwords.
- ▶ Only predefined user *admin* can change the passwords of other users.
- ▶ Passwords are changeable from both the XCLI and the GUI.
- ▶ Passwords are case-sensitive.
- ▶ User password assignment is mandatory at the time new user account is created.
- ▶ Creating user accounts with empty password or removing password from an existing user account is not permitted.

User roles

There are four predefined user roles (in the XIV GUI and the XCLI. Roles are referred to as *categories* and are used for day to day operation of the XIV Storage System. The following section describes predefined roles, their level of access, and applicable use:

- ▶ *storageadmin*:

The *storageadmin* (Storage Administrator) role is the user role with highest level of access available on the system. A user assigned to this role has an ability to perform changes on any system resource except for maintenance of physical components or changing the status of physical components.

- ▶ *applicationadmin*:

The *applicationadmin* (Application Administrator) role is designed to provide flexible access control over volume snapshots. User assigned to the *applicationadmin* role can create snapshots of specifically assigned volumes, perform mapping of their own snapshots to a specifically assigned host, and delete their own snapshots. The user group to which an application administrator belongs determines the set of volumes that the application administrator is allowed to manage. If a user group is defined with *access_all*="yes", application administrators who are members of that group can manage all volumes on the system. For more details on user group membership and group to host association, see "User groups" on page 135.

- ▶ *readonly*:

As the name implies, users assigned to the *readonly* role can only view system information. Typical use for the *readonly* role is a user who is responsible for monitoring system status, system reporting, and message logging, and who must not be permitted to make any changes on the system.

- ▶ *technician*:

The *technician* role has a single predefined user name (*technician*) assigned to it, and is intended to be used by IBM support personnel for maintaining the physical components of the system. The technician is limited to the following tasks: physical system maintenance and phasing components in or out of service. The technician has extremely restricted access to the system and is unable to perform any configuration changes to pools, volumes, or host definitions on the XIV Storage System.

- ▶ *xiv_development*:

The *xiv_development* role has a single predefined user name (*xiv_development*) assigned to it and is intended to be used by IBM development personnel.

- ▶ *xiv_maintenance*:

The *xiv_maintenance* role has a single predefined user name (*xiv_maintenance*) assigned to it and is intended to be used by IBM maintenance personnel.

User roles: There is no capability to add new user roles or to modify predefined roles. In native authentication mode, after a user is assigned a role, the only way to assign a new role is to first delete the user account and then recreate it.

Table 5-2 Predefined user role assignment

Predefined user	User role
admin	storageadmin
technician	technician
xiv_development	xiv_development
xiv_maintenance	xiv_maintenance

Native authentication mode implements user role mechanism as a form of *Role Based Access Control* (RBAC). Each predefined user role determines the level of system access and associated functions that a user is allowed to use.

RBAC: The XIV Storage System implements Role Based Access Control (RBAC) based authentication and authorization mechanisms.

All user accounts must be assigned to a single user role. Assignment to multiple roles is not permitted. Deleting or modifying role assignment of natively authenticated users is also not permitted.

User groups

A user group is a group of application administrators who share the same set of snapshot creation permissions. The permissions are enforced by associating the user groups with hosts or clusters. User groups have the following characteristics:

- ▶ Only users assigned to the *applicationadmin* role can be members of a user group.
- ▶ A user can be a member of a single user group.
- ▶ A maximum of eight user groups can be created.
- ▶ Group names are case-sensitive
- ▶ In native authentication mode, a user group can contain up to eight members.
- ▶ If a user group is defined with *access_all="yes"*, users assigned to the *applicationadmin* role who are members of that group can manage all snapshots on the system.
- ▶ A user must be assigned to the *storageadmin* role to be permitted to create and manage user groups.

Important: A user group membership can only be defined for users assigned to the *applicationadmin* role.

User group and host associations

Hosts and clusters can be associated with only a single user group. When a user is a member of a user group that is associated with a host, that user is allowed to manage snapshots of the volumes mapped to that host.

User group and host associations have the following properties:

- ▶ User groups can be associated with both hosts and clusters. This enables limiting group member access to specific volumes.
- ▶ A host that is part of a cluster can only be associated with a user group through user group to cluster association. Any attempts to create user group association for that host will fail.
- ▶ When a host is added to a cluster, the association of that host is removed. Limitations on the management of volumes mapped to the host is controlled by the association of the cluster.
- ▶ When a host is removed from a cluster, the association of that cluster remains unchanged. This enables continuity of operations so that all scripts relying on this association will continue to work.

Optional account attributes

In this topic we discuss optional attributes for email and phone numbers:

- ▶ *Email*: Email is used to notify specific users about events through email messages. Email addresses must follow standard formatting procedures.

Acceptable value: Any valid email address. Default value is not defined.

- ▶ *Phone and area code*: Phone numbers are used to send SMS messages to notify specific users about system events. Phone numbers and area codes can be a maximum of 63 digits, hyphens (-) and periods (.)

Acceptable value: Any valid telephone number. Default value is not defined.

5.2.2 Managing user accounts with XIV GUI

This section illustrates the use of the XIV GUI in native authentication mode for creating and managing user accounts, as well for creating user groups and defining group membership.

Adding users with the GUI

The following steps require that you initially log on to the XIV Storage System with storage administrator access rights (*storageadmin* role). If this is the first time that you are accessing the system, use the predefined user *admin* (default password *adminadmin*):

1. Open the XIV GUI and log on as shown in Figure 5-1.




Figure 5-1 GUI Login

2. Users are defined per system. If you manage multiple systems and they have been added to the GUI, select the particular system with which you want to work.

3. In the main Storage Manager GUI window, move the mouse pointer over the padlock icon to display the Access menu. All user access operations can be performed from the Access menu (see Figure 5-2). There are three choices:
 - *Users*: Define or change single users.
 - *Users Groups*: Define or change user groups, and assign application administrator users to groups.
 - *Access Control*: Define or change user groups, and assign application administrator users or user group to hosts.
4. Move the mouse over the **Users** menu item (it is now highlighted in yellow) and click it, as shown in Figure 5-2.




Figure 5-2 GUI Access menu

5. The Users window is displayed.

If the storage system is being accessed for the first time, the window displays the predefined users only. See Figure 5-3 for an example. The default columns are Name, Category, Group, Phone, and Email.

Name	Category	Group	Phone	Email
admin	storageadmin			
technician	technician			
xiv_development	xiv_development			
xiv_maintenance	xiv_maintenance			

Figure 5-3 GUI Users management

An additional column called Full Access can be displayed (this only applies to users assigned to the *applicationadmin* role). To add the Full Access column, right-click the blue heading bar to display the Customize Columns dialog shown in Figure 5-4.




Figure 5-4 Customize columns dialog

- a. Change the default password for the admin user, which can be accomplished by right-clicking the user name and selecting **Change Password** from the context menu, as illustrated in Figure 5-5. You must be logged in as *admin* to change the admin password.




Figure 5-5 GUI admin user change password

6. To add a new user, you can either click the Add icon in the menu bar or right-click the empty space to get the context menu. Both options are visible in Figure 5-6. Click **Add User**. Remember that there is a maximum of 32 users in the regular XIV access control.




Figure 5-6 GUI Add User option

7. The Define User dialog is displayed. A *user* is defined by a unique name and a password (see Figure 5-7). The default role (denoted as Category in the dialog panel) is Storage Administrator. Category must be assigned. Optionally, enter the email address and phone number for the user. Click **Add** to create the user and return to the Users window.

Add User

Name: * lab_admin

New Password: (6-12) *

Retype New Password: *

Category: Application Administrator

User Group: None

Email Address: lab_admin@domain.de

Phone Number: 0049 0815

Add **Cancel**

Figure 5-7 GUI Define User attributes

8. If you need to test the user that you just defined, click the current user name shown in the upper right corner of the IBM XIV Storage Manager window (Figure 5-8), which allows you to log in as a new user.




Figure 5-8 GUI quick user change

Defining user groups with the GUI

The IBM XIV Storage system can simplify various user management tasks with the capability to create user groups.

Groups: User groups only apply to users assigned to the *applicationadmin* role.

A user group can also be associated with one or multiple hosts or clusters.

The following steps illustrate how to create user groups, add users (with application administrator role) to the group, and how to define host associations for the group:

1. Be sure to log in as *admin* (or another user with storage administrator rights). From the Access menu, click **Users Groups** as shown in Figure 5-9. In our scenario, we create a user group called EXCHANGE CLUSTER 01. As shown in Figure 5-9, the user groups can be accessed from the *Access* menu (padlock icon).





Figure 5-9 Select Users Groups

2. From the Users Groups window that displays, to add a new user group, either click the **Add User Group** icon (shown in Figure 5-10) in the menu bar, or right-click in an empty area of the User Groups table and select **Add User Group** from the context menu as shown in Figure 5-10.



Figure 5-10 Add User Group

- From the Create User Group dialog that displays, enter a meaningful group name and click **Add** (see Figure 5-11).



The 'Add User Group' dialog box contains three input fields: 'Name:' with a red asterisk and the text 'application02_group', 'LDAP Role:' which is empty, and 'Full Access:' with an unchecked checkbox. At the bottom are 'Add' and 'Cancel' buttons.

Figure 5-11 Enter New User Group Name

LDAP role: The “LDAP role” field is not applicable to user group definition in native authentication mode and will have no effect even if a value is entered.

If a user group has the *Full Access* flag turned on, all members of that group will have unrestricted access to all snapshots on the system.

At this stage, the user group EXCHANGE CLUSTER 01 still has no members and no associations defined. Next, we create an association between a host and the user group.

- Select **Access Control** from the Access menu as shown in Figure 5-12. The Access Control window displays.




Figure 5-12 Access Control

- Right-click the name of the user group that you have created to bring up a context menu and select **Updating Access Control** as shown in Figure 5-13.




Figure 5-13 Update Access Control for a user group

The *User Group Access Control* dialog shown in Figure 5-14 is displayed. The panel contains the names of all the hosts and clusters defined to the XIV Storage System. The left pane displays the list of Unauthorized Hosts/Clusters for this particular user group and the right pane shows the list of hosts that have already been associated to the user group.

6. Add or remove hosts from either list by selecting a host and clicking the appropriate arrow. Finally, click **Update** to save the changes.




Figure 5-14 Access Control Definitions panel

After a host (or multiple hosts) have been associated with a user group, you can define user membership for the user group (remember that a user must have the application administrator role to be added to a user group).

7. Go to the *Users* window and right-click the user name to display the context menu. From the context menu (see Figure 5-15), select **Add to Group** to add this user to a group.




Figure 5-15 Add a user to a group

- From the *Select User Group* dialog that displays, select the desired group from the pull-down list and click **OK** (see Figure 5-16).




Figure 5-16 *Select User Group*

The user *lab_admin* has been added as a member to the user group *application01_group* in this example.

- Verify this group membership in the Users panel as shown in Figure 5-17.

Name	Category	Group
admin	storageadmin	
lab_admin	applicationadmin	application01_group
technician	technician	
xiv_development	xiv_development	
xiv_maintenance	xiv_maintenance	

Figure 5-17 *View user group membership*

The user *lab_admin* is an *applicationadmin* with the Full Access right set to *no*. This user can now perform snapshots of the “itso-v480R-tic1” volumes. Because “itso-v480R-tic1” is the only cluster associated with the group, *adm_pmeier01* is only allowed to map those snapshots to “itso-v480R-tic1”. However, you can add another host or a cluster, such as a test or backup host, to allow *lab_admin* to map a snapshot volume to a test or backup server.

5.2.3 Managing user accounts using XCLI

This section summarizes the commands and options available to manage user accounts, user roles, user groups, group memberships, and user group to host associations through the XCLI session.

Table 5-3 shows the various commands and a brief description for each command. The table also indicates the user role required to issue specific commands.

Table 5-3 XCLI access control commands

Command	Description	Role required to use command
access_define	Defines an association between a user group and a host.	storageadmin
access_delete	Deletes an access control definition.	storageadmin
access_list	Lists access control definitions.	storageadmin, readonly, and applicationadmin
user_define	Defines a new user.	storageadmin
user_delete	Deletes a user.	storageadmin
user_group_add_user	Adds a user to a user group.	storageadmin
user_group_create	Creates a user group.	storageadmin
user_group_delete	Deletes a user group.	storageadmin
user_group_list	Lists all user groups or a specific one.	storageadmin, readonly, and applicationadmin
user_group_remove_user	Removes a user from a user group.	storageadmin
user_group_rename	Renames a user group.	storageadmin
user_group_update	Updates a user group.	storageadmin
user_list	Lists all users or a specific user.	storageadmin, readonly, and applicationadmin
user_rename	Renames a user.	storageadmin
user_update	Updates a user. You can rename the user, change password, modify the Access All setting, modify email, area code, and/or phone number.	storageadmin, and applicationadmin

Adding users with the XCLI

Before performing the following steps, the XCLI component must be installed on the management workstation, and a *storageadmin* user is required. The following examples assume a Windows-based management workstation.

To add users, follow these steps:

1. Use the **user_list** command to obtain the list of predefined users and categories as shown in Example 5-1. This example assumes that no users, other than the default users, have been added to the system.

Example 5-1 XCLI user_list

```
>>user_list
```

Name	Category	Group	Active...	Access All
xiv_development	xiv_development	yes		
xiv_maintenance	xiv_maintenance	yes		
admin	storageadmin	yes		
technician	technician	yes		
lab_admin	applicationadmin	application01_group	yes	... no

2. If this is a new system, change the default password for the admin user, by using the **update_user** command as shown in Example 5-2.

Example 5-2 XCLI user_update

```
>>user_update user=admin password=New_Passw0rd password_verify=New_Passw0rd
Command executed successfully.
```

3. Add a new user as shown in Example 5-3. This is straightforward; you define a user by a unique name, password, and role (designated here as *category*).

Example 5-3 XCLI user_define

```
>>user_define user=lab_user2 password=Go4XIV password_verify=Go4XIV
category=storageadmin
Command executed successfully.
```

Defining user groups with the XCLI

To define a user group, proceed as in the following examples:

1. Use the **user_group_create** command as shown in Example 5-4 to create a user group called EXCHANGE_CLUSTER_01.

Example 5-4 XCLI user_group_create

```
>>user_group_create user_group=application_Mainz
Command executed successfully.
```

Spaces: Avoid spaces in user group names. If spaces are required, the group name must be placed between single quotation marks, such as "name with spaces."

The user group *application_Mainz* is empty and has no associated host. The next step is to associate a host or cluster.

2. Associate the user group *application_Mainz* to the host *application_Mainz_05*, as shown in Example 5-5.

Example 5-5 XCLI access_define

```
>>access_define user_group="application_Mainz" host="application_Mainz_05"
Command executed successfully.
```

A host has been assigned to the user group. The user group still does not have any users included.

3. Add the first user as shown in Example 5-6.

Example 5-6 XCLI user_group_add_user

```
>>user_group_add_user user_group=application_Mainz user=lab_user3  
Command executed  
successfully.
```

The user `lab_user3` has been assigned to the user group `application_Mainz`. This user is an *applicationadmin* with the Full Access right set to *no*.

5.2.4 Password management and resets

Password management in native authentication mode is internal to the XIV Storage System. The XIV system has no built-in password management rules such as password expiration, preventing reuse of the same passwords and/or password strength verification. Furthermore, if you want to log on to multiple systems at any given time through the GUI, you must be registered with the same password on all the IBM XIV Storage Systems.

In native authentication mode, as long as users can log in, they can change their own passwords.

The predefined user *admin* is the only user is authorized to change other users passwords. Direct access to user credential repository is not permitted. System security is enforced by allowing password changes only through XIV GUI and XCLI.

Figure 5-18 shows how to change a password by right-clicking the selected user in the Users window. Then, you select **Change Password** from the context menu.




Figure 5-18 GUI change password context menu

The Change Password dialog shown in Figure 5-19 is displayed. Enter the New Password and then retype it for verification in the appropriate field (remember that only alphanumeric characters are allowed). Click **Update**.




Figure 5-19 GUI Change Password window

5.2.5 Managing multiple systems

Managing multiple XIV Storage Systems is straightforward in native authentication mode. Due to the fact that user credentials are stored locally on every XIV system, it is key to keep the same user name and password on separate XIV Storage Systems to allow for quick transitions between systems in the XIV GUI. This approach is especially useful in Remote Mirror configurations, where the storage administrator is required to switch from source to target system.

Figure 5-20 illustrates the GUI view of multiple systems when using non-synchronized passwords. For this example, the system named “XIV 02 7802969” has a user account *lab_admin* that provides the storage admin level of access. Because the userid is not configured for the all XIV Storage Systems, only the two systems are currently shown as accessible.

The user can see the other systems, but is unable to access them with the *lab_admin* user (the unauthorized systems appear greyed out). They also state that the user is unknown. If the system has the *lab_admin* defined with another password, the systems are displayed in the same state.



Figure 5-20 Single User Login

To allow simultaneous access to multiple systems, the simplest approach is to have corresponding passwords manually synchronized among those systems. Figure 5-21 illustrates the use of user account with passwords synchronized among four XIV systems. The storage administrator can easily switch between these systems for the activities without having to log on each time with another password. The XIV system where the user was successfully authenticated is now displayed in color with an indication of its status.




Figure 5-21 Manual user password synchronization among multiple XIV systems

5.3 LDAP managed user authentication

The XIV Storage System offers the capability to use LDAP server based user authentication.

When LDAP authentication is enabled, the IBM XIV Storage System accesses a specified LDAP directory to authenticate users whose credentials are maintained in the LDAP directory (with the exception of the admin, technician, maintenance, and development users, which remain locally administered and maintained).

The benefits of an LDAP based centralized user management can be substantial when considering the size and complexity of the overall IT environment. Maintaining local user credentials repositories is relatively straightforward and convenient when only dealing with a small number of users and a small number storage systems. However, as the number of users and interconnected systems grows, the complexity of user account management rapidly increases and managing such an environment becomes a time consuming task.

In this section, we review various benefits of this approach. Although the benefits from utilizing LDAP are significant, you must also evaluate the considerable planning effort and complexity of deploying LDAP infrastructure, if it is not already in place.

5.3.1 Introduction to LDAP

The Lightweight Directory Access Protocol (LDAP) is an open industry standard that defines a standard method for accessing and updating information in a directory. It is being supported by a growing number of software vendors and is being incorporated into a growing number of products and applications.

A directory is a listing of information about objects arranged in an order that gives details about each object. Common examples are a city telephone directory and a library card catalog. In computer terms, a directory is a specialized database, also called a data repository, that stores typed and ordered information about objects. A particular directory might list information about users (the objects) consisting of typed information such as user names, passwords, email addresses and so on. Directories allow users or applications to find resources that have the characteristics needed for a particular task.

Directories in LDAP are accessed using the client/server model. An application that wants to read or write information in a directory does not access the directory directly, but uses a set of programs or APIs that cause a message to be sent from LDAP client to LDAP server. LDAP server retrieves the information requested on behalf of the client application and returns the requested information if the client has permission to see the information. LDAP defines a message protocol used between the LDAP clients and the LDAP directory servers. This includes methods to search for information, read information, and to update information based on permissions.

LDAP-enabled directories have become a popular choice for storing and managing user access information. LDAP provides a centralized data repository of user access information that can be securely accessed over the network. It allows the system administrators to manage the users from multiple XIV Storage Systems in one central directory.

5.3.2 LDAP directory components

An LDAP directory is a collection of objects organized in a tree structure. The LDAP naming model defines how objects are identified and organized. Objects are organized in a tree-like structure called the Directory Information Tree (DIT). Objects are arranged within the DIT based on their distinguished name (DN). Distinguished name (DN) defines location of an object within DIT. Each object is also referred to as an entry in a directory belonging to an object classes. An object class describes the content and purpose of the object. It also contains a list of attributes, such as a telephone number or surname, that can be defined in an object of that object class.

As shown in Figure 5-22, the object with the distinguished name (DN) `cn=mbarlen, ou=Marketing, o=IBM` belongs to object class `objectClass=ePerson`.

Object class `ePerson` contains attributes:

`cn` (common name),
`mail`,
`sn` (surname),
`givenName`
`telephoneNumber`.

Each attribute has a value assigned to it:

`cn=mbarlen`,
`mail=marion@ibm.com`,
`sn=Barlen`,
`givenName=Marion`,
`telephoneNumber=112`.

In this example, the object represents a single employee record. If a record for a new employee in organizational unit (ou) Marketing of organization (o) IBM needs to be created, the same location in DIT will be the same, `ou=Marketing, o=IBM`, and the same set of attributes defined by objectClass `ePerson` will also be used. The new object will be defined using its own set of attribute values because the new employee will have their own name, email address, phone number, and so on.

For more information about the directory components, see the IBM Redbooks publication, *Understanding LDAP - Design and Implementation*, SG24-4986.




Figure 5-22 Example of a Directory Information Tree (DIT)

All the objects and attributes with their characteristics are defined in a schema. The schema specifies what can be stored in the directory.

5.3.3 LDAP product selection

LDAP authentication for Version 10.2.2 of the XIV Storage System supports two LDAP server products:

- ▶ Microsoft Active Directory
- ▶ Sun Java Services Directory Server Enterprise Edition

The current skill set of your IT staff is always an important consideration when choosing any products for centralized user authentication. If you have skills in running a particular directory server, then it might be a wise choice to standardize on this server because your skilled people will best be able to customize and tune the server as well as to provide the most reliable and highly-available implementation for the LDAP infrastructure.

Microsoft Active Directory might be a better choice for an enterprise with most of its infrastructure components deployed using Microsoft Windows operating system. Sun Java Services Directory Server Enterprise Edition on the other hand provides support for UNIX®-like operating systems including Linux, as well as Microsoft Windows.

All LDAP servers share many basic characteristics because they are based on the industry standard Request for Comments (RFCs). However, because of implementation differences, they are not always entirely compatible with each other. For more information about RFCs, particularly regarding LDAP RFC 4510-4533, see the following Web page:

<http://www.ietf.org/rfc.html>

Current implementation of LDAP based user authentication for XIV does not support connectivity to multiple LDAP server of various types. However, it is possible to configure an XIV Storage System to use multiple LDAP servers of the same type to eliminate a single point of failure. The IBM XIV system will support communication with a single LDAP server at a time. The LDAP authentication configuration allows specification of multiple LDAP servers that the IBM XIV Storage System can connect to if a given LDAP server is inaccessible.

5.3.4 LDAP login process overview

The XIV login process when LDAP authentication is enabled is depicted in Figure 5-23:

1. The XIV user login process starts with the user launching a new XCLI or GUI session and submitting credentials (user name and password) to the IBM XIV Storage System (step “1”).
2. In step “2” the IBM XIV Storage System logs into a previously defined LDAP server using the credentials provided by the user. If login to the LDAP server fails, a corresponding error message is returned to the user and the login process terminates.
3. If XIV successfully logs into the LDAP server, it retrieves attributes for establishing LDAP role mapping (step “3”). If the IBM XIV Storage System cannot establish LDAP role mapping, the user login process terminates and a corresponding error message is returned to the user.
4. If LDAP role mapping is successfully established, XIV creates a new session and returns a prompt to the user (step “4”). For more information about LDAP role mapping, see 5.3.5, “LDAP role mapping”




Figure 5-23 LDAP authentication process overview

5.3.5 LDAP role mapping

Before any LDAP user can be granted access to XIV, the user must be assigned to one of the supported user roles. XIV uses the *storageadmin*, *readonly*, and *applicationadmin* predefined roles. The mechanism used for determining what role a particular user is assigned to is called *role mapping*. In native mode, a role is explicitly assigned to a user at the time of *user account* creation. In LDAP mode, the role of a specific user is determined at the time the user logs in to the IBM XIV Storage System.

Planning considerations

When initially planning to use LDAP based authentication with XIV, the LDAP server administrator has to make a decision as to which LDAP attribute can be used for role mapping. As discussed in 5.3.2, “LDAP directory components” on page 148 each LDAP object has a number of associated attributes. The type of LDAP object classes used to create user account for XIV authentication depends on the type of LDAP server being used. The SUN Java Directory server uses the *inetOrgPerson* LDAP object class, and Active Directory uses the *organizationalPerson* LDAP object class for definition of user accounts for XIV authentication.

For a detailed definition of the *inetOrgPerson* LDAP object class and list of attributes, see the Internet FAQ Archive website:

<http://www.faqs.org/rfcs/rfc2798.html>

For a detailed definition of the *organizationalPerson* LDAP object class and list of attributes, see the Microsoft website:

[http://msdn.microsoft.com/en-us/library/ms683883\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/ms683883(VS.85).aspx)

The designated attribute (as established by the LDAP administrator) will be used for storing a value that will map a user to a role. In our examples we used the *memberOf* attribute for the purpose of role mapping. Ultimately the decision on what attribute is to be used for role mapping must be left to the LDAP administrator.

If the *memberOf* attribute is already used for something else, then the LDAP administrator has to designate another attribute. For the purpose of illustration, we assume that the LDAP administrator agrees to use the *memberOf* attribute.

The XIV administrator now can make the necessary configuration change to instruct the IBM XIV Storage System to use *memberOf* as the attribute name. This is done by assigning the *memberOf* value to the *xiv_group_attr* configuration parameter, using the **ldap_config_set** XCLI command:

```
ldap_config_set xiv_group_attr=memberOf
```

Next, the XIV administrator defines two names that will be used as role names in LDAP. In our example, the XIV administrator uses the “xiv_admin” and “xiv_ro” names for mapping to *storageadmin* and *readonly* roles. The XIV administrator sets corresponding parameters in the IBM XIV Storage System:

```
ldap_config_set storage_admin_role="xiv_admin"  
ldap_config_set read_only_role="xiv_ro"
```

Case sensitivity: The LDAP server uses case-insensitive string matching for the *memberOf* attribute value. For example, “xiv_admin”, Xiv_Admin”, or “XIV_ADMIN” will be recognized as equal strings. To simplify XIV system administration, however, treat both the XIV configuration parameter and LDAP attribute value as if they were case-sensitive and assign the “xiv_admin” value to both.

The “xiv_admin” and “xiv_ro” names were used simply because both strings can be easily associated with their corresponding XIV roles: *storageadmin* and *readonly* respectively. It is not necessary to use the same names in your XIV system configuration.

However, if you were to change these parameters, consider using names that are self descriptive and easy to remember, in order to simplify the LDAP server administration tasks. Every time the LDAP server administrator creates a new XIV account, one of the names must be entered as a description attribute value (except for the *applicationadmin* role, as we explain later in “LDAP role mapping for applicationadmin” on page 154). After being configured in both XIV and LDAP, changing these parameters, although possible, can potentially be time consuming due to the fact that each existing LDAP account must be changed individually, to reflect the new attribute value.

These configuration tasks can also be done from the XIV GUI. On the main XIV Storage Management panel, select **Configure System**, and select the **LDAP** panel on the left panel. Enter *memberOf* in the XIV Group Attribute field, xiv_ro in the Read Only Role field, and xiv_admin in the Storage Admin Role field as shown in Figure 5-24. Finally, click **Update** to save the changes.




Figure 5-24 LDAP definitions in XIV GUI

With XIV Software V10.2.2, there is also a GUI wizard available to help you with configuring LDAP and role mappings. See 5.4, “Defining LDAP on the XIV system” on page 155.

The XIV administrator informs the LDAP administrator that the *xiv_admin* and *xiv_ro* groups have to be used for role mapping:

- The LDAP administrator creates a user account in LDAP and assigns the “user to the *memberOf* attribute. When the newly created user logs in to the system, XIV performs the role mapping as depicted in Figure 5-25.




Figure 5-25 Assigning LDAP authenticated user to storageadmin role

- The LDAP administrator creates an account and assigns *xiv_ro* to the *memberOf* attribute. The newly created user is assigned to the *readonly* role. When this user logs into the IBM XIV Storage System, it performs the role mapping as depicted in Figure 5-26.




Figure 5-26 Assigning LDAP authenticated user to readonly role

LDAP role mapping for applicationadmin

The LDAP account can also be assigned to an *applicationadmin* role, but the mechanism of creating role mapping in this case differs from the one used for *storageadmin* and *readonly* role mapping.

The XIV system will assign a user to the *applicationadmin* role if it can match the value of the *MemberOf* attribute with the *ldap_role* parameter of any user groups defined in XIV. If an account is assigned the *applicationadmin* role, it also becomes a membership of the user group whose *ldap_role* parameter matches the value of the user's *MemberOf* attribute.

The user group must be created before the user logs in to the system, otherwise the login will fail. The XIV administrator creates a user group, with the **user_group_create** XCLI command, as follows:

```
user_group_create user_group=app01_group ldap_role=app01_administrator
```

After the LDAP administrator has created the user account and assigned the *app01_administrator* value to the *MemberOf* attribute, the user can be authenticated by the IBM XIV Storage System. The role assignment and group membership inheritance for a newly created user is depicted in Figure 5-27.

If the IBM XIV Storage System cannot find a match for the value assigned to the *MemberOf* attribute of a user, then the user is denied system access.




Figure 5-27 Assigning LDAP authenticated user to applicationadmin role

Table 5-4 summarizes the LDAP to XIV role mappings.

Table 5-4 LDAP role mapping summary

XIV role	XIV configuration parameter name	XIV configuration parameter setting	LDAP attribute value
storageadmin	storage_admin_role	xiv_admin	xiv_admin
readonly	read_only_role	xiv_ro	xiv_ro
applicationadmin	N/A	N/A	arbitrary (must match <i>ldap_role</i> field in XIV user group)

5.4 Defining LDAP on the XIV system

This section describes how to configure XIV in an existing LDAP environment.

Important: By upgrading to the XIV software V10.2.2, any prior LDAP configuration on XIV will be reset. Contact IBM if you need to migrate an existing LDAP implementation.

Be aware that all users (but not the predefined users) that you had eventually created for native authentication will be deactivated (not deleted though) when you activate the LDAP mode in XIV. Defined users will not show anymore in the XIV user settings. However they are not deleted and will be shown and activated again if you eventually disable the LDAP authentication later on.

There are three ways to configure the LDAP on the XIV, the GUI LDAP Wizard, through the regular XIVGUI panels and the XCLI. You can see examples for all in the following section.

In any case, the first step for the XIV administrator is to verify Domain Name Server (DNS) name resolution as illustrated in Example 5-7.

Example 5-7 DNS name resolution verification

```
>> dns_test name=xiv.domain
Name                                     IP (Primary DNS)  IP (Secondary DNS)
xiv.domain                             9.155.66.59
```

If the **dns_test** command returns an unexpected result, do not proceed further with the configuration steps until the DNS name resolution issue is resolved.

The XIV Storage System software Version 10.2.2 brings various changes to the LDAP implementation:

- ▶ It supports special character and white spaces in the login name.
- ▶ There is a new LDAP Configuration Wizard in the XIV GUI.
- ▶ A Directory Service User, for use by XIV, must be configured. This user is used to check the LDAP connection and functionality.

Important: As best practice, the LDAP server and XIV system must have their clocks synchronized to the same time source, be registered, and be configured to use the same Domain Name Server servers.

5.4.1 Using the XIV GUI LDAP Wizard to configure LDAP

Follow these steps to configure LDAP by using the GUI LDAP Wizard.

1. Go to the Access icon and click **Users**, as seen in Figure 5-28.



Figure 5-28 Users

2. Click the **LDAP Wizard** icon on the toolbar as shown in Figure 5-29.




Figure 5-29 LDAP Wizard

The LDAP Wizard opens the LDAP Configuration Welcome panel.

3. As shown in Figure 5-30, click **Next** to start the configuration.



Figure 5-30 Wizard - LDAP Configuration

In the Server Type tab, as shown in Figure 5-31, you can choose the Server type from the corresponding pull-down. You can use either a Microsoft Active Directory or a Sun Directory server.

4. Click **Next** to choose **Microsoft Active Directory** as we did here.




Figure 5-31 Server Type

5. From the LDAP Server tab, click the “+” icon located on the right side to add an LDAP server to the configuration, as illustrated in Figure 5-32.





Figure 5-32 Adding LDAP Servers

6. In the new dialog window that displays (Figure 5-33), enter the FQDN (Fully Qualified Domain Name), which is the DNS name of the LDAP server (such as xiv.domain in our example).

Also enter the LDAP server IP address and the Search DN, which in our Active Directory implementation is CN=Users,DC=xiv,DC=domain, as illustrated in Figure 5-33.

For example, in a default Active Directory implementation, a server with the Domain name ldap.domain.de will map to CN=Users,DC=ldap,DC=domain,DC=de.

7. If you need a certification file for the connection to your LDAP server, add it here for the connection.
8. Click **Define** to save the settings.




The 'Define LDAP Server' dialog box contains the following fields and controls:

- FQDN:** A text field with the value 'xiv.domain' and a red asterisk indicating it is required.
- Server Address:** A text field with the value '9.155.66.59' and a red asterisk indicating it is required.
- Search DN:** A text field with the value 'CN=Users,DC=xiv,DC=domain' and a red asterisk indicating it is required.
- Certificate File:** A text field with a 'Browse...' button next to it.
- Buttons:** 'Define' and 'Cancel' buttons at the bottom.

Figure 5-33 Define LDAP Server

Now you can see your added server reflected in the table, as shown in Figure 5-34.

9. Add more LDAP servers if necessary, or click **Next** to go on with the configuration process.



The 'Wizard - LDAP Configuration - Directory Services Server Configuration' window shows the following components:

- Left Navigation Pane:** A list of steps: Welcome, Server Type, **LDAP Servers** (highlighted), XIV User, Unique ID's, Groups, Secure LDAP, Bind & Cache, Test & Activation, and Finish.
- Main Content Area:**
 - Title:** Directory Services Server Configuration
 - Description:** Add, Remove or modify your Directory Services servers. If you plan on using Secure LDAP, please make sure you attach the appropriate LDAP Server certificate ("pem" format) for each server you configure.
 - Table:** A table with two columns: FQDN and Address. It contains one entry: 'xiv.domain' with address '9.155.66.59'. There are '+' and '-' icons to the right of the table for adding and removing rows.
- Bottom Buttons:** 'Back', 'Next', 'Finish', and 'Cancel' buttons.

Figure 5-34 LDAP Servers

10. From the XIV User tab, as shown in Figure 5-35, define one user (this must be an already defined LDAP user) that will be used by XIV to validate the LDAP settings. That user does not need any special permissions.
11. As seen in Figure 5-35, enter the complete user DN for LDAP, which in our example is CN=XIV,CN=Users,DC=xiv,DC=domain. Enter the password for this user as well. The password will not be saved; it will just be used for this wizard process. Click **Next**.



Figure 5-35 XIV User

12. In the Unique ID's tab now displayed, set the parameters for the login (see Figure 5-36). In our example the ID Attribute was set to ObjectSid, the Name Attribute to displayName and the Group Attribute to memberOf. In SUN Directory environments, for the Name Attribute, use uid; for Group Attribute in XIV, use isMemberOf.
13. Ask your LDAP admin for the value of those parameters in your environment. Click **Next** to save the settings and proceed.




Figure 5-36 Unique ID's

In the Groups tab shown in Figure 5-37, you have to define the group attributes that XIV will consider for the permissions. These groups must have been first defined in LDAP. Our example has an LDAP group named xiv_admin and another one named xiv_ro.

14. Define the role with the complete DN path. In our example, the Admin role is CN=xiv_admin,CN=Users,DC=xiv,DC=domain and the Read Only role is CN=xiv_ro,CN=Users,DC=xiv,DC=domain, as shown in Figure 5-37. Click **Next** until you reach the **Test and Activation** tab.

Figure 5-37 Groups

15. From the Test and Activation tab shown in Figure 5-38, enable LDAP by setting “Use LDAP” to “YES”. There is also a possibility to test your LDAP with a user already defined in LDAP.

Figure 5-38 Test & Activation

16. At this stage, click **Finish** to complete the setup or click **Next** to go to the exit and summary panel shown in Figure 5-39.

LDAP mode considerations:

- ▶ After you activate the LDAP mode, it will deactivate all previously defined XIV users. Only the LDAP user with the right group role will have access to the system. You will also be logged out from the XIV GUI.
- ▶ All users and passwords will be reactivated again, after you have disabled the LDAP mode and return to the native mode. This means that the users are just deactivated and not deleted.




Figure 5-39 Wizard summary

17. Check the summary and click **Finish**. You will be asked if you are sure that you want to enable LDAP now, because it will log you off the system afterwards. See Figure 5-40.




Figure 5-40 Activation confirmation

Click **OK**.

18. When you get a confirmation that LDAP was configured, as shown in Figure 5-41, click **OK** again to dismiss the dialog.




Figure 5-41 End of LDAP configuration

5.4.2 Using the XIV GUI directly to configure LDAP

Configuring LDAP in XIV can also be done directly through XIV GUI dialogs, rather than using the LDAP wizard. Follow these steps:

1. Go to the **Access** icon and click **Users**, as shown in Figure 5-28.



Figure 5-42 Users

2. Click the **Configure LDAP** icon located in the toolbar, as shown in Figure 5-43.



Figure 5-43 Configure LDAP

3. In the LDAP Configuration menu shown in Figure 5-44, set the value of “Use LDAP” to “Yes” and choose the LDAP Server Type that you want to use. In our example, we selected **Microsoft Active Directory**. Click **Servers** to go to the next step.




Figure 5-44 LDAP Configuration - General

4. In the Servers page in the LDAP configuration, click the green plus (+) icon, as shown in Figure 5-45.





Figure 5-45 LDAP Configuration - Servers

5. In the new dialog window that displays, enter the FQDN (Fully Qualified Domain Name), which is the DNS name of the LDAP server (such as xiv.domain in our example).
Also enter the LDAP server IP address and the Search DN, which in our Active Directory implementation is CN=Users,DC=xiv,DC=domain, as illustrated in Figure 5-46.
In a default Active Directory implementation, for example, a server with the Domain Name ldap.domain.de will map to CN=Users,DC=ldap,DC=domain,DC=de.
6. Click **Define** to save these settings.



The 'Define LDAP Server' dialog box contains the following fields and values:


Field	Value
FQDN:	xiv.domain
Server Address:	9.155.66.59
Search DN:	CN=Users,DC=xiv,DC=domain
Certificate File:	[Empty] Browse...

Buttons: Define, Cancel

Figure 5-46 Define LDAP Servers

- Click **User Credentials** to define the Service User DN, which is just used to verify your LDAP settings.

In Figure 5-47 you can see that a service user named XIV is defined using a complete Service User DN, which in this example is “CN=XIV,CN=Users,DC=xiv,DC=domain”. Enter and confirm the user password.



The 'LDAP Configuration' dialog box has a sidebar with the following tabs: General, Servers, User Credentials (selected), Role Mapping, Secure LDAP, and Parameters. The 'User Credentials' section contains the following fields and values:

Field	Value
Service User DN	CN=XIV,CN=Users,DC=xiv,DC=domain
Password	
Confirm Password	

Buttons: Update, Cancel

Figure 5-47 LDAP Configuration - User Credentials

- Click **Role Mapping** to enter the LDAP attributes you want to use for XIV login. See Figure 5-48 for an illustration of the values for Microsoft Active Directory we used in our example.

User ID Attribute : ObjectSid
 User Name Attribute: displayName
 XIV Group Attribute: memberOf
 Storage Admin Role: CN=xiv_admin,CN=Users,DC=xiv,DC=domain
 Read Only role: CN=xiv_ro,CN=Users,DC=xiv,DC=domain

Figure 5-48 LDAP Configuration - Role Mapping

LDAP mode: When you activate the LDAP mode, it will deactivate all defined users on the XIV and only LDAP users with the right group role will have access to the system. You will also be logged out from the GUI.

- Click **Update** and the configuration is now finished.

5.4.3 Using XCLI

The quickest method for configuring LDAP in XIV is to use the XCLI.

First you have to define your LDAP Server in the XIV, as you can see in Example 5-8.

Example 5-8 Xcli add server

```
XIV_1300203>>ldap_add_server fqdn=xiv.domain address=9.155.66.59
base_dn=CN=Users,DC=xiv,DC=domain
Command executed successfully.
```

```
XIV LAB 3 1300203>>ldap_list_servers
FQDN      Address Base DN ...
xiv.domain 9.155.66.59  CN=Users,DC=xiv,DC=domain ...
```

After adding one or more LDAP servers, you can now define the login attributes in the XIV. In Example 5-9 you can see an example for a possible Microsoft Active Directory implementation.

Example 5-9 Xcli LDAP configuration

```
XIV_1300203>>ldap_config_set server_type="Microsoft Active Directory"
xiv_group_attr=memberOf
storage_admin_role=CN=xiv_admin,CN=Users,DC=xiv,DC=domain
read_only_role=CN=xiv_ro,CN=Users,DC=xiv,DC=domain
user_id_attr=ObjectSid
xiv_user=CN=XIV,CN=Users,DC=xiv,DC=domain user_name_attr=displayName
```

Command executed successfully.

For a detailed description of the LDAP role mapping, see 5.3.5, “LDAP role mapping” on page 151. The XIV configuration parameters `storage_admin_role`, `read_only_role` and `xiv_group_attr` must have values assigned for LDAP role mapping to work.

After the configuration was submitted successfully, you can check your values with the command `ldap_config_get`.

At this stage, user accounts must be created in the LDAP Directory. See Appendix A, “Additional LDAP information” on page 279.

Now that all configuration and verification steps are completed, the LDAP mode can be activated, using the `ldap_mode_set` command as illustrated in Example 5-10.

LDAP mode: When you activate the LDAP mode, it will deactivate all defined users on the XIV system, and only LDAP users with the right group role will have access to the system.

Example 5-10 Enabling the XIV LDAP mode

```
XIV_1300203>>ldap_mode_set mode=Active
```

```
Warning: ARE_YOU_SURE_YOU_WANT_TO_ENABLE_LDAP_AUTHENTICATION y/n: y
Command executed successfully.
```

The LDAP authentication mode is now fully configured, activated, and ready to be tested. A simple test that can validate the authentication result is to open an XCLI session using credentials of a newly created Active Directory account `xivtestuser1` and run `ldap_user_test`. This command can only be successfully executed by a user authenticated through LDAP.

5.5 LDAP managed user authentication

As previously explained, when the IBM XIV Storage System is configured for LDAP authentication, user credentials are stored in the centralized LDAP repository. The LDAP repository resides on an LDAP server and is accessed by the IBM XIV Storage System, using an LDAP protocol.

5.5.1 LDAP repository credential objects

The LDAP repository maintains the following types of credential objects:

- ▶ LDAP user name
- ▶ LDAP user password
- ▶ LDAP user role
- ▶ User groups

LDAP user name

The XIV system limitations for acceptable user names, such as number of characters and character set, no longer apply when user names are stored in an LDAP repository. Each LDAP product has its own set of rules and limitations that applies to user names. Generally, do not use very long names and non-alpha-numeric characters even if your LDAP product of choice supports it. If you eventually decide to migrate user credentials between local and LDAP repositories or vice-versa, the task can be greatly simplified if the same set of rules is applied to both local and centralized repositories. In fact, the set of rules enforced by the IBM XIV Storage System for local user names must be used for LDAP as well, because it is the most restrictive of the two. For details about XIV system limitations for user names, see “User name” on page 132.

Special consideration must be given to using the “space” character in user names. Although this feature is supported with LDAP, it has a potential for making certain administrative tasks more difficult because the user names in this case will have to be enclosed in quotes to be interpreted correctly.

The same set of locally stored predefined user names exist on the IBM XIV Storage System regardless of the authentication mode. Users *technician* and *admin* are always authenticated locally even on a system with activated LDAP authentication mode. Avoid creating LDAP *user accounts* with the same names.

If a user account with the same user name is registered in both local and LDAP repositories, and LDAP authentication mode is in effect, then LDAP authentication takes precedence, and the IBM XIV Storage System will perform authentication using LDAP account credentials. The only exception to this rule are the predefined user names listed in the previous paragraph. To reduce complexity and simplify maintenance, it is generally not desirable to have the same user names registered in local and LDAP repositories.

If a user account was registered in the local repository on the IBM XIV Storage System before the LDAP authentication mode was activated, then this account will not be accessible while LDAP authentication is in effect. The account will become accessible again upon deactivation of the LDAP authentication mode.

LDAP user passwords

User passwords are stored in the LDAP repository when the IBM XIV Storage System is in LDAP authentication mode. Password management becomes a function of the LDAP server. The XIV system relies entirely on the LDAP server to provide functionality such as enforcing initial password resets, password strength, password expiration, and so on. Various LDAP server products provide their own sets of tools and policies for password management.

The following examples in Figure 5-49 and Example 5-11 provide an illustration of techniques that can be used for password management, and by no means represent a complete list of product password management capabilities.




Figure 5-49 Default Active Directory Password Policy settings

Example 5-11 Default SUN Java Directory Password Policy settings

```
# /opt/sun/ds6/bin/dsconf get-server-prop | grep ^pwd
Enter "cn=Directory Manager" password:
pwd-accept-hashed-pwd-enabled      : N/A
pwd-check-enabled                  : off
pwd-compat-mode                    : DS5-compatible-mode
pwd-expire-no-warning-enabled      : on
pwd-expire-warning-delay           : 1d
pwd-failure-count-interval         : 10m
pwd-grace-login-limit              : disabled
pwd-keep-last-auth-time-enabled    : off
pwd-lockout-duration               : 1h
pwd-lockout-enabled                : on
pwd-lockout-repl-priority-enabled  : on
pwd-max-age                        : disabled
pwd-max-failure-count              : 3
pwd-max-history-count              : disabled
pwd-min-age                        : disabled
pwd-min-length                     : 6
pwd-mod-gen-length                 : 6
pwd-must-change-enabled            : off
pwd-root-dn-bypass-enabled         : off
pwd-safe-modify-enabled            : off
pwd-storage-scheme                 : SSHA
pwd-strong-check-dictionary-path    : /opt/sun/ds6/plugins/words-english-big.txt
pwd-strong-check-enabled           : off
pwd-strong-check-require-charset   : lower
pwd-strong-check-require-charset   : upper
pwd-strong-check-require-charset   : digit
pwd-strong-check-require-charset   : special
pwd-supported-storage-scheme       : CRYPT
pwd-supported-storage-scheme       : SHA
pwd-supported-storage-scheme       : SSHA
```

```
pwd-supported-storage-scheme      : NS-MTA-MD5
pwd-supported-storage-scheme      : CLEAR
pwd-user-change-enabled           : on
```

In the event of a user's password expiration or account lockout, the user will get the message shown in Example 5-12, while attempting to login to XCLI.

Example 5-12 XCLI authentication error due to account lockout

```
>> ldap_user_test
Error:  USER_NOT_AUTHENTICATED_BY_LDAP_SERVER
Details: User xivtestuser2 was not authenticated by LDAP server
'xivhost2.storage.tucson.ibm.com'.
```

The XIV GUI in this situation will also fail with the error message seen in Figure 5-50.




Figure 5-50 XIV GUI authentication failure due to account lockout

Although password policy implementation will greatly enhance overall security of the system, all advantages and disadvantages of such implementation need to be carefully considered. One of the possible disadvantages is increased management overhead for account management as a result of implementing complex password management policies.

Passwords: A comprehensive solution for user password policy implementation is beyond the scope of this book.

LDAP user roles

There are predefined user roles (also referred to as *categories*) used for day to day operation of the XIV Storage System. In the following section we describe predefined roles, their level of access, and applicable use:

► *storageadmin:*

The *storageadmin* (Storage Administrator) role is the user role with the highest level of access available on the system. A user assigned to this role has an ability to perform changes on any system resource except for maintenance of physical components or changing the status of physical components. The assignment of the *storageadmin* role to an LDAP user is done through LDAP role mapping process. For a detailed description, see 5.3.5, “LDAP role mapping” on page 151.

► *applicationadmin:*

The *applicationadmin* (Application Administrator) role is designed to provide flexible access control over volume snapshots. Users assigned to the *applicationadmin* role can create snapshots of specifically assigned volumes, perform mapping of their own snapshots to a specifically assigned host, and delete their own snapshots. The user group to which an application administrator belongs determines the set of volumes that the application administrator is allowed to manage. If a user group is defined with *access_all*=“yes”, application administrators who are members of that group can manage

all volumes on the system. The assignment of the *applicationadmin* role to an LDAP user is done through the *LDAP role mapping process*. For a detailed description, see 5.3.5, “*LDAP role mapping*” on page 151. Detailed description of user group to host association is provided in “User group membership for LDAP users”.

► *readonly*:

As the name implies, users assigned to *readonly* role can only view system information. Typical use for *readonly* role is a user responsible for monitoring system status, system reporting and message logging and must not be permitted to make any changes on the system. The assignment of *readonly* role to an *LDAP user* is done through *LDAP role mapping process*. For a detailed description, see 5.3.5, “*LDAP role mapping*” on page 151.

Roles: There is no capability to add new user roles or to modify predefined roles. In LDAP authentication mode, role assignment can be changed by modifying the LDAP attribute (*memberOf* in our example).

LDAP authentication mode implements user role mechanism as a form of *Role Based Access Control* (RBAC). Each predefined user role determines the level of system access and associated functions a user is allowed to use.

RBAC: The XIV Storage System implements Role Based Access Control (RBAC) based authentication and authorization mechanisms.

All user accounts must be assigned to a single user role. Any LDAP user assigned to multiple roles will not be authenticated by the IBM XIV Storage System. Deleting role assignment (by removing the description attribute value in the LDAP object) of LDAP users will also lead to XIV’s inability to authenticate that user.

User group membership for LDAP users

A user group is a group of application administrators who share the same set of snapshot management permissions. The permissions are enforced by associating the user groups with hosts or clusters. User groups are defined locally on the IBM XIV Storage System.

User group membership for an LDAP user is established during the login process by matching the designated LDAP attribute value with the *ldap_role* parameter assigned to a user group. A user group is associated with host volumes through access definition. An LDAP user, member of the user group, is permitted to manage snapshots of volumes mapped to the host associated with the user group.

User groups have the following characteristics in LDAP authentication mode:

- Only users assigned to the *applicationadmin* role can be members of a user group.
- An LDAP user can only be a member of a single user group.
- A maximum of eight user groups can be created.
- In LDAP authentication mode, there is no limit on the number of members in a user group.
- If a user group is defined with *access_all*= “yes”, users assigned to the *applicationadmin* role who are members of that group can manage all snapshots on the system.
- The user group parameter *ldap_role* can only be assigned a single value.
- The *ldap_role* parameter must be unique across all defined user groups.
- Only users assigned to the *storageadmin* role can create, modify and delete user groups.

- Only users assigned to the *storageadmin* role can modify *ldap_role* parameter of a user group.

Important: A user group membership can only be defined for users assigned to the *applicationadmin* role.

Figure 5-51 illustrates the relationship between LDAP user, LDAP role, XIV role, user group membership, associated host, mapped volumes, and attached snapshots.




Figure 5-51 User group membership for LDAP user

Example 5-13 is a Korn shell script that can be used to demonstrate the relationship between LDAP user, LDAP role, XIV role, user group membership, associated host, mapped volumes and attached snapshots.

Example 5-13 query_snapshots.ksh listing user's role, group membership, volumes and snapshots

```
#!/bin/ksh

# XIV customer-configurable LDAP server parameters:

LDAPHOSTNAME=xivhost2.storage.tucson.ibm.com
XIV_GROUP_ATTRIB=MemberOf
READ_ONLY_ROLE="xiv_ro"
STORAGE_ADMIN_ROLE="xiv_admin"
BASE_DN="dc=xivauth"

echo "Enter username: "
read USERNAME

echo "Enter password: "

stty -echo
read USERPASSWORD
stty echo

LDAP_ROLE=$(/opt/sun/dsee6/bin/ldapsearch -x -b $BASE_DN -D uid=$USERNAME, \
$BASE_DN -w $USERPASSWORD -h $LDAPHOSTNAME uid=$USERNAME $XIV_GROUP_ATTRIB \
| grep ^$XIV_GROUP_ATTRIB | awk -F': ' '{ print $2 }')

if [[ $? -ne 0 ]]
then
    echo Failed to query LDAP account details
    exit -1
fi

xcli -c ARCXIVJEMT1 -u $USERNAME -p $USERPASSWORD ldap_user_test
if [[ $? -ne 0 ]]
then
    exit -1
fi

echo "-----"
if [[ $LDAP_ROLE == $READ_ONLY_ROLE ]]
then
    echo XIV Role: \"readonly\"
    exit 0
elif [[ $LDAP_ROLE == $STORAGE_ADMIN_ROLE ]]
then
    echo XIV Role: \"storageadmin\"
    exit 0
else
    echo User: $USERNAME
    echo LDAP role: $LDAP_ROLE
    echo XIV Role: \"applicationadmin\"
    USER_GROUP_NAME=`xcli -c ARCXIVJEMT1 -u $USERNAME -p $USERPASSWORD
user_group_list | grep -w $LDAP_ROLE | awk '{ print $1 }'`
    echo Member of user group: \"$USER_GROUP_NAME\"

```

```

        for HOSTNAME in `xcli -c ARCXIVJEMT1 -u $USERNAME -p $USERPASSWORD
access_list user_group=$USER_GROUP_NAME | grep -v "Type   Name           User Group"
| awk '{ print $2 }'`
        do
            echo Host: $HOSTNAME associated with $USER_GROUP_NAME user group
        done

        for HOSTNAME in `xcli -c ARCXIVJEMT1 -u $USERNAME -p $USERPASSWORD
access_list user_group=$USER_GROUP_NAME | grep -v "Type   Name" | awk '{ print $2
}'`
        do
            for VOLUME in `xcli -c ARCXIVJEMT1 -u $USERNAME -p $USERPASSWORD
mapping_list host=app01_host | grep -v LUN | awk '{ print $2 }'`
            do
                SNAPSHOT=`xcli -c ARCXIVJEMT1 -u $USERNAME -p $USERPASSWORD
snapshot_list vol=$VOLUME | grep -v ^Name | awk '{ print $1 }'`
                echo VOLUME: $VOLUME '->' SNAPSHOT: $SNAPSHOT
            done
        done
fi

```

The sample output of the *query_snapshots.ksh* script shown in Example 5-14 provides an illustration of the configuration described in Figure 5-51 on page 171.

Example 5-14 Sample run output of query_snapshots.ksh script

```

$ query_snapshots.ksh
Enter username:
app01_administrator
Enter password:
-----
User: xivtestuser3
LDAP_ROLE: app01_administrator
XIV Role: "applicationadmin"
Member of user group: "app01_group"
Host: app01_host associated with app01_group user group
VOLUME: app01_vol01 -> SNAPSHOT: app01_snap01
VOLUME: app01_vol02 -> SNAPSHOT: app01_snap02

```

5.5.2 Managing LDAP user accounts

Managing user accounts in LDAP authentication mode is done using LDAP management tools. The XCLI commands and XIV GUI tools cannot be used for creating, deleting, modifying, or listing LDAP user accounts. The set of tools for LDAP account management is specific to the LDAP server type. Examples of LDAP account creation are provided in Appendix A in the topics, “Creating user accounts in Microsoft Active Directory” on page 280 and “Creating user accounts in SUN Java Directory” on page 286. The same set of LDAP management tools can also be used for account removal, modification, and listing.

Generating LDAP account lists

To generate a list of all LDAP user accounts registered under the Base_DN (XIV system configuration parameter specifying the location of LDAP accounts in the directory information tree DIT), you can use the **ldapsearch** queries shown in Example 5-15 and Example 5-16.

Example 5-15 Generating list of LDAP accounts registered in SUN Java Directory

```
# ldapsearch -x -b dc=xivauth -H ldap://xivhost2.storage.tucson.ibm.com:389 -D
cn="Directory Manager" -w passwd uid | grep ^uid:
```

```
uid: xivsunproduser3
uid: xivsunproduser4
uid: xivsunproduser5
```

Example 5-16 Generating list of LDAP accounts registered in Active Directory

```
# ldapsearch -x -H "ldap://xivhost1.xivhost1ldap.storage.tucson.ibm.com:389" -D
'CN=Administrator,CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com' -w
Passw0rd -b 'CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com'
cn=xivtestuser1 | grep ^cn:
```

```
cn: xivadproduser10
cn: xivadproduser11
cn: xivadproduser12
```

The queries generating LDAP account lists are provided as a demonstration of LDAP tools capabilities to perform a search of information stored in LDAP directory and generate simple reports. Note that both queries are issued on behalf of the LDAP administrator account, cn="Directory Manager" and cn="Administrator" for SUN Java Directory and Active Directory respectively. A privileged account such as LDAP administrator has the authority level allowing it not only listing but also creating, modifying, and removing other user accounts.

The Active Directory management interface also allows you to build custom views based on LDAP search queries. The following example builds a query that generates the list of XIV accounts whose names start with xiv, and whose description is one of the following three: "Storage Administrator", "Read Only", or starts with "app":




Figure 5-52 Active Directory query listing XIV accounts

For generating the "XIV Storage Accounts" view, we used the LDAP query shown in Example 5-17.

Example 5-17 LDAP query for generating list of XIV user accounts

```
(&(&(objectCategory=user)(cn=xiv*)(|(description=Read Only)(description=Storage Administrator)(description=app*))))
```

To create this query, select **Saved Queries** → **New** → **Query** → XIV Storage Accounts (query name in this example) → **Define Query** → Select **Custom Search** in the **Find** scroll down list → **Advanced** and paste the LDAP query from Example 5-17 into the *Enter LDAP Query* field.

When a new user account is created and its name and attributes satisfy the search criterion, this user account will automatically appear in the *XIV Storage Accounts* view. A similar technique can be applied for managing user accounts in the SUN Java Directory. Any LDAP GUI front-end supporting LDAP version 3 protocol can be used for creating views and managing LDAP entries (XIV user accounts). An example of such an LDAP front end for the SUN Java Directory is the *Directory Editor* product. This product is part of the SUN Java Directory product suit. Demonstrating Directory Editor capabilities for managing LDAP objects is outside of the scope for this book.

Table 5-5 provides a list of commands that *cannot* be used for user account management when LDAP authentication mode is active.

Table 5-5 XIV commands unavailable in LDAP authentication mode

XIV command
user_define
user_update
user_rename
user_group_add_user
user_group_remove_user

Authentication: When the IBM XIV Storage System operates in LDAP authentication mode, user account creation, listing, modification, and removal functionality is provided by the LDAP server.

Note that the **user_list** command can still operate when LDAP authentication mode is active. However, this command will only show locally defined XIV user accounts and not LDAP accounts. See Example 5-18.

Example 5-18 user_list command output in LDAP authentication mode

```
>> user_list show_users=all
Name                Category                Group    Active
xiv_development     xiv_development         Group    yes
xiv_maintenance     xiv_maintenance         Group    yes
admin               storageadmin             Group    yes
technician          technician                Group    yes

>> ldap_mode_get
Mode
```

Active
>>

As shown in Example 5-18, the Active parameter is set to “no” for user ITSO. The parameter specifies whether a user can login in current authentication mode. All predefined local XIV users can still login when LDAP authentication mode is active.

Defining user groups with the GUI in LDAP authentication mode

User group information is stored locally on the IBM XIV Storage System regardless of the authentication mode. The user group concept only applies to users assigned to an application_administrator role.

A user group can also be associated with one or multiple hosts or clusters.

The following steps illustrate how to create user groups, how to add users (with application administrator role) to the group, and how to define host associations for the group:

1. Be sure to log in as *admin* (or another user with storage administrator rights). From the Access menu, click **Users Groups** as shown in Figure 5-53. In our scenario, we create a user group called app01_group. The user groups can be selected from the Access menu (padlock icon).



Figure 5-53 Select Users Groups

2. In the Users Groups window, to add a new user group, either click the **Add User Group** icon (shown in Figure 5-54) in the menu bar, or right-click in an empty area of the User Groups table and select **Add User Group** from the context menu.




Figure 5-54 Add User Group

3. In the Create User Group dialog, enter a meaningful group name, specify role for LDAP role mapping described in 5.3.5, “LDAP role mapping” on page 151 and click **Add** (Figure 5-55). To avoid potential conflicts with already registered user groups, the IBM XIV Storage System verifies the uniqueness of the group name and the role. If a user group with the same name or the same role exists in the XIV repository, the attempt to create a new user group will fail and an error message is displayed.

The *Full Access* flag has the same significance as in native authentication mode. If a user group has the Full Access flag turned on, all members of that group will have unrestricted access to all snapshots on the system.




Figure 5-55 Enter New User Group Name and Role for LDAP role mapping

At this stage, the user group app01_group is still empty.

- Next, add a host to the user group by selecting **Access Control** from the Access menu as shown in Figure 5-56. The Access Control window displays.




Figure 5-56 Access Control

- Right-click the name of the user group that you have created to bring up a context menu and select **Update Access Control** as shown in Figure 5-57.




Figure 5-57 Updating Access Control for a user group

The User Group Access Control dialog shown in Figure 5-58 is displayed. The panel contains the names of all the hosts and clusters defined on the XIV Storage System.

The left pane displays the list of Unauthorized Hosts/Clusters for this particular user group and the right pane shows the list of hosts that have already been associated with the user group.

- Add or remove hosts from either list by selecting a host and clicking the appropriate arrow.
- Finally, click **Update** to save the changes (Figure 5-58).




Figure 5-58 Access Control Definitions panel

Unlike in native authentication mode, in LDAP authentication mode, user group membership cannot be defined using the XIV GUI or XCLI. The group membership is determined at the time the LDAP authenticated user logs into the IBM XIV Storage System, based on the information stored in LDAP directory. A detailed description of the process of determining user group membership can be found in 5.3.5, “LDAP role mapping” on page 151.

Starting with XIV Storage System Software 10.2.2, it is now possible to delete user groups when LDAP authentication is enabled.

5.5.3 Managing user groups using XCLI in LDAP authentication mode

This section describes how to manage user groups using the XCLI.

Commands and options

A variety of commands and options are available to manage user groups, roles, and associated host resources through the XCLI.

Defining user groups with the XCLI

To use the GUI to define user groups:

1. Use the **user_group_create** command as shown in Example 5-19 to create a user group called `app01_group` with corresponding LDAP role `app01_administrator`.

Example 5-19 XCLI user_group_create in LDAP authentication mode

```
>> user_group_create user_group=app01_group ldap_role=app01_administrator
Command completed successfully
```

Spaces: Avoid spaces in user group names. If spaces are required, the group name must be placed between single quotation marks, such as “name with spaces.”

The user group app01_group is empty and has no associated hosts or clusters. The next step is to associate a host or cluster with the group.

2. Associate the user group app01_group to app01_host, as shown in Example 5-20,

Example 5-20 XCLI access_define

```
>> access_define user_group=app01_group host=app01_host  
Command completed successfully
```

5.5.4 Active Directory group membership and XIV role mapping

In Active Directory, a user can be a member of a single group or multiple groups. An LDAP group is a collection of users with common characteristics. *Group* is defined in the Active Directory container *Users*. A group is defined first as an empty container, and then existing users can be assigned as members of this group. A group is represented as a separate object in the LDAP Directory Information Tree (DIT) and gets a distinguished name (DN) assigned to it.

Groups defined in the Active Directory can be used for XIV role mapping. When a user becomes a member of a group in the Active Directory, it gets a new attribute assigned. The value of the new attribute points to the DN of the group. *MemberOf* is the name of that attribute. The *MemberOf* attribute value determines the Active Directory group membership.

To create a group in Active Directory, follow these steps:

1. Start Active Directory Users and Computer by selecting **Start** → **Administrative Tools** → **Active Directory Users and Computers**
2. Right-click the “Users” container, and select **New** → **Group**.
3. Enter a group name and click **OK**.

The new Active Directory group creation dialog is shown in Figure 5-59.




Figure 5-59 Creating Active Directory group

To assign an existing user to the new group, follow these steps:

1. Start Active Directory Users and Computer by selecting **Start** → **Administrative Tools** → **Active Directory Users and Computers**.
2. Expand the *Users* container, right-click the user name that you want to make a member of the new group, and select **Properties**.
3. In the Properties panel, select the **Member Of** tab and click **Add** → **Advanced** → **Find Now**. From the presented list of existing user groups, select **XIVReadOnly** and click **OK**.
4. You can now see a group selection dialog as shown in Figure 5-60. Confirm your choice by clicking **OK**.




Figure 5-60 Active Directory group selection dialog

To illustrate the new *memberOf* attribute in the existing LDAP user object and the new LDAP object representing the “XIVReadOnly” group, we run the **ldapsearch** queries against the Active directory LDAP server as shown in Example 5-21.

Example 5-21 Active Directory group membership ldapsearch queries

```
# ldapsearch -x -H "ldap://xivstorage.org:389" -b CN=Users,dc=xivstorage,dc=org
-D cn=xivadproduser10,CN=Users,dc=xivstorage,dc=org -w pass2remember
cn=xivadproduser10 memberOf
dn: CN=xivadproduser10,CN=Users,DC=xivstorage,DC=org
memberOf: CN=XIVReadOnly,CN=Users,DC=xivstorage,DC=org
```

```
# ldapsearch -x -H "ldap://xivstorage.org:389" -b CN=Users,dc=xivstorage,dc=org
-D cn=xivadproduser10,CN=Users,dc=xivstorage,dc=org -w pass2remember
cn=XIVReadOnly
```

```
dn: CN=XIVReadOnly,CN=Users,DC=xivstorage,DC=org
objectClass: top
objectClass: group
cn: XIVReadOnly
member: CN=xivadproduser10,CN=Users,DC=xivstorage,DC=org
distinguishedName: CN=XIVReadOnly,CN=Users,DC=xivstorage,DC=org
instanceType: 4
whenCreated: 20090712021348.0Z
whenChanged: 20090712021451.0Z
uSNCreated: 61451
uSNChanged: 61458
name: XIVReadOnly
objectGUID:: Ai3j9w7at02cEjV1lpk/fQ==
```

```
objectSid:: AQUAAAAAAAAUAAAA1a5i5i2p5CXmfc4aQAAA==
sAMAccountName: ReadOnly
sAMAccountType: 268435456
groupType: -2147483646
objectCategory: CN=Group,CN=Schema,CN=Configuration,DC=xivstorage,DC=org
```

In the first **ldapsearch** query, we intentionally limited our search to the *memberOf* attribute (at the end of the **ldapsearch** command) so that the output is not obscured with unrelated attributes and values. The value of the *memberOf* attribute contains the DN of the group.

The second **ldapsearch** query illustrates the CN=XIVReadOnly LDAP object content. Among other attributes, it contains the *member* attribute that points at the DN of the user defined as a member. The attribute *member* is a multivalued attribute; there can be more than one user assigned to the group as a member. *MemberOf* is also a multivalued attribute, and a user can be a member of multiple groups.

The XIV system can now be configured to use the *memberOf* attribute for role mapping. In Example 5-22 we are mapping the Active Directory group XIVReadOnly to the XIV `read_only_role`, XIVStorageadmin to `storage_admin_role`, and XIV user group `app01_group` to Active Directory group `XIVapp01_group`. You must be logged on as *admin*.

Example 5-22 Configuring XIV to use Active Directory groups for role mapping

```
>> ldap_config_set xiv_group_attr=memberOf
Command executed successfully.
>> ldap_config_set read_only_role=CN=XIVReadOnly,CN=Users,DC=xivstorage,DC=org
Command executed successfully.
>> ldap_config_set
storage_admin_role=CN=XIVStorageadmin,CN=Users,DC=xivstorage,DC=org
Command executed successfully.
>> ldap_config_get
Name                                Value
base_dn                            CN=Users,dc=xivstorage,dc=org
xiv_group_attr                      memberOf
third_expiration_event              7
version                             3
user_id_attr                       objectSid
current_server                     xivstorage.org
use_ssl                             no
session_cache_period               10
second_expiration_event            14
read_only_role                     CN=XIVReadOnly,CN=Users,DC=xivstorage,DC=org
storage_admin_role                  CN=XIVStorageadmin,CN=Users,DC=xivstorage,DC=org
first_expiration_event              30
bind_time_limit                     30

>> user_group_update user_group=app01_group
ldap_role=cn=XIVapp01_group,CN=Users,DC=xivstorage,DC=org
Command executed successfully.
>> user_group_list user_group=app01_group
Name      Access All  LDAP Role                                Users
app01_group  no          cn=XIVapp01_group,CN=Users,DC=xivstorage,DC=org
```

Alternatively, the same configuration steps can be accomplished through the XIV GUI. To change the LDAP configuration settings in the XIV GUI, open the “Tools” menu at the top of

the main XIV Storage Manager panel, select **Configure** → **LDAP** → **Role Mapping**, and change the configuration parameter settings.

Important: The XIV configuration parameters “Storage Admin Role” and “Read Only Role” can only accept a string of up to 64 characters long. In certain cases the length of the domain name might prevent you from using the `memberOf` attribute for role mapping because the domain name is encoded in the attribute value (“DC=xivstorage,DC=org” in this example represents the xivstorage.org domain name).

Now, by assigning Active Directory group membership, you can grant access to the IBM XIV Storage System.

A user in Active Directory can be a member of multiple groups. If this user is a member of more than one group with corresponding role mapping, XIV fails authentication for this user due to the fact that the role cannot be uniquely identified. In Example 5-23, user *xivadproducer10* can be mapped to Storage Admin and Read Only roles, hence the authentication failure followed by the `USER_HAS_MORE_THAN_ONE_RECOGNIZED_ROLE` error message.

Example 5-23 LDAP user mapped to multiple roles authentication failure

```
$ xcli -c "ARCXIVJEMT1" -u xivadproducer10 -p pass2remember ldap_user_test
Error:  USER_HAS_MORE_THAN_ONE_RECOGNIZED_ROLE
Details: User xivadproducer10 has more than one recognized LDAP role.

$ ldapsearch -x -H "ldap://xivstorage.org:389" -b CN=Users,dc=xivstorage,dc=org
-D cn=xivadproducer10,CN=Users,dc=xivstorage,dc=org -w pass2remember
cn=xivadproducer10 memberOf

dn: CN=xivadproducer10,CN=Users,DC=xivstorage,DC=org
memberOf: CN=XIVReadOnly,CN=Users,DC=xivstorage,DC=org
memberOf: CN=XIVStorageadmin,CN=Users,DC=xivstorage,DC=org
```

An LDAP user can be a member of multiple Active Directory groups and successfully authenticate to XIV as long as only one of those groups is mapped to an XIV role. As illustrated in Example 5-24, the user *xivadproducer10* is a member of two Active Directory groups *XIVStorageadmin* and *nonXIVgroup*. Only *XIVStorageadmin* is mapped to an XIV role.

Example 5-24 LDAP user mapped to a single roles authentication success

```
$ xcli -c "ARCXIVJEMT1" -u xivadproducer10 -p pass2remember ldap_user_test
Command executed successfully.

$ ldapsearch -x -H "ldap://xivstorage.org:389" -b CN=Users,dc=xivstorage,dc=org
-D cn=xivadproducer10,CN=Users,dc=xivstorage,dc=org -w pass2remember
cn=xivadproducer10 memberOf

dn: CN=xivadproducer10,CN=Users,DC=xivstorage,DC=org
memberOf: CN=nonXIVgroup,CN=Users,DC=xivstorage,DC=org
memberOf: CN=XIVStorageadmin,CN=Users,DC=xivstorage,DC=org
```

After all Active Directory groups are created and mapped to corresponding XIV roles, the complexity of managing LDAP user accounts will be significantly reduced because the role mapping can now be done through Active Directory group membership management.

The easy to use point and click interface leaves less room for error when it comes to assigning group membership, as opposed to entering text into description field as previously described.

5.5.5 SUN Java Directory group membership and XIV role mapping

SUN Java Directory group membership can be used for XIV role mapping similar to the way described in 5.5.4, “Active Directory group membership and XIV role mapping” on page 179.

In SUN Java Directory, a user can be a member of a single or multiple groups. A group is a collection of users with common characteristics. Groups can defined anywhere in DIT in SUN Java Directory. A group is represented as a separate object in LDAP Directory Information Tree (DIT) and gets a distinguished name (DN) assigned to it.

Groups defined in SUN Java Directory can be used for XIV role mapping. When a user becomes member of a group in SUN Java Directory, he or she gets a new attribute assigned. The value of the new attribute points to the DN of the group. *IsmemberOf* is the name of that attribute. The *IsmemberOf* attribute value determines the SUN Java Directory group membership.

To create a group in SUN Java Directory using SUN Java Web Console tool:

1. Point your Web browser to HTTPS port 6789, in our case, “https://xivhost2.storage.tucson.ibm.com:6789”
2. Login to the system and select the “Directory Service Control Center (DSCC)” application and Authenticate to Directory Service Manager.
3. Select **Directory Servers** tab, click **xivhost2.storage.tucson.ibm.com:389** and select **Entry Management** tab. Verify that dc=xivauth suffix is highlighted in the left panel and click **New Entry** in the Selected Entry panel.
4. Accept dc=xivauth in Entry Parent DN: field → **Next** → **Entry Type** “Static Group - (groupOfUniqueNames)” → **Next** → and enter attribute values as shown in Figure 5-61.

Required Attributes * Indicates required field

Naming Attribute:

* Full Name (cn):

Allowed Attributes

Group Member (uniqueMember):

Organization (o):

Organizational Unit (ou):

businessCategory:

description:

owner:

seeAlso:

Figure 5-61 Creating new group in SUN Java Directory

Confirm your choice by clicking **Finish** in the panel that is partially shown in Figure 5-62.

Review your settings and click finish if they are correct.

Entry DN:	cn=XIVReadOnly,dc=xivauth
Object Class:	Static Group - (groupOfUniqueNames)
Full Name (cn):	XIVReadOnly
Group Member (uniqueMember):	uid=xivsunproduser3,dc=xivauth

Figure 5-62 Confirming group creation in SUN Java Directory

To illustrate the new *IsmemberOf* attribute in the existing LDAP user object and the new LDAP object representing the “XIVReadOnly” group, we run the **ldapsearch** queries against SUN Java Directory LDAP server as shown in Example 5-25.

Example 5-25 SUN Java Directory group membership ldapsearch queries

```
$ ldapsearch -x -H ldap://xivhost2.storage.tucson.ibm.com:389 -D
uid=xivsunproduser3,dc=xivauth -w passwd -b dc=xivauth uid=xivsunproduser3
ismemberof
```

```
dn: uid=xivsunproduser3,dc=xivauth
ismemberof: cn=XIVReadOnly,dc=xivauth
```

```
$ ldapsearch -x -H ldap://xivhost2.storage.tucson.ibm.com:389 -D
uid=xivsunproduser3,dc=xivauth -w passwd -b dc=xivauth cn=XIVReadOnly
dn: cn=XIVReadOnly,dc=xivauth
objectClass: groupOfUniqueNames
objectClass: top
uniqueMember: uid=xivsunproduser3,dc=xivauth
cn: XIVReadOnly
```

In the first **ldapsearch** query, we intentionally limited our search to the *ismemberof* attribute (at the end of the **ldapsearch** command) so that the output is not obscured with unrelated attributes and values. The value of the *ismemberof* attribute contains the DN of the group. The second **ldapsearch** query illustrates the CN=XIVReadOnly LDAP object content. Among other attributes, it contains the *uniqueMember* attribute, which points at the DN of the user defined as a member. The attribute *uniqueMember* is a multivalued attribute, and there can be more than one user assigned to the group as a member. *Ismemberof* is also a multivalued attribute, and a user can be a member of multiple groups.

XIV can now be configured to use the *ismemberof* attribute for role mapping. In example Example 5-26, we are mapping the SUN Java Directory group XIVReadOnly to the XIV read_only_role, XIVStorageadmin to the storage_admin_role, and the XIV user group app01_group to the SUN Java Directory group XIVapp01_group. You must be logged in to the XCLI as *admin*.

Example 5-26 Configuring XIV to use SUN Java Directory groups for role mapping

```
$ xcli -c "ARCXIVJEMT1" -u admin -p s8cur8pwd ldap_config_set
xiv_group_attr=ismemberof
Command executed successfully.
>> ldap_config_set read_only_role=cn=XIVReadOnly,dc=xivauth
Command executed successfully.
```

```
>> ldap_config_set storage_admin_role=cn=XIVStorageAdmin,dc=xivauth
Command executed successfully.
```

```
>> ldap_config_get
```

Name	Value
base_dn	dc=xivauth
xiv_group_attr	ismemberof
third_expiration_event	7
version	3
user_id_attr	uid
current_server	
use_ssl	no
session_cache_period	10
second_expiration_event	14
read_only_role	cn=XIVReadOnly,dc=xivauth
storage_admin_role	cn=XIVStorageAdmin,dc=xivauth
first_expiration_event	30
bind_time_limit	30

```
>> user_group_update user_group=app01_group ldap_role=cn=XIVapp01_group,dc=xivauth
Command executed successfully.
```

```
>> user_group_list user_group=app01_group
```

Name	Access All	LDAP Role	Users
app01_group	no	cn=XIVapp01_group,dc=xivauth	

Alternatively, the same configuration steps can be accomplished through the XIV GUI. To change the LDAP configuration settings in XIV GUI open “Tools” menu at the top of main XIV Storage Manager panel, select **Configure** → **LDAP** → **Role Mapping** and change the configuration parameter settings.

Important: The XIV configuration parameters “Storage Admin Role” and “Read Only Role” can only accept a string of up to 64 characters long. In certain cases, the length of the distinguished name (DN) might prevent you from using the *ismemberof* attribute for role mapping because the DN is encoded in the attribute value (“dc=xivauth”) in this example.

Now, by assigning SUN Java Directory group membership, you can grant access to the IBM XIV Storage System.

A user in the SUN Java Directory can be a member of multiple groups. If this user is a member of more than one group with corresponding role mapping, XIV fails authentication for this user because the role cannot be uniquely identified. In Example 5-27, user xivsunproduser3 can be mapped to the Storage Admin and Read Only roles, hence the authentication failure followed by the USER_HAS_MORE_THAN_ONE_RECOGNIZED_ROLE error message.

Example 5-27 LDAP user mapped to multiple roles authentication failure

```
$ xcli -c "ARCXIVJEMT1" -u xivsunproduser3 -p pass2remember ldap_user_test
```

```
Error: USER_HAS_MORE_THAN_ONE_RECOGNIZED_ROLE
```

```
Details: User xivsunproduser3 has more than one recognized LDAP role.
```

```
$ ldapsearch -x -H ldap://xivhost2.storage.tucson.ibm.com:389 -D
```

```
uid=xivsunproduser3,dc=xivauth -w passw0rd -b dc=xivauth uid=xivsunproduser3
ismemberof
```

```
dn: uid=xivsunproduser3,dc=xivauth
ismemberof: cn=XIVReadOnly,dc=xivauth
ismemberof: cn=XIVStorageAdmin,dc=xivauth
```

An LDAP user can be a member of multiple SUN Java Directory groups and successfully authenticate in XIV as long as only one of those groups is mapped to an XIV role. As illustrated in Example 5-28, the user xivsunproduser3 is a member of two SUN Java Directory groups, XIVStorageadmin and nonXIVgroup. Only XIVStorageadmin is mapped to an XIV role.

Example 5-28 LDAP user mapped to a single roles authentication success

```
$ xcli -c "ARCXIVJEMT1" -u xivsunproduser3 -p pass2remember ldap_user_test
Command executed successfully.

$ ldapsearch -x -H ldap://xivhost2.storage.tucson.ibm.com:389 -D
uid=xivsunproduser3,dc=xivauth -w passw0rd -b dc=xivauth uid=xivsunproduser3
ismemberof

dn: uid=xivsunproduser3,dc=xivauth
ismemberof: cn=XIVReadOnly,dc=xivauth
ismemberof: cn=nonXIVgroup,dc=xivauth
```

After all SUN Java Directory groups are created and mapped to corresponding XIV roles, the complexity of managing LDAP user accounts will be significantly reduced because the role mapping can now be done through SUN Java Directory group membership management. The easy to use point and click interface leaves less room for error when it comes to assigning group membership, as opposed to entering text into the description field as previously described.

5.5.6 Managing multiple systems in LDAP authentication mode

The task of managing multiple XIV Storage Systems can be significantly simplified by using LDAP authentication mode. Because user credentials are stored centrally in the LDAP directory, it is no longer necessary to synchronize user credentials among multiple XIV systems. After a user account is registered in LDAP, multiple XIV systems can use credentials stored in LDAP directory for authentication.

Because the user's password is stored in the LDAP directory, all connected XIV systems will authenticate the user with this password, and if the password is changed, all XIV systems will automatically accept the new password. This mode of operation is often referred to as "Single Sign-On". This mode allows for quick transitions between systems in the XIV GUI because the password has to be entered only once. This approach is especially useful in Remote Mirror configurations, where the storage administrator is required to frequently switch from source to target system.




Figure 5-63 LDAP Single Sign-On

Important: To allow single sign-on in LDAP authentication mode, all XIV systems must be configured to use the same set of LDAP configuration parameters for role mapping. If role mapping is set up differently on any two XIV systems, it is possible that a user can login to one but not the other XIV system.

5.6 Securing LDAP communication with SSL

In any authentication scenario, information is exchanged between the LDAP server and XIV system where access is being sought. Security Socket Layer (SSL) can be used to implement secure communications between the LDAP client and server. LDAPS (LDAP over SSL, the secure version of LDAP protocol) allows secure communications between the IBM XIV Storage System and LDAP server with encrypted SSL connections. This allows a setup where user passwords never appear on the wire in clear text.

SSL provides methods for establishing identity using X.509 certificates and ensuring message privacy and integrity using encryption. In order to create an SSL connection, the LDAP server must have a digital certificate signed by a trusted certificate authority (CA). Companies have the choice of using a trusted third-party CA or creating their own certificate authority. In this scenario, the xivauth.org CA will be used for demonstration purposes.

To be operational, SSL has to be configured on both the client and the server. Server configuration includes generating a certificate request, obtaining a server certificate from a certificate authority (CA), and installing the server and CA certificates.

See Appendix A, "Additional LDAP information" on page 279 for guidance on how to configure Windows Server and SUN Java Directory for SSL support. You can then proceed to "Configuring XIV to use LDAP over SSL".

5.6.1 Configuring XIV to use LDAP over SSL

To be operational, SSL has to be configured on both the IBM XIV Storage System and the LDAP server. Client (XIV system) configuration includes uploading CA certificates to XIV and enabling SSL mode. The cacert.pem file is ready to be uploaded to the IBM XIV Storage System.

When a new LDAP server is added to the IBM XIV Storage System configuration, a security certificate can be entered in the optional certificate field. If the LDAP server was originally added without a certificate, you will need to remove that definition first and add new definition with the certificate.

LDAP server: When defining the LDAP server with a security certificate in XIV, the fully qualified name of the LDAP server must match the “issued to name” in the client’s certificate.

For registering the LDAP server with security certificate, it might be easier to use the XIV GUI as it has file upload capability (see Figure 5-64). XCLI can also be used, but in this case you need to cut and paste a very long string containing the certificate into the XCLI session. To define the LDAP server in the XIV GUI, open the *Tools* menu at the top of the main XIV Storage Manager panel, and select **Configure** → **LDAP** → **Servers** (green ‘+’ sign on the right panel).




Figure 5-64 Defining Active Directory LDAP server with SSL certificate

In Figure 5-64, the Server Type selected must correspond to your specific LDAP directory, either Microsoft Active Directory as shown, or Sun Directory.

To verify that the LDAP server is defined with the correct certificate, compare the certificate expiration date as it is registered in XIV with the “Valid to” date as shown in Figure A-12 on page 297 (for Microsoft Active Directory) or Figure A-18 on page 303 (for SUN Java Directory).

To view the expiration date of the installed certificate in XIV GUI, open the *Tools* menu at the top of the main XIV Storage Manager panel, select **Configure** → **LDAP** → **Servers**, and double-click on a name of the LDAP server. In our example, the expiration date of the certificate as shown by XIV system in Figure 5-65 matches the “Valid to” date as shown in Figure A-12 on page 297 (for Microsoft Active Directory) or Figure A-18 on page 303 (for SUN Java Directory).




Figure 5-65 Viewing Active Directory server certificate expiration date

By default, LDAP authentication on XIV is configured to use non-SSL communication. To enable the use of SSL in the XIV GUI, open the *Tools* menu at the top of main XIV Storage Manager panel, select **Configure** → **LDAP** → **Parameters** → **Use SS**, and change it from “No” to “Yes” as shown in Figure 5-66.

The image shows the 'LDAP Configuration' dialog box with the 'Parameters' tab selected. The 'Use SSL' dropdown menu is set to 'Yes' and is circled in red. Other fields include LDAP Version (3), Base Domain Name (storage,DC=tucson,DC=ibm,DC=com), Bind Timeout (30), Session Cache Period (10), First Expiration Event (30), Second Expiration Event (14), and Third Expiration Event (7). Buttons for 'Update' and 'Cancel' are at the bottom.

Field	Value
LDAP Version	3
Base Domain Name	storage,DC=tucson,DC=ibm,DC=com
Use SSL	Yes
Bind Timeout (seconds)	30
Session Cache Period (minutes)	10
First Expiration Event	30
Second Expiration Event	14
Third Expiration Event	7

Figure 5-66 Enabling SSL for Active Directory LDAP communication

5.6.2 Maintaining the SSL certificate

A new SSL certificate must be installed before the existing one expires. If you let your SSL certificate expire, XIV LDAP authentication will no longer be possible until you either disable SSL or install the new certificate on both the LDAP server and the XIV. Before the SSL certificate expires, XIV will issue three notification events. The first “SSL Certificate is About to Expire” event can be seen in Figure 5-67.

The image shows an event notification window with the following details:

Severity:	Warning
Date:	2009-06-17 17:16:50
Index:	1592
Event Code:	LDAP_SSL_CERTIFICATE_IS_ABOUT_TO_EXPIRE
T. Shooting:	None
Description:	SSL Certificate of LDAP server 'xivhost2.storage.tucson.ibm.com' is about to expire on 2009-07-17 00:03:48 (first notification).

Figure 5-67 First notification of SSL Certificate of LDAP server expiration

5.7 XIV audit event logging

The XIV Storage System uses a centralized event log. For any command that has been executed that leads to a change in the system, an event entry is generated and recorded in the event log. The object creation time and the user are also as object attributes.

The event log is implemented as a circular log and is able to hold a set number of entries. When the log is full, the system wraps back to the beginning. If you need to save the log entries beyond what the system will normally hold, you can issue the XCLI command **event_list** and save the output to a file.

Event entries can be viewed by the GUI, XCLI commands, or by notification. A flexible system of filters and rules allows you to generate customized reports and notifications. For details about how to create customized rules, see 5.7.4, “Defining notification rules” on page 195.

5.7.1 Viewing events in the XIV GUI

The XIV GUI provides a convenient and easy to use view of the event log. To get to the view shown in Figure 5-68, click the Monitor icon from the main GUI window and the select **Events** from the context menu.

The window is split into two sections:

- ▶ The top part contains the management tools, such as wizards in the menu bar and a series of input fields and drop-down menus that act as selection filters.
- ▶ The bottom part is a table displaying the events according to the selection criteria. Use the table tile bar or headings to enable or change sort direction.

Date	Event Code	User	Description
2009-07-12 16:22:28	USER_FAILED_TO_RUN_COMMAND		User 'admin' failed authentication when trying to run c...
2009-07-12 16:22:23	LDAP_AUTHENTICATION_DEACTIVATED	admin	LDAP authentication deactivated.
2009-07-12 16:20:50	LDAP_SERVER_REMOVED	admin	LDAP server 'xivstorage.org' was removed from the sy...
2009-07-12 16:20:44	LDAP_SERVER_ADDED		LDAP server 'xivstorage' was added to the system.
2009-07-12 16:19:57	LDAP_SERVER_IS_INACCESSIBLE		LDAP server xivstorage.org is inaccessible.
2009-07-12 16:08:37	HOST_NO_MULTIPATH_ONLY_ONE_MOD...		Host 'itso_win_node1' is connected to the system thro...
2009-07-12 15:26:43	HOST_NO_MULTIPATH_ONLY_ONE_MOD...		Host 'itso_win_node1' is connected to the system thro...

Figure 5-68 GUI events main view

The system will progressively load the events into the table. A progress indicator is visible at the bottom right of the table, as shown in Figure 5-69.




Figure 5-69 Loading events into the table

5.7.2 Event attributes

This section overviews all available event types, event codes, and their severity levels.

Severity levels

You can select one of six possible severity levels as the minimal level to be displayed:

- ▶ **none:** Includes all severity levels
- ▶ **informational:** Changes, such as volume deletion, size changes, or host multipathing
- ▶ **warning:** Volume usage limits reach 80%, failing message sent
- ▶ **minor:** Power supply power input loss, volume usage over 90%, component TEST failed
- ▶ **major:** Component failed (disk), user system shutdown, volume and pool usage 100%, UPS on battery or Simple Mail Transfer Protocol (SMTP) gateway unreachable
- ▶ **critical:** Module failed or UPS failed

Event codes

See the *XCLI Reference Guide*, GC27-2213-00, for a list of event codes.

Event types

The following event types can be used as filters (specified with the parameter **object_type** in the XCLI command):

- ▶ **cons_group:** consistency group
- ▶ **destgroup:** event destination group
- ▶ **dest:** event notification address
- ▶ **dm:** data migration
- ▶ **host:** host
- ▶ **map:** volume mapping
- ▶ **mirror:** mirroring
- ▶ **pool:** pool
- ▶ **rule:** rule
- ▶ **smsgw:** sms gateway
- ▶ **smtpgw:** smtp gateway
- ▶ **target:** fc/iSCSI connection
- ▶ **volume:** volume
- ▶ **cluster:** cluster
- ▶ **user:** user
- ▶ **user_group:** user group
- ▶ **ip_interface:** ip interface
- ▶ **ldap_conf:** ldap configuration

5.7.3 Viewing events in the XCLI

Table 5-6 provides a list of all the event-related commands available in the XCLI. This list covers setting up notifications and viewing the events in the system. See Chapter 7, “Monitoring” on page 203 for a more in-depth discussion of system monitoring.

Table 5-6 XCLI: All event commands

Command	Description
custom_event	Generates a custom event.
dest_define	Defines a new destination for event notifications.

Command	Description
dest_delete	Deletes an event notification destination.
dest_list	Lists event notification destinations.
dest_rename	Renames an event notification destination.
dest_test	Sends a test message to an event notification destination.
dest_update	Updates a destination.
destgroup_add_dest	Adds an event notification destination to a destination group.
destgroup_create	Creates an event notification destination group.
destgroup_delete	Deletes an event notification destination group.
destgroup_list	Lists destination groups.
destgroup_remove_dest	Removes an event notification destination from a destination group.
destgroup_rename	Renames an event notification destination group.
event_clear	Clears alerting events.
event_list	Lists system events.
event_list_uncleared	Lists uncleared alerting events.
event_redefine_threshold	Redefines the threshold of a parameterized event.
smsgw_define	Defines a Short Message Service (SMS) gateway.
smsgw_delete	Deletes an SMS gateway.
smsgw_list	Lists SMS gateways.
smsgw_prioritize	Sets the priorities of the SMS gateways for sending SMS messages.
smsgw_rename	Renames an SMS gateway.
smsgw_update	Updates an SMS gateway.
smtpgw_define	Defines an SMTP gateway.
smtpgw_delete	Deletes a specified SMTP gateway.
smtpgw_list	Lists SMTP gateways.
smtpgw_prioritize	Sets the priority of which SMTP gateway to use to send emails.
smtpgw_rename	Renames an SMTP gateway.
smtpgw_update	Updates the configuration of an SMTP gateway.
rule_activate	Activates an event notification rule.
rule_create	Creates an event notification rule.
rule_deactivate	Deactivates an event notification rule.
rule_delete	Deletes an event notification rule.
rule_list	Lists event notification rules.
rule_rename	Renames an event notification rule.
rule_update	Updates an event notification rule.

Event_list command and parameters

The syntax of the event_list command is:

```
event_list
[max_events=MaxEventsToList]
[after=<afterTimeStamp|ALL>]
[before=<beforeTimeStamp|ALL>]
[min_severity=<informational|warning|minor|major|critical> ]
[alerting=<yes|no>]
[cleared=<yes|no>]
[code=EventCode]
[object_type=<cons_group | destgroup | dest|dm | host | map | mirror | pool | rule
smgsw | smtpgw | target | volume> ]
[ beg=BeginIndex ]
[ end=EndIndex ]
[ internal=<yes|no|all> ]
```

XCLI examples

To illustrate how the commands operates, the **event_list** command displays the events currently in the system. Example 5-29 shows the first few events logged in our system.

Example 5-29 XCLI viewing events

```
C:\XIV>xcli -c -c "ARCXIVJEMT1" event_list
Timestamp          Severity          Code           User   Description
2009-07-09 12:49:44 Informational    UNMAP_VOLUME   ITS0   Volume with name 'Win_Vol_2' was
unmapped from cluster with name 'Win2008Cluster01'.
2009-07-09 12:49:48 Informational    UNMAP_VOLUME   ITS0   Volume with name 'Win_Vol_1' was
unmapped from cluster with name 'Win2008Cluster01'.
2009-07-09 12:50:05 Informational    UNMAP_VOLUME   ITS0   Volume with name 'Win_Vol_3' was
unmapped from host with name 'Win2008C1H1'.
2009-07-09 12:50:06 Informational    MAP_VOLUME     ITS0   Volume with name 'Win_Vol_3' was
mapped to LUN '1' for host with name 'Win2008C1H1'.
2009-07-09 12:50:24 Informational    UNMAP_VOLUME   ITS0   Volume with name 'Win_Vol_4' was
unmapped from host with name 'Win2008C1H2'.
.....
```

Example 5-30 illustrates the command for listing all instances when the user was updated. The USER_UPDATED event is generated when a user's password, email, or phone number is modified. In this example, the -t option is used to display specific fields, such as index, code, description of the event, time stamp, and user name. The description field provides the ID that was modified, and the user field is the ID of the user performing the action.

Example 5-30 View USER_UPDATED event with the XCLI

```
C:\XIV>xcli -c "ARCXIVJEMT1" -t index,code,description,timestamp,user_name event_list
code=USER_UPDATED
Index Code          Description                                     Timestamp          User
1089  USER_UPDATED  User with name 'admin' was updated.             2009-07-10 21:45:50 admin
1091  USER_UPDATED  User with name 'admin' was updated.             2009-07-10 21:45:51 admin
1092  USER_UPDATED  User with name 'technician' was updated.         2009-07-10 21:45:52
xiv_development
1093  USER_UPDATED  User with name 'technician' was updated.         2009-07-10 21:45:53
xiv_development
```

5.7.4 Defining notification rules

Example 5-31 describes how to set up a rule in the XCLI to notify the storage administrator when a user's access control has changed. The rule itself has four event codes that generate a notification. The events are separated with commas with no spaces around the commas. If any of these four events are logged, the XIV Storage System uses the "relay" destination to issue the notification.

Example 5-31 Setting up an access notification rule using the XCLI

```
C:\XIV>xcli -c -c "ARCXIVJEMT1" rule_create rule=test
codes=ACCESS_OF_USER_GROUP_TO_CLUSTER_REMOVED,ACCESS_OF_USER_GROUP_TO_HOST_REMOVED
,ACCESS_TO_CLUSTER_GRANTED_TO_USER_GROUP,ACCESS_TO_HOST_GRANTED_TO_USER_GROUP
dests=relay
Command executed successfully.
```

A simpler example is setting up a rule notification for when a user account is modified. Example 5-32 creates a rule on the XIV Storage System called ESP that sends a notification whenever user account is modified on the system. The notification is transmitted through the relay destination.

Example 5-32 Create a rule for notification with the XCLI

```
C:\XIV>xcli -c -c "ARCXIVJEMT1" rule_create rule=user_update codes=USER_UPDATED
dests=relay
Command executed successfully.
```

The same rule can be created in the GUI. See Chapter 7, "Monitoring" on page 203 for more details about configuring the system to provide notifications and setting up rules.



Performance characteristics

In Chapter 2, “IBM XIV logical architecture and concepts” on page 13, we have described the XIV Storage System’s parallel architecture, disk utilization, and unique caching algorithms. These characteristics, inherent to the system design, deliver optimized and consistent performance regardless of the workload the XIV Storage System endures.

In the current chapter, we further explore the concepts behind this high performance. We discuss best practices to follow when connecting to an XIV Storage System, and explain how to extract statistics provided by the XIV Storage System Graphical User Interface (GUI) and Extended Command Line Interface (XCLI). In addition to the XIV GUI Performance Statistics options, we also show how to use XIV TOP for additional and more granular system performance statistics.

6.1 Performance concepts

The XIV Storage System maintains a high level of performance by leveraging all the disk, memory, and I/O resources in the system at all times. The system offers advanced caching features, employing effective data mirroring, and integrating power copy services functionality such as snapshots and remote mirroring. These characteristics are the basis for many unique features that distinguish XIV from its competition.

In this chapter we cover these functions as they pertain to the XIV Storage System:

- ▶ Full disk resource utilization
- ▶ Caching mechanisms
- ▶ Data mirroring
- ▶ Snapshots

6.1.1 Full disk resource utilization

Utilization of all disk resources improves the performance by minimizing the bottlenecks within the system. The XIV Storage System stripes and mirrors data into 1 MB partitions across all the disks in the system; it then disperses the 1 MB partitions in a pseudo-random distribution. This pseudo-random distribution results in a lower access density, which is measured by throughput divided by the total disk capacity. See Chapter 2, “IBM XIV logical architecture and concepts” on page 13 for further details about the architecture of the system.

Various benefits result from fully utilizing all of the disk resources. Each disk drive performs an equal workload as the data is balanced across the entire system. The pseudo-random distribution ensures load balancing at all times and eliminates *hot-spots* in the system.

6.1.2 Caching mechanisms

The XIV Storage System caching management is unique, by dispersing the cache into each module as opposed to a central memory cache. The distributed cache enables each module to concurrently service host I/Os and cache to disk access, as opposed to the central memory caching algorithm, which implements memory locking algorithms that generate access contention.

To improve memory management, each Data Module uses a PCI Express (PCI-e) bus between the cache and the disk modules, which provides a sizable interconnect between the disk and the cache. This design aspect allows large amounts of data to be quickly transferred between the disks and the cache by the bus.

Having a large bus “pipe” permits the XIV Storage System to have small cache pages. More so, a large bus “pipe” between the disk and the cache allows the system to perform many small requests in parallel, again improving the performance.

A Least Recently Used (LRU) algorithm is the basis for the cache management algorithm. This feature allows the system to generate a high hit ratio for frequently utilized data. In other words, the efficiency of the cache usage for small transfers is very high, when the host is accessing the same data set.

The cache algorithm starts with a single 4 KB page and gradually increase the number of pages prefetched until an entire partition, 1 MB, is read into cache. If the access results in a cache hit, the algorithm doubles the amount of data prefetched into the system.

The prefetching algorithm continues to double the prefetch size until a cache miss occurs, or the prefetch size maximum of 1 MB is obtained. Because the modules are managed independently if a prefetch crosses a module boundary, then the logically adjacent module (for that volume) is notified in order to begin pre-staging the data into its local cache.

6.1.3 Data mirroring

The XIV Storage System maintains two copies of each 1 MB data partition, referred to as the *primary partition* and *secondary partition*. The primary partition and secondary partition for the same data are also kept on separate disks in separate modules. We call this *data mirroring*.

By implementing data mirroring, the XIV Storage System performs a single disk access on reads and two disk accesses on writes; one access for the primary copy and one access for the mirrored secondary copy. Other storage systems that use RAID architecture might translate the I/O into two disk writes and two disk reads for RAID 5 and three disk writes and three disk reads for RAID 6. This allows the XIV Storage System data mirroring algorithm reduce the disk access times and provide quicker responses to requests.

A 1 MB partition is the amount of data stored on a disk with a contiguous address range. Because the cache operates on 4 KB pages, the smallest chunk of data that can be staged into cache is a single cache page, or 4 KB. The data mirroring only mirrors the data that has been modified. By only storing modified data, the system performs at maximum efficiency.

It is also important to note that the data mirroring scheme is not the same as a data stripe in a RAID architecture. Specifically, a *partition* refers to a contiguous region on a disk. The partition is how the XIV Storage System tracks and manages data. This large 1 MB partition is not the smallest workable unit. Because the cache uses 4 KB pages, the system can stage and modify smaller portions of data within the partition, which means that the large partition assists the sequential workloads and the small cache page improves performance for the random workloads.

Disk rebuilds in which data is populated to a new disk after replacement of a phased out or failed disk gains an advantage due to the mirroring implementation. When a disk fails in the RAID 5 or RAID 6 system, a spare disk is added to the array and the data is reconstructed on that disk using the parity data. The process can take a few hours based on the size of the disk drive. With the XIV Storage System, the rebuild is not focused on one disk. After the disk is replaced, the system enters a redistribution phase. In the background, the work is spread across all the disks in the system as it slowly moves data back onto to the new disk. By spreading the work across all the disks, each disk is performing a small percentage of work, therefore the impact to the host is minimal.

6.1.4 Snapshots

Snapshots complete nearly instantly within the XIV Storage System. When a snapshot is issued, no data is copied but rather the snapshot creates system pointers to the original data. As the host writes modified data in the master volume, the XIV Storage System redirects the write data to a new partition. Only the data that was modified by the host is copied into the new partition, which prevents moving the data multiple times and simplifies the internal management of the data. See the IBM Redbooks publication: *IBM XIV Storage System Copy Services and Migration*, SG24-7759 for details about how the snapshot function is implemented.

6.2 Best practices

Tuning of the XIV Storage System is not required by design. Because the data is balanced across all the disks, the performance is at maximum efficiency. This section is dedicated to external considerations that enable maximum performance. The settings given in this section are host-agnostic and are general rules when operating the XIV Storage System.

6.2.1 Distribution of connectivity

The main goal for the host connectivity is to create a balance of the resources in the XIV Storage System. Balance is achieved by distributing the physical connections across the Interface Modules. A host usually manages multiple physical connections to the storage device for redundancy purposes by a SAN connected switch. It is ideal to distribute these connections across each of the Interface Modules. This way the host utilizes the full resources of each module that is connected to and can obtain maximum performance. It is important to note that it is not necessary for each host instance to connect to each Interface Module. However, when the host has more than one physical connection, it is beneficial to have the cables divided across the modules.

Similarly, if multiple hosts have multiple connections, make sure to spread the connections evenly across the Interface Modules. See 3.2.5, “Patch panel” on page 63.

6.2.2 Host configuration considerations

There are various key points when configuring the host for optimal performance. Because the XIV Storage System is distributing the data across all the disks an additional layer of volume management at the host, such as Logical Volume Manager (LVM), might hinder performance for workloads. Multiple levels of striping can create an imbalance across a specific resource. Therefore, it is best to disable host striping of data for XIV Storage System volumes and allow the XIV Storage System to manage the data.

Based on your host workload, you might need to modify the maximum transfer size that the host generates to the disk to obtain the peak performance. For applications with large transfer sizes, if a smaller maximum host transfer size is selected, the transfers are broken up, causing multiple round-trips between the host and the XIV Storage System. By making the host transfer size as large or larger than the application transfer size, fewer round-trips occur, and the system experiences improved performance. If the transfer is smaller than the maximum host transfer size, the host only transfers the amount of data that it has to send.

Due to the distributed data features of the XIV Storage System, high performance is achieved by parallelism. Specifically, the system maintains a high level of performance as the number of parallel transactions occur to the volumes. Ideally, the host workload can be tailored to use multiple threads or spread the work across multiple volumes.

Changing the queue depth

The XIV Storage architecture was designed to perform under real-world customer production workloads, with lots of I/O requests at the same time. Queue depth is an important host bus adapter (HBA) setting because it essentially controls how much data is allowed to be “in flight” onto the SAN from the HBA. A queue depth of 1 requires that each I/O request be completed before another is started. A queue depth greater than one indicates that multiple host I/O requests might be waiting for responses from the storage system. So, the higher the host HBA queue depth, the more parallel I/O goes to the XIV Storage System.

The XIV Storage architecture eliminates the legacy storage concept of a large central cache. Instead, each component in the XIV grid has its own dedicated cache. The XIV algorithms that stage data between disk and cache work most efficiently when multiple I/O requests are coming in parallel - this is where the queue depth host parameter becomes an important factor in maximizing XIV Storage I/O performance.

Sample queue depth comparison

Figure 6-1 shows a queue depth comparison for a database I/O workload (70 percent reads, 30 percent writes, 8k block size, DBO = Database Open).

Note that the performance numbers in this example are valid for this special test at an IBM lab only. The numbers do not describe the general capabilities of IBM XIV Storage System as you might observe them in your environment.




Figure 6-1 Host side queue depth comparison

A good practice is starting with a queue depth of 64 per HBA, to ensure exploitation of the XIV's parallel architecture.

Nevertheless the initial queue depth value might need to be adjusted over time. While higher queue depth in general yields better performance with XIV one must consider the limitations per port on the XIV side. Each HBA port on the XIV Interface Module is designed and set to sustain up to 1400 concurrent I/Os (except for port 3 when port 4 is defined as initiator, in which case port 3 is set to sustain up to 1000 concurrent I/Os). With a queue depth of 64 per host port as suggested, one XIV port is limited to 21 concurrent host ports given that each host will fill up the entire 64 depth queue for each request.

If in a very large environment, the relationship of 21 host ports connected to one XIV port is not sufficient, lower queue depth values have to be configured. This method can also be used as a "poor man's" Quality of Service (QoS) mechanism.

6.2.3 XIV sizing validation

Currently, your IBM Representative provides sizing settings based on the workload.

If I/O workload data of an existing or comparable application environment is available, this information can be used in the IBM storage subsystem sizing tool Disk Magic™. This Windows-based disk subsystem performance modeling tool is licensed exclusively to IBM and IBM Business Partners. Updated versions of Disk Magic are available at regular intervals (typically one per month). Each version of Disk Magic will expire at a pre-assigned date.



Monitoring

In this chapter we describe the various methods and functions that are available to monitor the IBM XIV Storage System. We also show how you can gather information from the system in real time, in addition to the self-monitoring, self-healing, and automatic alerting function implemented within the XIV software.

We also discuss the Call Home function, Secure Remote Support, and repair procedures.

7.1 System monitoring

The IBM XIV Storage System software includes features that allow you to monitor the system:

- ▶ You can review or request the current system status and performance statistics at any time.
- ▶ You can set up alerts to be triggered when specific error conditions or problems arise in the system. Alerts can be conveyed as messages to a user, an email, or a Short Message Service (SMS) text to a mobile phone.
- ▶ Depending on the nature or severity of a problem, the system can automatically alert the IBM Remote Support Center, who will immediately initiate the necessary actions to promptly repair the system. Note that CallHome must be configured and working properly for this feature to work.
- ▶ You can configure IBM Tivoli Storage Productivity Center (TPC) to communicate with the IBM XIV Storage System.
- ▶ In addition, the Secure Remote Support feature allows remote monitoring and repair by IBM support personnel.

7.1.1 Monitoring with the GUI

The monitoring functions available from the IBM XIV Storage System GUI allow the user to easily display and review the overall system status, as well as events and various statistics.

These functions are accessible from the Monitor menu as shown in Figure 7-1.




Figure 7-1 GUI monitor functions

Monitoring the system

Selecting **System** from the Monitor menu shown in Figure 7-1 takes you to the System view, shown in Figure 7-2 (note that this view is also the default or main GUI window for the selected system).

The System view shows a graphical representation of the IBM XIV Storage System rack with its components. You can click the curved arrow located at the bottom right of the picture of the rack to display a view of the patch panel.

You get a quick overview in real time of the system's overall condition and the status of its individual components. The display changes dynamically to provide details about a specific component when you position the mouse cursor over that component.



Figure 7-2 Monitoring the IBM XIV

Status bar indicators located at the bottom of the window indicate the overall operational levels of the IBM XIV Storage System:

- The first indicator on the left shows the amount of soft or hard storage capacity currently allocated to Storage Pools and provides alerts when certain capacity thresholds are reached. As the physical, or hard, capacity consumed by volumes within a Storage Pool passes certain thresholds, the color of this meter indicates that additional hard capacity might need to be added to one or more Storage Pools.

Clicking the icon on the right side of the indicator bar that represents up and down arrows will toggle the view between hard and soft capacity.

Our example indicates that the system has a usable hard capacity of 79113 GB, of which 84% or 66743 GB are actually used.

You can also get more detailed information and perform more accurate capacity monitoring by looking at Storage Pools (see 4.3.2, “Managing storage pools with the XIV GUI” on page 104).

- The second indicator in the middle, displays the number of I/O operations per second (IOPS).
- The third indicator on the far right shows the general system status and, for example, indicates when a redistribution is underway.

In our example, the general system status indicator shows that the system is undergoing a Rebuilding phase, which was triggered by a failing disk (Disk 7 in Module 7) as shown in Figure 7-2.

Monitoring events

To get to the Events window, select Events from the Monitor menu as shown in Figure 7-3.

Extensive information and many events are logged by the IBM XIV Storage System. The system captures entries for problems with various levels of severity, including warnings and other informational messages. These informational messages include detailed information about logins, configuration changes, and the status of attached hosts and paths. All of the collected data can be reviewed in the Events window that is shown in Figure 7-3.




Figure 7-3 Events window

Because many events are logged, the number of entries is typically huge.

To get a more useful and workable view, there is an option to filter the events logged. Without filtering the events, it might be difficult to find the entries for a specific incident or information. Figure 7-4 shows the possible filter options for the events.




Figure 7-4 Event filter

If you double-click a specific event in the list, you can get more detailed information about that particular event, along with the best eventual actions to take.

Figure 7-5 shows details for a critical event where a switch link failed. For this type of event, contact IBM XIV Support immediately. Notification happens automatically if call home and remote support are enabled as described in 7.5, “Call Home and Remote support” on page 267. A problem record is generated within IBM XIV Support, and an IBM support person will contact the customer to report the problem and will follow up with the repair actions.



Figure 7-5 Event properties

Event severity

The events are classified into a level of severity depending on their impact on the system. Figure 7-6 gives an overview of the criteria and meaning of the various severity levels.

<u>Severity:</u>		The Events are categorized in these five categories.
	= Critical	“Critical” an event have occurred where one ore more parts have failed and the redundancy and machine operation can be affected.
	= Major	“Major” an event have occurred where a part have failed and the redundancy is temporary affected. (ex: failing disk)
	= Minor	“Minor” an event occurred where a part have failed but system is still fully redundant and have no operational impact
	= Warning	“Warning” information for the user that something in the system have changed but no impact for the system
	= Informational	“Informational” event is for information only without any impact or danger for system operation

Figure 7-6 Event severity

Event notification configuration

The events window offers a Toolbar (Figure 7-7) which contains a Setup wizard, with the ability to view and modify gateways, destinations, and rules, as well as modify the email addresses for the IBM XIV Storage system and to export the events to a comma delimited CSV file for use in an external spreadsheet program.

Clicking the **Setup** icon starts the Events Configuration Wizard, which guides you through the process to create gateways, add destinations, and define rules for event notification.



Figure 7-7 Event rules configuration

For further information about event notification rules, see 7.3, “XIV event notification” on page 237.

In addition to the standard system generated events, a storage administrator can generate a custom_event with their own custom event description and desired severity level using an XCLI command. This is useful when creating scripts to automate a process. The user can generate an informational event stating a process has started or finished as shown in Example 7-1.

Example 7-1 Using custom_event

```
>> custom_event description="User Initiated Backup Has Started" severity=INFORMATIONAL
command 0:
administrator:
  command:
    code = "SUCCESS"
    status = "0"
    status_str = "Command completed successfully"
aserver = "DELIVERY_SUCCESSFUL"
```

```
>> event_list code=CUSTOM_EVENT
Timestamp      Severity      Code          User          Description
-----
2010-10-13 14:15:48  Informational  CUSTOM_EVENT  itso          User Initiated Backup Has Started
```

Monitoring statistics

The Statistics monitor, shown in Figure 7-8, provides information about the performance and workload of the IBM XIV.




Figure 7-8 Monitor statistics

There is flexibility in how you can visualize the statistics. Options are selectable from a control pane located at the bottom of the window, which is shown in Figure 7-9.




Figure 7-9 Statistics filter

For detailed information about performance monitoring, see 7.1, “System monitoring” on page 204.

7.1.2 Monitoring with XCLI

The XIV Command Line Interface (XCLI) provides various commands to monitor the IBM XIV Storage System and gather real-time system status, monitor events, and retrieve statistics. Refer also to 4.1, “IBM XIV Storage Management software” on page 84 for more information about how to set up and use the XCLI.

System monitoring

Various XCLI commands are available for system monitoring as illustrated next. For complete information about these commands, see the *XCLI Users Guide*, which is available at:

<http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp>

The `state_list` command, shown Example 7-2, gives an overview of the general status of the system. In the example, the system is operational, data is fully redundant, and no shutdown is pending.

Example 7-2 The `state_list` command

```
>> state_list
Category          Value
shutdown_reason   No Shutdown
target_state      on
off_type          off
redundancy_status Full Redundancy
system_state      on
safe_mode         no
```

In Example 7-3, the `system_capacity_list` command shows an overview of used and free capacity system-wide. In the example, both the hard and soft usable capacity is 79113 GB, with 54735 GB of free hard capacity, 54302 GB of free soft capacity. It also shows that the all spare capacity is still available.

Example 7-3 The `system_capacity_list` command

```
>> system_capacity_list
Soft  Hard  Free Hard  Free Soft  Spare Modules  Spare Disks  Target Spare Modules  Target Spare Dis
79113 79113 54735    54202     1           3             1               3
```

In Example 7-4, the **version_get** command displays the current version of the XIV code installed on the system. Knowing the current version of your system assists you in determining when upgrades are required.

Example 7-4 The version_get command

```
>> version_get
Version
10.2.2
```

In Example 7-5, the **time_list** command is used to retrieve the current time from the IBM XIV Storage System. This time is normally set at the time of installation. Knowing the current system time is required when reading statistics or events. In certain cases, the system time might differ from the current time (at the user's location), and therefore, knowing when something occurred according to the system time assists with debugging issues.

Example 7-5 The time_list command

```
>> time_list
Time      Date      Time Zone  Daylight Saving Time
08:57:15  2008-08-19  GMT       no
```

System components status

In this section, we present various XCLI commands that are available to get the status of specific system components, such as disks, modules, or adapters.

The **component_list** command that is shown In Example 7-6 gives the status of all hardware components in the system. The filter options **filter=<FAILED | NOTOK>** is used to only return failing components. The example shows a failing disk in module 4 on position 9.

Example 7-6 The component_list command

```
>> component_list filter=NOTOK
Component ID  Status  Currently Functioning
1:Disk:4:9    Failed  no
```

As shown in Example 7-7, the **disk_list** command provides more in-depth information for any individual disk in the IBM XIV Storage System, which might be helpful in determining the root cause of a disk failure. If the command is issued without the disk parameter, all the disks in the system are displayed.

Example 7-7 The disk_list command

```
>> disk_list disk=1:Disk:13:10
Component ID  Status  Currently Functioning  Capacity (GB)  Target Status  Model      Size  Serial
1:Disk:13:11  Failed  yes                    1 TB           Hitachi...    942772  PAJ1W2XF
```

```
>> disk_list disk=1:Disk:13:11
Component ID  Status  Currently Functioning  Capacity (GB)  Target Status  Model      Size  Serial
1:Disk:13:11  Failed  yes                    1 TB           Hitachi...    942772  PAJU02YF
```

In Example 7-8, the **module_list** command displays details about the modules themselves. If the module parameter is not provided, all the modules are displayed. In addition to the status of the module, the output describes the number of disks, number of FC ports, and number of iSCSI ports.

Example 7-8 The module_list command

```
>> module_list module=1:Module:4
```

Component ID	Status	Currently Functioning	Target Status	Type	Data Disks	FC Ports	iSCSI P
1:Module:4	OK	yes		p10hw_auxiliary	12	4	0

It is possible to use the XCLI parameter -x with commands to provide even more information about a component. The -x parameter will display the output of a command in XML format. You can see the additional information displayed in Example 7-9.

Example 7-9 Using -x parameter to display additional information in XML format

```
>> module_list module=1:Module:4 -x
```

```
<command id="0">
<administrator>
  <command>
    <changes_session_id value="1285625276148030:4:1285773388015388:4"/>
    <code value="SUCCESS"/>
    <last_change_index value="1217792"/>
    <machine_session_id value="1285625276148030"/>
    <status value="0"/>
    <status_str value="Command completed successfully"/>
    <return>
      <module id="100225">
        <component_id value="1:Module:4"/>
        <status value="OK"/>
        <currently_functioning value="yes"/>
        <requires_service value="None"/>
        <target_status value=""/>
        <type value="p10hw_auxiliary"/>
        <disk_bay_count value="12"/>
        <fc_port_count value="4"/>
        <ethernet_port_count value="0"/>
        <io_allowed value="yes"/>
        <io_enabling_priority value="0"/>
        <serial value="SHU09296990073A"/>
        <original_serial value="SHU09296990073A"/>
        <part_number value="45W1980"/>
        <original_part_number value="45W1980"/>
        <bmc_version value="0.63-343"/>
        <sdr_version value="SDR Package 46"/>
        <bios_version value="S5000.86B.10.40.0091.090520081447"/>
        <ses_version value="60"/>
        <sas_version value="1.27.00"/>
        <memory_gb value="8"/>
      </module>
    </return>
  </command>
</administrator>
<aserver status="DELIVERY_SUCCESSFUL"/>
</command>
```

In Example 7-10, the **ups_list** command describes the current status of the Uninterruptible Power Supply (UPS) component. It provides details about when the last test was performed and the results. Equally important is the current battery charge level. A non-fully charged battery can be a cause of problems in case of power failure.

The output of the **ups_list** command is broken into two lines for easier reading.

Example 7-10 The ups_list command

```
>> ups_list
```

Component ID	Status	Currently Functioning	Input Power On	Battery Charge Level	Last Self Test Date
1:UPS:1	OK	yes	yes	100	06/23/2009
1:UPS:2	OK	yes	yes	100	06/24/2009
1:UPS:3	OK	yes	yes	100	06/24/2009

Last Self Test Result	Monitoring Enabled	UPS Status
Passed	yes	ON_LINE
Passed	yes	ON_LINE
Passed	yes	ON_LINE

Example 7-11 shows the **switch_list** command that is used to display the current status of the switches.

Example 7-11 The switch_list command

```
>> switch_list
```

Component ID	Status	Currently Functioning	AC Power State	DC Power State	Interconnect	Failed Fans
1:Switch:1	OK	yes	OK	OK	Up	0
1:Switch:2	OK	yes	OK	OK	Up	0

The **psu_list** command shown in Example 7-12 lists all the power supplies in each of the modules. You can display an individual Power Supply Unit (PSU) by using the **psu=** parameter.

Example 7-12 The psu_list command

```
>> psu_list
```

Component ID	Status	Currently Functioning	Hardware Status
1:PSU:1:1	OK	yes	OK
1:PSU:1:2	OK	yes	OK
1:PSU:2:1	OK	yes	OK
1:PSU:2:2	OK	yes	OK
1:PSU:3:1	OK	yes	OK
1:PSU:3:2	OK	yes	OK
.			
.			
.			
.			
1:PSU:12:1	OK	yes	OK
1:PSU:12:2	OK	yes	OK
1:PSU:13:1	OK	yes	OK
1:PSU:13:2	OK	yes	OK
1:PSU:14:1	OK	yes	OK
1:PSU:14:2	OK	yes	OK
1:PSU:15:1	OK	yes	OK
1:PSU:15:2	OK	yes	OK

Events

Events can also be managed with XCLI commands. Various commands are available to list, filter, close, and send notifications for the events. There are many commands and parameters available. You can obtain detailed and complete information in the *IBM XIV XCLI User Manual*.

Next we illustrate just a few of the options of the **event_list** command. Various parameters can be used to sort and filter the output of the **event_list** command. See Table 7-1 for a list of the most commonly used parameters.

Table 7-1 The event_list command parameters

Name	Description	Syntax and example
max_events	Lists a specific number of events	<event_list max_events=100>
after	Lists events after the specified date and time	<event_list after=2008-08-11 04:04:27>
before	Lists events before specified date and time	<event_list before 2008-08-11 14:43:47>
min_severity	Lists events with the specified and higher severities	<event_list min_severity=major>
alerting	Lists events for which an alert was sent or for which no alert was sent	<event_list alerting=no> <event_list alerting=yes>
cleared	Lists events for which an alert was cleared or for which the alert was not cleared	<event_list cleared=yes> <event_list cleared=no>

These parameters can be combined for better filtering. In Example 7-13, two filters were combined to limit the amount of information displayed. The first parameter **max_events** only allows three events to be displayed. The second parameter is the date and time that the events must not exceed. In this case, the event occurred approximately 1.5 minutes before the cutoff time.

Example 7-13 The event_list command with two filters combined

```
>> event_list max_events=5 before=2009-06-29.11:00:00
```

Timestamp	Severity	Code	User	Description
2009-06-29 09:40:43	Informational	HOST_MULTIPATH_OK		Host 'itso_2008' has redundant connection
2009-06-29 09:53:20	Informational	HOST_MULTIPATH_OK		Host 'isabella' has redundant connections
2009-06-29 10:32:28	Informational	HOST_MULTIPATH_OK		Host 'isabella' has redundant connections
2009-06-29 10:58:25	Informational	HOST_MULTIPATH_OK		Host 'isabella' has redundant connections
2009-06-29 10:58:35	Informational	HOST_MULTIPATH_OK		Host 'isabella' has redundant connections

The event list can also be filtered on severity. Example 7-14 displays all the events in the system that contain a severity level of Major and all higher levels, such as Critical.

Example 7-14 The event_list command filtered on severity

```
>> event_list min_severity=Major max_events=5
```

Timestamp	Severity	Code	User	Description
2009-06-26 17:11:36	Major	TARGET_DISCONNECTED		Target named 'XIV MN00035' is no longer accessi
2009-06-26 22:30:33	Major	TARGET_DISCONNECTED		Target named 'XIV MN00035' is no longer accessi
2009-06-28 08:07:41	Major	SWITCH_INTERCONNECT_DOWN		Inter-switch connection lost connectivity on 1
2009-06-28 08:07:41	Major	SWITCH_INTERCONNECT_DOWN		Inter-switch connection lost connectivity on 1
2009-06-29 07:26:58	Critical	MODULE_FAILED		1:Module:9 failed.

Certain events generate an alert message and do not stop until the event has been cleared. These events are called *alerting events* and can be viewed by the GUI or XCLI with a separate command. After the alerting event is cleared, it is removed from this list, but it is still visible with the **event_list** command. See Example 7-15 on page 214.

Example 7-15 The event_list_uncleared command

```
>> event_list_uncleared
No alerting events exist in the system
```

Monitoring statistics

The statistics gathering mechanism is a powerful tool. The IBM XIV Storage System continually gathers performance metrics and stores them internally. Using the XCLI, data can be retrieved and filtered by using many metrics. Example 7-16 provides an example of gathering the statistics for 10 days, with each interval covering an entire day. The system is given a time stamp as the ending point for the data. Due to the magnitude of the data being provided, it is best to redirect the output to a file for further post-processing. See Chapter 6, “Performance characteristics” on page 197 for a more in-depth view of performance.

Example 7-16 Statistics for 10 days

```
>> statistics_get count=10 interval=1 resolution_unit=day end=2009-06-29.17:15:00 > perf.out
```

The **usage_get** command is a useful tool to provide details about the current utilization of pools and volumes. The system saves the usage every hour for later retrieval. This command works the same as the **statistics_get** command. You specify the time stamp to begin or end the collection and the number of entries to collect. In addition, you need to specify the pool name or the volume name. See Example 7-17.

Example 7-17 The usage_get command by pool

```
>> usage_get pool=ITS0_SVC max=10 start=2009-06-22.11:00:00
Time                Volume Usage (MB)    Snapshot Usage (MB)
2009-06-22 11:00:00  0                      0
2009-06-22 12:00:00  0                      0
2009-06-22 13:00:00  0                      0
2009-06-22 14:00:00  262144                 0
2009-06-22 15:00:00  262144                 0
2009-06-22 16:00:00  262144                 0
2009-06-22 17:00:00  262144                 0
2009-06-22 18:00:00  262144                 0
2009-06-22 19:00:00  262144                 0
2009-06-22 20:00:00  262144                 0
```

Note that the usage is displayed in MB. Example 7-18 shows that the volume ITS0_SVC_MDISK_01 is utilizing 46 MB of space. The time when the data was written to the device is also recorded. In this case, the host wrote data to the volume for the first time on 22 June 2009.

Example 7-18 The usage_get command by volume

```
>> usage_get vol=ITS0_SVC_MDISK_01 max=10 start=2009-06-22.11:00:00
Time                Volume Usage (MB)    Snapshot Usage (MB)
2009-06-22 11:00:00  0                      0
2009-06-22 14:00:00  46                     0
2009-06-22 15:00:00  46                     0
2009-06-22 16:00:00  46                     0
2009-06-22 17:00:00  46                     0
2009-06-22 18:00:00  46                     0
2009-06-22 19:00:00  46                     0
```

7.1.3 Using Tivoli Storage Productivity Center

Starting with Version 10.2.2 of the XIV Storage System software, integration with Tivoli Storage Productivity Center (TPC) v4.2 or higher also allows XIV performance monitoring from TPC.

For detailed information about TPC 4.2, see the IBM Redbooks publication, *IBM Tivoli Storage Productivity Center V4.2 Release Guide*, SG24-7725.

TPC is an integrated suite for managing storage systems, capacity, storage networks, even replication.

A storage system that supports and includes a Common Information Model Object Manager (CIMOM) agent (SMI-S compliant / also called SMI-S Agent) can be used directly by TPC. The CIM agent can provide information regarding the configuration of the storage device, such as storage pools, volumes, disks, hosts and host mappings as well as other details or performance characteristics of the storage system. TPC collects that information and stores in the TPC database by using the discovery- and probing process.

Starting with TPC 4.2 a second process for data collection by way of a new, native API is also available. The native API is an internal, native interface to any IBM Storage Systems. Figure 7-10 is a schematic view of both methods.




Figure 7-10 TPC Data collection (prior to version 4.2 on the left and with 4.2 on the right)

The Native API (NAPI) enables additional TPC functions and performance data collection. It can be used with Version 10.2.2 or higher of the BM XIV Storage System software.

Setting up and discovering XIV systems in TPC 4.2

While TPC 4.1.x manages and monitors the XIV through its CIM agent (embedded in the XIV code), TPC 4.2 is using a new data collection process and the Native API to interlace with XIV.

Next we describe the sequence of steps to add an IBM XIV Storage System into your TPC configuration:

1. Create an XIV user and password that will be used by TPC.
2. Note the IBM XIV Storage System management IP address.
3. Start the TPC configure devices wizard to discover and probe the XIV.
4. Create a performance monitor.

XIV user creation

Initially, create a normal XIV user and password on the IBM XIV Storage System by using the XIV GUI as shown in Figure 7-11. This user and password will later on be used in TPC.


The image shows a 'Add User' dialog box with a title bar containing a close button. The dialog contains several input fields: 'Name:' with the value 'tpcadmin', 'New Password: (6-12)' with masked characters, 'Retype New Password:' with masked characters, 'Category:' with a dropdown menu showing 'Read Only', 'User Group:' with a dropdown menu showing 'None', 'Email Address:' with an empty text box, and 'Phone Number:' with two empty text boxes. There are red asterisks next to the Name, New Password, and Retype New Password fields. At the bottom, there are 'Add' and 'Cancel' buttons.

Figure 7-11 User creation

Note the IBM XIV Storage System management IP address

To add any IBM XIV Storage System into TPC, you need to know the management IP addresses of the XIV modules 4, 5, and 6. Figure 7-12 shows how to display the management IP addresses. Simply right-click a particular XIV system in the “All Systems” view from the XIV GUI and select “Modify IP Addresses” from the pop-up menu.


The image shows an 'Edit System XIV LAB 01 EBC' dialog box with a title bar containing a close button. The dialog contains several input fields: 'System Name:' with the value 'XIV LAB 01 EBC', 'IP/Hostname 1:' with the value '9.155.53.250', 'IP/Hostname 2:' with the value '9.155.53.251', 'IP/Hostname 3:' with the value '9.155.53.252', and 'Connect Directly:' with an unchecked checkbox. There is a red asterisk next to the IP/Hostname 1 field. At the bottom, there are 'Update' and 'Cancel' buttons.

Figure 7-12 Displaying the Management IP Addresses

Discovery phase

To add the IBM XIV Storage System into TPC, start the **Configure Devices** Wizard. Just select **IBM Tivoli Storage Productivity Center** → **Configure Devices** as displayed in Figure 7-13. and proceed with the steps to add, discover, and probe an XIV into TPC.




Figure 7-13 Configure Devices Wizard

The discovery usually takes a few minutes and can be run on a schedule. How often you run a discovery depends on the dynamic of your environment. It must be run to detect a new subsystem, also to perform a basic health checks of all CIMOMs and other IBM Storage subsystems.

TPC configuration wizard

Perform the following steps:

1. Select **IBM Tivoli Storage Productivity Center** → **Configure Devices** in order to start the wizard.
2. Select **Storage Subsystem** and click **Next**.
3. Select **Add and configure new storage subsystem** and click **Next**.
4. Select **IBM XIV** as **Device Type** and enter the required information:
 - IP Address: The IP address of the CIMOM. For XIV, this corresponds to the 3 management IP address of the IBM XIV Storage System - start with the first IP address
 - Username: The XIV user created for TPC
 - Password: The password of the XIV
5. Click **Add** and redo these steps for the other IP addresses as illustrated in Figure 7-14.

IP Address	Device Type
9.155.53.250	IBM XIV
9.155.53.251	IBM XIV

Figure 7-14 Configure storage subsystem connections

- After all three IPs are added, click **Next** and the wizard automatically performs a discovery.
- Click **Next** again and the discovered XIV will be displayed as unconfigured storage as shown in Figure 7-15

IP Address	Name	Device ...
9.155.53.251	XIV-2810-6000105-IBM	XIV

Figure 7-15 New XIV discovered in TPC

Probing phase

The newly added XIV storage system must be probed for TPC to collect detailed information. Probes use agents to collect statistics, including data about drives, pools, volumes. The results of the probe jobs are stored in the repository and are used in TPC to supply the data necessary for generating a number of reports, including Asset, Capacity, and Storage Subsystem reports.

To configure the probe for XIV, carry on with the following steps within the Wizard:

- Select the XIV and click **Next**.
- Now you have various options to specify the probe details as shown in Figure 7-16.




Figure 7-16 Specifying probe details

- Use a monitoring group or template: Select the predefined **Monitoring Group** or create your own new group by selecting **Monitoring Template**
- Select monitoring group: Specify **Subsystem Advanced Group**, **Subsystem Basic Group** or **Subsystem Standard Group**

Tip: Configure individual probes for every XIV system, but set them to run at various times.

After clicking **Next**, you get an overview of your configured probe. Figure 7-17 displays the details of the Subsystem Standard Group including the scheduling of the probe and the Alerts.




Figure 7-17 Probe Summary

3. Confirm the summary by clicking **Next**.
4. Click **Finish** and a background probe will be started automatically.
5. Select **IBM Tivoli Storage Productivity Center** → **Configuration Utility** → **Disk Manager** to check the status of your IBM XIV Storage System. You can also use the Disk Manager to run a new Simon Discovery, to create a new probe or to modify the Alerting.




Figure 7-18 Disk Manager

In case you want to modify the probe you just have created, you have to select **IBM Tivoli Storage Productivity Center** → **Monitoring** → **Probes** and choose the group, in our case, **Subsystem Standard Probe**. Click **When to Run** to adjust the schedule or **Alert** in case you want to modify the alerting as illustrated in Figure 7-19.




Figure 7-19 Verifying the probe details

Performance monitoring

The IBM XIV Storage System with software Version 10.2.2 or higher can use the performance monitoring feature in TPC 4.2. and you can set it up with the disk manager, as follows:

1. Click **IBM Tivoli Storage Productivity Center** → **Configuration Utility** → **Disk Manager** → **Storage Subsystem Performance Monitor** → **Create Subsystem Performance Monitor** to create a new performance monitor as shown in Figure 7-20.




Figure 7-20 Subsystem Performance Monitor

- Optionally, you can also directly select **Disk Manager** → **Monitoring** → **Subsystem Performance Monitor** to get into the **Subsystem Performance Monitors** view.
2. Afterwards, select the XIV and click the double arrow icon to move the selected system from the **Available subsystems** into the **Selected subsystems** section.
 3. Click **Save** and specify the subsystem performance monitor name.
 4. Afterwards, the Subsystem performance monitor is automatically submitted.
 5. In case you want to review or change the sampling and scheduling, just select the performance monitor within **Disk Manager** → **Monitoring** → **Subsystem Performance Monitor** and click **Sampling and scheduling** as seen in Figure 7-21.




Figure 7-21 Performance Monitor Details

Configuration information and reporting

In Figure 7-22 you can see a list of various XIV subsystems as reported in TPC.




Figure 7-22 List storage subsystems

XIV Storage System TPC reports

Tivoli Storage Productivity Center version 4.2 includes basic capacity and asset information in tabular reports as well as in Topology Viewer. In addition, LUN correlation information and performance data is available.

TPC probes collect the following information from XIV systems:

- ▶ Storage pools
- ▶ Volumes
- ▶ Disks
- ▶ Ports
- ▶ Host definitions, LUN Mapping & Masking information

Capacity: Space is calculated differently in the XIV Graphical User Interface (GUI) and the Command Line Interface (CLI) than in TPC. XIV defines 1 Gigabyte as $10^9 = 1,000,000,000$ Bytes, while TPC defines 1 Gigabyte as $2^{30} = 1,073,741,824$ Bytes.

This is why capacity information might appear to differ (and seem wrong) when comparing the XIV GUI with the TPC GUI, when in fact it is the exact same size.

Because the IBM XIV Storage System also provides thin provisioning, additional columns for the thin provisioning storage pools were introduced to the TPC GUI.

Note that the TPC terminology of *configured space* accords with XIV's terminology of soft capacity, while TPC terminology of *real space* accords with XIV's terminology of hard space.

Additional *Configured Real Space* and *Available Real Space* columns were introduced to report on the hard capacity of a subsystem, while the pre-existent *Consumed Space* and *Available Space* columns now report on the soft capacity of a subsystem in the following reports:

- ▶ Storage Subsystem list under **Disk Manager** → **Storage Subsystems**
- ▶ Storage Subsystem Details panel under **Disk Manager** → **Storage Subsystems**
- ▶ Storage Subsystem Details panel under **Data Manager** → **Reporting** → **Storage Subsystem**
- ▶ **Data Manager** → **Reporting** → **Asset** → **System-wide** → **Storage Subsystems**
- ▶ **Data Manager** → **Reporting** → **TPC-wide Storage Space** → **Disk Space** → **By Storage Subsystem (Group)**

Configured Real Space and *Available Real Space* columns, reporting on the hard capacity of a storage pool, were also added to the following report:

- ▶ Storage Pool Details panel under **Data Manager** → **Reporting** → **Asset** → **By Storage Subsystem** → **<Subsystem Name>** → **Storage Pools**

See Figure 7-23 for an example of the Storage Pool Details panel.

Detail for Storage Pool	itso
Storage Subsystem	XIV-2810-1300203-IBM
Type	RAID 10
Status	OK
Storage Pool Space	9.11 TB
Available Storage Pool Space	6.55 TB
Configured Real Space	9.11 TB
Available Real Space	8.95 TB
Is Space Efficient	true
Is pool encryptable	false
Is pool encrypted	false
Solid State	Non-Solid State
Number of Disks	N/A
Number of Volumes	18
Surfaced Volume Space	1008.00 GB
Un-surfaced Volume Space	202.00 GB
Backend Volume Space	0

Figure 7-23 Storage Pool Details panel

A *Volume Real Space* column was added to report on the hard capacity of a volume, while the pre-existent *Volume Space* columns report on the soft capacity of a volume in these reports:

- ▶ Volume Details panel under **Disk Manager** → **Storage Subsystems** → **Volumes**
- ▶ **Disk Manager** → **Reporting** → **Storage Subsystems** → **Volumes**
- ▶ **Disk Manager** → **Reporting** → **Storage Subsystems** → **Volume to HBA Assignment**
- ▶ Added **Backend Volume Real Space** for XIV volumes as backend volumes under **Disk Manager** → **Reporting** → **Storage Subsystems** → **Volume to Backend Volume Assignment**
- ▶ Volume Details panel under **Data Manager** → **Reporting** → **Asset** → **By Storage Subsystem** → **<Subsystem Name>** → **Volumes**
- ▶ **Data Manager** → **Reporting** → **Asset** → **System-wide** → **Volumes**

See Figure 7-24 for an example of the Volume Details panel.

Volume	Host Paths
Storage Subsystem	XIV-2810-1300203-IBM
Volume Space	22.00 GB
Volume Real Space	14.90 GB
Unavailable Volume Space	N/A
Type	Raw
RAID Level	RAID 10
Flash Copy	None
Attributes	Surfaced
"tpc_thin_vol1 (ID:05A8)"	

Figure 7-24 Volume Details panel

Due to the XIV architecture and the fact that each volume resides on all disks, various reports in the TPC GUI will not necessarily provide meaningful information for IBM XIV Storage Systems. Correlation of disks and volumes, for example under the **Data Manager → Reporting → Asset → By Storage Subsystem → Select a Storage Subsystem → Disks** branch, is not possible - TPC will not report any volumes under the branch of a particular disk.

Also, because XIV storage pools are used to group volumes but not disks, no disks will be reported for a particular storage pool under the reporting branch mentioned before.

Finally, the following reports will not contain any information for XIV Storage Subsystems:

- ▶ **Disk Manager → Reporting → Storage Subsystems → Computer Views → By Computer** (Relate Computers to Disks)
- ▶ **Disk Manager → Reporting → Storage Subsystems → Computer Views → By Computer Group** (Relate Computers to Disks)
- ▶ **Disk Manager → Reporting → Storage Subsystems → Computer Views → By Filesystem/Logical Volume** (Relate Filesystems/Logical Volumes to Disks)
- ▶ **Disk Manager → Reporting → Storage Subsystems → Computer Views → By Filesystem Group** (Relate Filesystems/Logical Volumes to Disks)
- ▶ **Disk Manager → Reporting → Storage Subsystems → Storage Subsystem Views → Disks** (Relate Disks to Computers)

Figure 7-25 illustrates how TPC can report on XIV storage pools.

Navigation Tree		Selection Storage Subsystems						
- Administrative Services - Services - Data Sources - Discovery - Configuration - IBM Tivoli Storage Productivity Center - Data Manager - Data Manager for Databases - Data Manager for Chargeback - Disk Manager - Storage Subsystems - Storage Optimizer - SAN Planner - Monitoring - Alerting - Profile Management - Reporting - Groups - Storage Subsystems - Disks - Volumes - Storage Pools - By Storage Subsystem - By Storage Subsystem Group		Storage Pools: By Storage Subsystem Number of Rows: 19						
		Storage Subsystem	Storage Pool	Type(s)	Track Format	Status	Storage Pool Space	Available Storage Pool Space
		XIV-2810-1300202-IBM	test_pool	RAID 10	Unknown	ok	45.48 TB	17.97
		XIV-2810-1300209-IBM	bfs_regular	RAID 10	Unknown	ok	944.00 GB	607.22
		XIV-2810-1300209-IBM	bfs_source	RAID 10	Unknown	ok	96.00 GB	47.50
		XIV-2810-1300209-IBM	bfs_target	RAID 10	Unknown	ok	96.00 GB	63.33
		XIV-2810-1300209-IBM	bfs_thin	RAID 10	Unknown	ok	1.83 TB	1.50
		XIV-2810-1300209-IBM	ESX_pool	RAID 10	Unknown	ok	3.64 TB	3.27
		XIV-2810-1300209-IBM	ITSQ_SVC	RAID 10	Unknown	ok	7.83 TB	79.16
		XIV-2810-1300209-IBM	ITSQ_test	RAID 10	Unknown	ok	64.00 GB	0
		XIV-2810-1300209-IBM	mirror_pool	RAID 10	Unknown	ok	288.00 GB	127.59
		XIV-2810-1300209-IBM	WindowsThinPool	RAID 10	Unknown	ok	1.38 TB	207.68
		XIV-2810-1300209-IBM	xivaix01_pool	RAID 10	Unknown	ok	9.19 TB	2.56
		XIV-2810-1300774-IBM	celI0002	RAID 10	Unknown	ok	384.00 GB	47.50
		XIV-2810-1300774-IBM	ITSQ_ESX	RAID 10	Unknown	ok	944.00 GB	879.17
		XIV-2810-1300774-IBM	ITSQ_Windows	RAID 10	Unknown	ok	9.22 TB	8.11
		XIV-2810-1300774-IBM	OSL_POOL_0	RAID 10	Unknown	ok	15.47 TB	15.45
		XIV-2810-1300774-IBM	OSL_POOL_THIN	RAID 10	Unknown	ok	30.94 TB	15.47
		XIV-2810-MN00019-IBM	test_pool	RAID 10	Unknown	ok	70.95 TB	21.72
		XIV-2810-MN00033-IBM	test_pool	RAID 10	Unknown	ok	14.75 TB	3.58
		XIV-2810-MN00035-IBM	test_pool	RAID 10	Unknown	ok	71.67 TB	22.94

Figure 7-25 XIV storage pools as seen in TPC

These queries, when combined with SCSI inquiry data TPC collects from the hosts, allow TPC to correlate LUNs reported by the IBM XIV Storage System to LUNs as seen by the host systems.

Also, when the IBM XIV Storage System is providing storage to the IBM System Storage SAN Volume controller (SVC), TPC can correlate LUNs reported by the IBM XIV Storage System to SVC managed disks (MDisks).

Element Manager Launch

If the XIV GUI is installed and configured, TPC provides the ability to launch the XIV management software.

TPC Performance Reports

The IBM XIV Storage System with software level 10.2.2. or higher enables TPC to create performance reports. Obviously, the corresponding XIV performance monitors need to be set up first, as already described in “Performance monitoring” on page 221.

To create a performance report, select **IBM Tivoli Storage Productivity Center** → **Configuration Utility** → **Disk Manager** → **Storage Subsystem Performance Management** → **Show Subsystems Performance System Report** as shown in Figure 7-26.




Figure 7-26 Show Performance Report

Afterwards you can drill down into the details of any storage subsystem or even deeper into the single volume view and create TPC reports or graphs.

The screenshot shows the 'Storage Subsystem Performance: By Volume' view. It includes a table with 5 columns: Subsystem, Volume, Time, Interval, and Read I/O. The table lists 16 rows of performance data for various subsystems and volumes, all dated Oct 5, 2010 6:00:02 PM. The left pane shows the navigation tree with 'Subsystem Performance' selected under 'Reporting' > 'System Reports' > 'Data' > 'Disk'.

Subsystem	Volume	Time	Interval	Read I/O
XIV-2810-6000105-IBM	*gpfs_1(ID:14093)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*gpfs_2(ID:14094)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*gpfs_3(ID:14095)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*gpfs_4(ID:14941)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*itso_aix_1(ID:7342)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*itso_aix_2(ID:7343)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*N5600A_Data_1(ID:15079)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*N5600A_Data_2(ID:15080)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*N5600A_Data_3(ID:15081)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*N5600B_data_m_1(ID:15082)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*N5600B_data_m_2(ID:15083)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*N5600B_data_m_3(ID:15084)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*Onstor_Core_1(ID:14034)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*Onstor_Core_2(ID:14035)	Oct 5, 2010 6:00:02 PM	300 s	
XIV-2810-6000105-IBM	*ONStor_Lun_1(ID:13790)	Oct 5, 2010 6:00:02 PM	300 s	

Figure 7-27 Subsystem Performance

TPC Volume Creation

After the **Configure Devices** tool was used to add an IBM XIV Storage System into TPC as described in Figure 7-13 on page 217, the new NAPI of TPC v.4.2 is used as native interface between TPC and the XIV. This interface also enables TPC directly to create a volume on the XIV and to map it to a host by using the **Volume Create Wizard**.

Simply choose the XIV you want to create a new volume by selecting **Disk Manager** → **Storage Subsystem** and click **Create Volumes** to start the **Volume Create Wizard**.

7.2 Performance monitoring

During normal operation, the IBM XIV Storage System constantly gathers statistical information. The data can then be processed using the GUI or XIV Command Line Interface (XCLI). This section introduces the techniques for processing the statistics data.

7.2.1 Using the GUI

The GUI provides a mechanism to gather statistics data. For a description of setting up and using the GUI, see Chapter 4, “Configuration” on page 83. When working with the statistical information, the IBM XIV Storage System collects and maintains the information internally. As the data ages, it is consolidated to save space. By selecting specific filters, the requested data is mined and displayed. This section discusses the functionality of the GUI and how to retrieve the required data.

The first item to note is the current IOPS for the system is always displayed in the bottom center of the window. This feature provides simple access to the current load on the system. Figure 7-28 illustrates the GUI and the IOPS display; it also shows how to start the statistics monitor.




Figure 7-28 Starting the statistics monitor on the GUI

Select **Statistics** from the Monitor menu as shown in Figure 7-28 to display the Monitor default view that is shown in Figure 7-29.

Figure 7-29 shows the system IOPS:

- The X-axis of the graph represents the time and can vary from minutes to months.
- The Y-axis of the graph is the measurement selected. The default measurement is IOPS.

The statistics monitor also illustrates latency and bandwidth.



Figure 7-29 Default statistics monitor view

The other options in the statistics monitor act as filters for separating data. These filters are separated by the type of transaction (reads or writes), cache properties (hits compared to misses), or the transfer size of I/O as seen by the IBM XIV Storage System. See Figure 7-30 for a better view of the filter pane.




Figure 7-30 Filter pane for the statistics monitor

The filter pane allows you to select multiple items within a specific filter, for example, if you want to see reads and writes separated on the graph. By holding down Ctrl on the keyboard and selecting the read option and then the write option, you can witness both items displayed on the graph.

As shown in Figure 7-31, one of the lines represents the reads and the other line represents the writes. On the GUI, these lines are drawn in separate colors to differentiate the metrics.

This selection process can be performed on the other filter items as well.



Figure 7-31 Multiple filter selection

In certain cases, the user needs to see multiple graphs at one time. On the right side of the filter pane, there is a selection to add graphs (see Figure 7-30 on page 227). Up to four graphs are managed by the GUI. Each graph is independent and can have separate filters.

Figure 7-32 illustrates this multi graph concept. The top graph is the IOPS for the day with the reads and writes separated. The second graph displays the bandwidth for a few minutes with reads and writes separated, which provides quick and easy access to multiple views of the performance metrics.



Figure 7-32 Multiple graphs using the GUI

There are various additional filters available, such as filtering by host, volumes, interfaces, or targets. These items are defined on the left side of the filter pane. When clicking one of these filters, a dialog window displays. Highlight the item, or select a maximum of four using the Ctrl key, to be filtered and then click **Click to select**. It moves the highlighted item to the lower half of the dialog box. In order to generate the graph, you must click the green check mark located on the lower right side of the dialog box. Your new graph is generated with the name of the filter at the top of the graph. See Figure 7-33 for an example of this filter.



Figure 7-33 Example of a host filter

On the left side of the chart in the blue bar, there are various tools to assist you in managing the data.

Figure 7-34 shows the chart toolbar in more detail.



Figure 7-34 Chart toolbar

The top two tools (magnifying glasses) zoom in and out for the chart, and the second set of two tools adjusts the X-axis and the Y-axis for the chart. Finally, the bottom two tools allow you to export the data to a comma-separated file or print the chart to a printer.

7.2.2 Using XIVTop

In addition to the XIV GUI Performance Statistics options, you can also use XIVTop for additional and more granular system performance statistics. The XIVTop application allows you to view and monitor performance information for defined volumes and hosts in real-time.

There are multiple ways to access the XIVTop utility.

To start XIVTop directly from the GUI desktop, right-click within the “All systems” view on the XIV system that you want to monitor.




Figure 7-35 Starting XIVTop from the main GUI desktop

Launch XIVTop from the XIVGUI desktop.




Figure 7-36 Launch XIVTop from the GUI Desktop

XIVTop can also be started within the statistic view from the XIV GUI Statistics Monitor Filter pane (see Figure 7-30 on page 227), or by clicking the icon as shown here in Figure 7-37.




Figure 7-37 Starting XIVTop from the statistics view

The XIVTop tool (Figure 7-38) was designed to quickly display the volumes or hosts that are consuming the most system resources at any given time. Information can be sorted in a variety of ways and exported to a file for use in spreadsheet applications.




Figure 7-38 XIVTop Monitoring Tool

The initial XIVTop screen is divided into two horizontal sections, the Volumes and Hosts section and the Performance Chart section. You can collapse or expand each section by clicking the Arrow icon on the right side of the header of each section (Figure 7-39).




Figure 7-39 XIVTop with Performance Chart Section collapsed

Each column of the Volumes & Hosts section can be sorted in ascending or descending order by clicking the column and toggle the direction of the sort indicator (Figure 7-40).




Figure 7-40 Sort Columns in ascending or descending order

The IOPS, Latency or BandWidth for up to four Volumes or 4 Hosts can be viewed in the Performance Chart section. The controls to select the Volumes, Hosts and Options operate in much the same way as the Statistics tool from within the GUI (Figure 7-41).




Figure 7-41 Display Performance information for up to 4 Volumes or Hosts with XIVTop

7.2.3 Using the XCLI

The second method to collect statistics is through the XCLI operation. In order to access the XCLI session, see Chapter 4, “Configuration” on page 83.

To retrieve the system time, issue the **time_list** command, and the system retrieves the current time. See Example 7-19 for an example of retrieving the IBM XIV Storage System time.

Example 7-19 Retrieving the XIV Storage System time

```
>> time_list
Time      Date      Time Zone  Daylight Saving Time
11:45:42   2010-09-27  GMT       no
```

After the system time is obtained, the **statistics_get** command can be formatted and issued.

The **statistics_get** command requires various parameters to operate. The command requires that you enter a starting or ending time point, a count for the number of intervals to collect, the size of the interval, and the units related to that size. The TimeStamp is modified from the previous **time_list** command. Example 7-20 provides a description of the command.

Example 7-20 The statistics_get command format

```
statistics_get [ host=H | host_iscsi_name=initiatorName | host_fc_port=WWPN |
target=RemoteTarget | remote_fc_port=WWPN | remote_ipaddress=IPAddress |
vol=VolName | ipinterface=IPInterfaceName | local_fc_port=ComponentId ]
< start=TimeStamp | end=TimeStamp > [ module=ComponentId ] count=N
interval=IntervalSize resolution_unit=<minute|hour|day|week|month>
```

To further explain this command, assume that you want to collect 10 intervals, and each interval is for one minute. The point of interest occurred June 16 2009 roughly 15 minutes 11:45:00. It is important to note the **statistics_get** command allows you to gather the performance data from any time period.

The time stamp is formatted as *YYYY-MM-DD:hh:mm:ss*, where the *YYYY* represents a four digit year, *MM* is the two digit month, and *DD* is the two digit day. After the date portion of the time stamp is specified, you specify the time, where *hh* is the hour, *mm* is the minute, and *ss* represents the seconds.

Example 7-21 shows a typical use of this command, and Figure 7-42 shows sample output of the statistics. The output displayed is a small portion of the data provided.

Example 7-21 The statistics_get command example

```
>> statistics_get end=2009-06-16.11:45:00 count=10 interval=1
resolution_unit=minute
```

Time	Read Hit Medium - IOps	Read Hit Medium - Latency	Read Hit Medium - Throughput
2009-06-16 11:35:00	3418	2851	111332
2009-06-16 11:36:00	3583	2769	119298
2009-06-16 11:37:00	3673	2797	121695
2009-06-16 11:38:00	3430	3121	113067
2009-06-16 11:39:00	3548	3108	117711
2009-06-16 11:40:00	3274	2933	109533
2009-06-16 11:41:00	2917	3270	96250
2009-06-16 11:42:00	2651	3262	87067
2009-06-16 11:43:00	2292	3168	74384
2009-06-16 11:44:00	2045	3449	65100

Figure 7-42 Output from `statistics_get` command

Extending this example, assume that you want to filter out a specific host defined in the XIV Storage System. By using the host filter in the command, you can specify for which host you want to see performance metrics, which allows you to refine the data that you are analyzing.

See Example 7-22 for an example of how to perform this operation, and see Figure 7-43 for a sample of the output for the command.

Example 7-22 The `statistics_get` command using the host filter

```
>> statistics_get host=adams end=2009-06-16.11:45:00 count=10 interval=1
resolution_unit=minute
```

Time	Read Hit Medium - IOps	Read Hit Medium - Latency	Read Hit Medium - Throughput
2009-06-16 11:35:00	149	2092	4489
2009-06-16 11:36:00	136	2689	4219
2009-06-16 11:37:00	141	3016	3897
2009-06-16 11:38:00	293	2927	8068
2009-06-16 11:39:00	418	3030	12574
2009-06-16 11:40:00	391	3101	12518
2009-06-16 11:41:00	445	3554	13858
2009-06-16 11:42:00	518	3536	15748
2009-06-16 11:43:00	490	3243	14352
2009-06-16 11:44:00	370	3628	11531

Figure 7-43 Output from the `statistics_get` command using the host filter

In addition to the filter just shown, the **`statistics_get` command** is capable of filtering iSCSI names, host worldwide port names (WWPNs), volume names, modules, and many more fields. As an additional example, assume you want to see the workload on the system for a specific module. The module filter breaks out the performance on the specified module. Example 7-23 pulls the performance statistics for module 5 during the same time period of the previous examples. Figure 7-44 shows the output.

Example 7-23 The `statistics_get` command using the module filter

```
>> statistics_get end=2009-06-16.11:45:00 module=5 count=10 interval=1
resolution_unit=minute
```

Time	Read Hit Medium - IOps	Read Hit Medium - Latency	Read Hit Medium - Throughput
2009-06-16 11:35:00	354	894	9980
2009-06-16 11:36:00	165	831	4485
2009-06-16 11:37:00	159	813	4194
2009-06-16 11:38:00	159	846	4166
2009-06-16 11:39:00	127	852	3374
2009-06-16 11:40:00	98	848	2493
2009-06-16 11:41:00	115	856	3080
2009-06-16 11:42:00	115	842	3124
2009-06-16 11:43:00	48	811	1262
2009-06-16 11:44:00	192	896	5036

Figure 7-44 Output from `statistics_get` command using the module filter

7.3 XIV event notification

The IBM XIV Storage System allows you to send alerts by email and SMS messages.

7.3.1 Configuring the XIV system

You can configure the XIV system using very flexible rules to ensure that notification is sent to the correct person, or group of people, according to the various parameters. This event notification is similar to, but not quite the same as XIV Call Home, discussed in 7.5, “Call Home and Remote support”

Setting up event notification

Configuration options are available from the XIV GUI. You have the flexibility to create a detailed events notification plan based on specific rules. This flexibility allows the storage administrator to decide, for instance, where to direct alerts for various event types. All these settings can also be done with XCLI commands.

Setup notification and rules with the GUI

To set up email or SMS notification and rules, proceed as follows:

1. From the XIV GUI main window, select the Monitor icon. From the Monitor menu, select **Events** to display the Events window as shown in Figure 7-45.




Figure 7-45 Setup notification and rules

2. From the toolbar, click **Setup** to invoke the Events Configuration wizard. The wizard will guide you through the configuration of Gateways, Destinations, and Rules.

The wizard Welcome window is shown in Figure 7-46.




Figure 7-46 Events Configuration wizard

Gateways

The wizard will first take you through the configuration of the gateways:

1. Click **Next** or **Gateway** to display the Events Configuration - Gateway dialog as shown in Figure 7-47.



Figure 7-47 Define Gateway panel

2. Click **Define Gateway**. The Gateway Create Welcome panel shown in Figure 7-48 displays. Click **Next**.



Figure 7-48 Configure Gateways

The Gateway Create - Select gateway type panel displays as shown in Example 7-49.




Figure 7-49 Select gateway type

3. When the wizard prompts for the type of the gateway, either SMTP for email notification or SMS if an alert or information will initiate an SMS, click either **SMTP** or **SMS**.

Gateways: An SMTP gateway needs to be defined before you can define an SMS gateway because SMS messages are sent from the IBM XIV Storage System in the form of an email.

The next steps differ for SMTP and SMS. Our illustration from now on is for SMTP. However, the steps to go through for SMS are similarly self-explanatory and are detailed in the next section, *Setting up notification and rules with the XCLI*.

4. To proceed using **SMTP**, click **Next**.
5. Enter the gateway name of the SMTP Gateway; click **Next**. Enter the IP address or DNS name of the SMTP gateway for the gateway address; click **Next**. The SMTP Sender Email address panel as shown in Figure 7-50 displays.

Figure 7-50 SMTP Sender Email Address

6. Set the sender email address. You can use the default, or enter a new address. In case of email problems, such as the wrong email address, a response email will be sent to this address. Depending on how your email server is configured, you might need to use an authorized address in order to ensure proper delivery of notifications. Click **Finish**.

The Create Gateway summary panel is displayed as shown in Figure 7-51.

Figure 7-51 Create the Gateway - Summary

7. Review the information you entered on this panel. If all is correct, click **Create**. If not, click **Back** until you are at the information that needs to be changed, or just click one of the buttons on the left, to take you directly to the information that needs to be changed.

Destinations

Next, the Events Configuration wizard will guide you through the setup of the destinations where you configure email addresses or SMS receivers. Figure 7-52 shows the Destination panel of the Events Configuration Wizard.




Figure 7-52 Add Destination

1. Click **Create Destination** to display the Welcome Panel, as shown in Figure 7-53.




Figure 7-53 Destination Create

2. Click **Next** to proceed. The Select Destination type panel, shown in Figure 7-54, is displayed.

3. On this panel, you configure the following settings:
 - Type: Event notification destination type can be either a destination group (containing other destinations), SNMP manager for sending SNMP traps, email address for sending email notification, or mobile phone number for SMS notification:
 - SNMP
 - EMAIL
 - SMS
 - Group of Destinations




Figure 7-54 Select destination type

- Depending on the selected type, the remaining configuration information required differs but is self-explanatory.

Rules

The final step in the Events Creation Wizard is creating a rule. A rule determines what notification is sent. It is based on event severity, event code or both:

1. Click **Create Rule** as shown in Figure 7-55.




Figure 7-55 Create Rule

2. On the Welcome panel, click **Next**. The Rule Create - Rule name panel shown in Figure 7-56 is displayed.



Figure 7-56 Rule name

3. To define a rule, configure the following settings:
 - Rule Name: Enter a name for the new rule. Names are case-sensitive and can contain letters, digits, or the underscore character (_). You cannot use the name of an already defined rule.
 - Rule condition setting: Select **Severity only** if you want the rule to be triggered by severity, select **Event Code only** if you want the rule be triggered by event, or both severity and event code for events that might have multiple severities depending on a threshold of certain parameters:
 - Severity only
 - Event Code only
 - Both severity and event code
 - Select the severity trigger: Select the minimum severity to trigger the rule's activation. Events of this severity or higher will trigger the defined rule.
 - Select the event code trigger: Select the event code to trigger the rule's activation from the drop down list. The list of event codes is pre-defined and can not be modified. See the *XCLI User Manual*, GC27-2213 for a list of all the available event codes and descriptions to chose from.
 - Rule destinations: Select destinations and destination groups to be notified when the event's condition occurs. Here, you can select one or more existing destinations or also define a new destination (see Figure 7-57).



Figure 7-57 Select destination

- Rule snooze: Defines whether the system repeatedly alerts the defined destination until the event is cleared. If so, a snooze time must be selected. Either:
 - Check **Use snooze timer**
 - Snooze time in minutes
 - Rule escalation: Allows the system to send alerts by other rules if the event is not cleared within a certain time. If so, an escalation time and rule must be specified:
 - Check **Use escalation rule**
 - Escalation Rule
 - Escalation time in minutes
 - Create Escalation Rule
4. On the summary panel shown in Figure 7-58, review the information you entered. Go back if you need to make changes, or if all is correct, click **Create**.



Figure 7-58 Rule Create

Setting up notification and rules with the XCLI

You use the same process to set up the IBM XIV Storage System for notification with the XCLI as you used with the GUI. The three-step process includes all the required configurations to allow the IBM XIV Storage System to provide notification of events:

- ▶ Gateway
- ▶ Destination
- ▶ Rules

The gateway definition is used for SMTP and SMS messages. There are various commands used to create and manage the gateways for the IBM XIV Storage System. Example 7-24 shows an SMTP gateway being defined. The gateway is named *test* and the messages from the IBM XIV Storage System are addressed to `xiv@us.ibm.com`.

When added, the existing gateways are listed for confirmation. In addition to gateway address and sender address, the port and reply to address can also be specified. There are various other commands that are available for managing a gateway.

Example 7-24 The `smtpgw_define` command

```
>> smtpgw_define smtpgw=test address=test.ibm.com from_address=xiv@us.ibm.com
Command executed successfully.
```

```
>> smtpgw_list
```

Name	Address	Priority
ITSO Mail Gateway	us.ibm.com	1
test	test.ibm.com	2

The SMS gateway is defined in a similar method. The difference is that the fields can use tokens to create variable text instead of static text. When specifying the address to send the SMS message, tokens can be used instead of hard-coded values. In addition, the message body also uses a token to have the error message sent instead of a hard-coded text.

Gateways: An SMTP gateway needs to be defined before you can define an SMS gateway because SMS messages are sent from the IBM XIV Storage System in the form of an email.

Example 7-25 provides an example of defining a SMS gateway. The tokens available to be used for the SMS gateway definition are:

- ▶ `{areacode}`: This escape sequence is replaced by the destination's mobile or cellular phone number area code.
- ▶ `{number}`: This escape sequence is replaced by the destination's cellular local number.
- ▶ `{message}`: This escape sequence is replaced by the text to be shown to the user.
- ▶ `\{`, `\}`, `\\`: These symbols are replaced by the `{`, `}` or `\` respectively.

Example 7-25 The `smsgw_define` command

```
>> smsgw_define smsgw=test email_address={areacode}{number}@smstest.ibm.com
subject_line="XIV System Event Notification" email_body={message}
Command executed successfully.
```

```
>> smsgw_list
```

Name	Email Address	SMTP Gateways
test	{areacode}{number}@smstest.ibm.com	all

When the gateways are defined, the destination settings can be defined. There are three types of destinations:

- ▶ SMTP or email
- ▶ SMS
- ▶ SNMP

Example 7-26 provides an example of creating a destination for all three types of notifications. For the email notification, the destination receives a test message every Monday at 12:00. Each destination can be set to receive notifications on multiple days of the week at multiple times.

Example 7-26 Destination definitions

```
>> dest_define dest=emailtest type=EMAIL email_address=test@ibm.com smtpgws=ALL
heartbeat_test_hour=12:00 heartbeat_test_days=Mon
Command executed successfully.
```

```
>> dest_define dest=smstest type=SMS area_code=555 number=5555555 msgws=ALL
Command executed successfully.
```

```
>> dest_define dest=snmpptest type=SNMP snmp_manager=9.9.9.9
Command executed successfully.
```

```
>> dest_list
```

Name	Type	Email Address	Area Code	Phone Number	SNMP Manager	User
ITSO_Catcher	SNMP					itsocatcher.us.ibm.com
smstest	SMS		555	5555555		
snmpptest	SNMP				9.9.9.9	
emailtest	EMAIL	test@ibm.com				

Finally, the rules can be set for which messages can be sent. Example 7-27 provides two examples of setting up rules. The first rule is for SNMP and email messages and all messages, even informational messages, are sent to the processing servers. The second example creates a rule for SMS messages. Only critical messages are sent to the SMS server, and they are sent every 15 minutes until the error condition is cleared.

Example 7-27 Rule definitions

```
>> rule_create rule=emailtest min_severity=informational dests=emailtest,snmpptest
Command executed successfully.
```

```
>>rule_create rule=smstest min_severity=critical dests=smstest snooze_time=15
Command executed successfully.
```

```
>> rule_list
```

Name	Minimum Severity	Event Codes	Except Codes	Destinations	Active	Escalation Only
ITSO_Major	Major	all		ITSO_Catcher	yes	no
emailtest	Informational	all		emailtest,snmpptest	yes	no
smstest	Critical	all		smstest	yes	no

Example 7-28 shows illustrations of deleting rules, destinations, and gateways. It is not possible to delete a destination if a rule is using that destination, and it is not possible to delete a gateway if a destination is pointing to that gateway.

Example 7-28 Deletion of notification setup

```
>> rule_delete -y rule=smstest
Command executed successfully.
>> dest_delete -y dest=smstest
Command executed successfully.
>> msgw_delete -y msgw=test
Command executed successfully.
```

7.3.2 SNMP-based monitoring

So far, we have discussed how to perform monitoring based on the XIV GUI and the XCLI. The IBM XIV Storage System also supports Simple Network Management Protocol (SNMP) for monitoring. SNMP-based monitoring tools, such as IBM Tivoli NetView® or the IBM Director, can be used to monitor the IBM XIV Storage System.

Simple Network Management Protocol (SNMP)

SNMP is an industry-standard set of functions for monitoring and managing TCP/IP-based networks and systems. SNMP includes a protocol, a database specification, and a set of data objects. A set of data objects forms a Management Information Base (MIB).

The SNMP protocol defines two terms, *agent* and *manager*, instead of the client and server terms that are used in many other TCP/IP protocols.

SNMP agent

An SNMP agent is a daemon process that provides access to the MIB objects on IP hosts on which the agent is running. An SNMP agent, or daemon, is implemented in the IBM XIV software and provides access to the MIB objects defined in the system. The SNMP daemon can send SNMP trap requests to SNMP managers to indicate that a particular condition exists on the agent system, such as the occurrence of an error.

SNMP manager

An SNMP manager can be implemented in two ways. An SNMP manager can be implemented as a simple command tool that can collect information from SNMP agents. An SNMP manager also can be composed of multiple daemon processes and database applications. This type of complex SNMP manager provides you with monitoring functions using SNMP. It typically has a graphical user interface for operators. The SNMP manager gathers information from SNMP agents and accepts trap requests sent by SNMP agents. In addition, the SNMP manager generates traps when it detects status changes or other unusual conditions while polling network objects. IBM Director is an example of an SNMP manager with a GUI interface.

SNMP trap

A *trap* is a message sent from an SNMP agent to an SNMP manager without a specific request from the SNMP manager. SNMP defines six generic types of traps and allows you to define enterprise-specific traps. The trap structure conveys the following information to the SNMP manager:

- ▶ Agent's object that was affected
- ▶ IP address of the agent that sent the trap

- ▶ Event description (either a generic trap or enterprise-specific trap, the including trap number)
- ▶ Time stamp
- ▶ Optional enterprise-specific trap identification
- ▶ List of variables describing the trap

SNMP communication

The SNMP manager sends SNMP get, get-next, or set requests to SNMP agents, which listen on UDP port 161, and the agents send back a reply to the manager. The SNMP agent can be implemented on any kind of IP host, such as UNIX workstations, routers, network appliances, and also on the IBM XIV Storage System.

You can gather various information about the specific IP hosts by sending the SNMP get and get-next requests, and you can update the configuration of IP hosts by sending the SNMP set request.

The SNMP agent can send SNMP trap requests to SNMP managers, which listen on UDP port 162. The SNMP trap requests sent from SNMP agents can be used to send warning, alert, or error notification messages to SNMP managers. Figure 7-59 illustrates the characteristics of SNMP architecture and communication.




Figure 7-59 SNMP communication

You can configure an SNMP agent to send SNMP trap requests to multiple SNMP managers.

Management Information Base (MIB)

The objects, which you can get or set by sending SNMP get or set requests, are defined as a set of databases called the *Management Information Base* (MIB). The structure of MIB is defined as an Internet standard in RFC 1155. The MIB forms a tree structure.

Most hardware and software vendors provide you with extended MIB objects to support their own requirements. The SNMP standards allow this extension by using the private sub-tree, which is called an enterprise-specific MIB. Because each vendor has a unique MIB sub-tree under the private sub-tree, there is no conflict among vendors' original MIB extensions.

The IBM XIV Storage System comes with its own specific MIB. Contact your IBM support representative and they will provide the file to you based on your systems software level.

XIV SNMP setup: Configuring a new destination

To effectively use SNMP monitoring with the IBM XIV Storage System, you must first set it up to send SNMP traps to an SNMP manager (such as the IBM Director server), which is defined in your environment. Figure 7-60 illustrates where to start to set up the SNMP destination. Also, see “Setup notification and rules with the GUI” on page 237.




Figure 7-60 Configure destination

Follow these steps to configure a new SNMP destination:

1. From the XIV GUI main window, select the Monitor icon.
2. From the Monitor menu, select **Events** to display the Events window as shown in Figure 7-60.

From the toolbar:

- a. Click **Destinations**. The Destinations dialog window opens.
- b. Select **SNMP** from the Destinations pull-down list.
- c. Click the green plus sign (+) and select **Destination** from the pop-up menu to add a destination as illustrated in Figure 7-61.




Figure 7-61 Add SNMP destination

3. From the Define Destination dialog, which is now open, enter a Destination Name (a unique name of your choice) and the IP or Domain Name System (DNS) of the server where the SNMP Management software is installed. See Figure 7-62.




Figure 7-62 Define SNMP destination

4. Click **Define** to effectively add the SNMP Manager as a destination for SNMP traps.

Also take care to set up the rules for the defined SNMP destination as described within the section “Rules” on page 242. Afterwards the IBM XIV Storage System is set up to send SNMP Traps to the defined SNMP manager. The SNMP Manager software will process the received information (SNMP traps) according to the MIB file.

Using IBM Director

In this section, we illustrate how to use IBM Director to monitor the IBM XIV Storage System. The IBM Director is an example of a possible SNMP manager for XIV. Other SNMP Managers can be used with XIV as well.

Attention: If you are already using IBM Systems Director, see “Using IBM Systems Director Version 6” on page 256.

IBM Director provides an integrated suite of software tools for a consistent, single point of management and automation. With IBM Director, IT administrators can view and track the hardware configuration of remote systems in detail and monitor the usage and performance of critical components, such as processors, disks, and memory.

All IBM clients can download the latest version of IBM Director code from the IBM Director Software Download Matrix page:

<http://www.ibm.com/systems/management/director/downloads.html>

For detailed information regarding the installation, setup, and configuration of IBM Director, see the documentation available at:

<http://www.ibm.com/systems/management/director/>

Compile MIB file

After you have completed the installation of your IBM Director environment, you can prepare it to manage the IBM XIV Storage System by compiling the provided MIB file.

Make sure to always use the latest MIB file provided. To compile the MIB file in your environment, proceed as follows:

1. At the IBM Director Console window, click **Tasks** → **SNMP Browser** → **Manage MIBs**, as shown in Figure 7-63.




Figure 7-63 Manage MIBs

2. In the MIB Management window, click **File** → **Select MIB to Compile**.
3. In the Select MIB to Compile window that is shown in Figure 7-64, specify the directory and file name of the MIB file that you want to compile, and click **OK**. A status window indicates the progress.




Figure 7-64 Compile MIB

When you compile a new MIB file, it is also automatically loaded in the *Loaded MIBs file directory* and is ready for use.

To load an already compiled MIB file:

1. In the MIB Management window, click **File** → **Select MIB** to load.
2. Select the MIB (that you to load) in the Available MIBs window, click **Add**, **Apply**, and **OK**.

This action will load the selected MIB file, and the IBM Director is ready to be configured for monitoring the IBM XIV Storage System.

Discover the IBM XIV Storage System

After loading the MIB file into the IBM Director, the next step is to discover the IBM XIV Storage Systems in your environment. Therefore, configure the IBM Director for auto-discover as follows:

1. From the IBM Director Console window, select **Options** → **Discovery Preferences**, as shown in Figure 7-65.




Figure 7-65 Discovery Preferences

2. In the Discovery Preferences window that is shown in Figure 7-66, follow these steps to discover IBM XIV Storage Systems:
 - a. Click the **Level 0: Agentless System** tab.
 - b. Click **Add** to bring up a window to specify whether you want to add a single address or an address range. Select **Unicast Range**.

Unicast: Because each XIV system is presented through three IP addresses, select Unicast Range when configuring the auto-discovery preferences.

- c. Next, enter the address range for the IBM XIV Storage Systems in your environment. You also set the Auto-discover period and the Presence Check period.




Figure 7-66 Discover Range

3. After you have set up the Discovery Preferences, the IBM Director will discover the IBM XIV Storage Systems and add them to the IBM Director Console as seen in Figure 7-67.




Figure 7-67 IBM Director Console

At this point, the IBM Director and IBM Director Console are ready to receive SNMP traps from the discovered IBM XIV Storage Systems.

With the IBM Director, you can display general information about your IBM XIVs, monitor the Event Log, and browse more information.

General System Attributes

Double-click the entry corresponding to your IBM XIV Storage System in the IBM Director Console window to display the General System Attributes as illustrated in Figure 7-68. This window gives you a general overview of the system status.

The screenshot shows the IBM Director Console interface. The top window, titled "All Managed Objects : SNMP System Object ID", displays a list of managed objects. The selected object is "XIV ESP 1 V10.0 1300203". Below this, a second window titled "Display System Attributes : XIV ESP 1 V10.0 1300203" is open, showing the "General Attributes" tab. This tab displays a table of system information.

System Information	
System Name	XIV ESP 1 V10.0 1300203
System Factory ID	SNMP Devices
System State	Online
System Presence Check Setting (minutes)	15
Secure/Unsecure supported	false
Access Denied	false
Encryption Enabled	false
TCP/IP Addresses	{ '9.155.53.252' }
TCP/IP Hosts	{ 'dyn-9-155-53-252.mainz.de.ibm.com' }
Operating System	Linux
OS Major Version	2
OS Minor Version	6
System Name (MIB2)	XIV ESP 1 V10.0 1300203
System Description (MIB2)	Linux nextra-1300203-module-6 2.6.16.46-253-xiv-154-x86_64-nextra #1 SMP Tue Jul 15 02:43:02 UTC 2008 x86_64
System Contact (MIB2)	Unknown
System Location (MIB2)	Unknown
System Object ID (MIB2)	enterprises.8072.3.2.10 (1.3.6.1.4.1.8072.3.2.10)
System Uptime (MIB2)	5 hours 2 minutes 6 seconds (1812669)

Figure 7-68 General System Attributes

Event Log

To open the Event Log, right-click the entry corresponding to your IBM XIV Storage System and select Event Log from the pop-up menu that is shown in Figure 7-69.




Figure 7-69 Select Event Log

The Event Log window can be configured to show the events for a defined time frame or to limit the number of entries to display. Selecting a specific event will display the Event Details in a pane on the right side of the window as shown in Figure 7-70.




Figure 7-70 IBM Director Event Log

Event actions

Based on the SNMP traps and the events, you can define various Event Actions with the Event Actions Builder as illustrated in Figure 7-71. Here, you can define various actions for the IBM Director to perform in response to specific traps and events.

IBM Director offers a wizard to help you define an Event Action Plan. Start the wizard by selecting **Tasks** → **Event Action Plans** → **Event Action Plan Wizard** in the IBM Director Console window. The Wizard will guide you through the setup.

The window in Figure 7-71 shows that the IBM Director will send, for all events, an email (to a predefined email address or to a group of email addresses).




Figure 7-71 Event Action Plan Builder

Using IBM Systems Director Version 6

In this section, we illustrate how to use IBM System Director Version 6 to monitor the IBM XIV Storage System. If you are still using IBM Director see section “Using IBM Director” on page 250 for details.

For detailed information regarding the installation, setup and configuration of IBM Systems Director, see the documentation available in the IBM Redbooks publication:

<http://www.redbooks.ibm.com/redbooks/pdfs/sg247694.pdf>

Compile MIB file

After you have completed the installation of your IBM Systems Director environment, you prepare it to manage the IBM XIV Storage System by compiling the provided MIB file.

Make sure to always use the latest MIB file provided. The most recent MIB file can also manage XIVs on older code levels. To compile the MIB file in your environment:

1. In the IBM Systems Director window, click **Settings** → **Manage MIBs**, as shown in Figure 7-72.




Figure 7-72 Import MIB

2. In the Managed MIBs window, click **Import MIB File**.
3. In the Import MIB Files window that is shown in Figure 7-73, specify the directory and file name of the MIB file that you want to compile, and click **OK**. A status window indicates the progress.




Figure 7-73 Browse Window

MIB file extension: IBM Systems Director accepts only MIB files including an extension of *.MIB, and if your MIB file uses another extension, you have to rename it to *.MIB prior to import.

Contact IBM Support to get the most recent MIB File of your XIV.

When you compile a new MIB file, it is also automatically loaded into the *Available MIB Modules directory* and is ready for use.

To load an already compiled MIB file:

1. Select the MIB (that you want to load) within the *Available MIBs Modules Selection* in the Manage MIBs window.
2. Click **Add**, **Apply**, and **OK**.

This action will load the selected MIB file, and the IBM Systems Director is ready to be configured for monitoring the IBM XIV Storage System.

Discover the IBM XIV Storage System

After loading the MIB file into the IBM Director, the next step is to discover the IBM XIV Storage Systems in your environment. Therefore, configure the IBM System Director for auto-discovery as follows:

1. From the IBM Systems Director, select **Inventory** → **System Discovery**, as shown in Figure 7-74 to adjust the discovery preferences.




Figure 7-74 Discovery Preferences

2. In the System Discovery window that is shown in Figure 7-75, follow these steps to discover IBM XIV Storage Systems:
 - a. Select **Range of IPv4 addresses**.
 - b. Next, enter the address range for the IBM XIV Storage Systems in your environment. You also set the Auto-discover period and the Presence Check period.
 - c. Click **Discover Now**.

IBM® Systems Director

Welcome administrator

View: All tasks

- Welcome
- My Startup Pages
- Find a Task
- Find a Resource
- Navigate Resources
- Automation
- Availability
- Inventory
 - System Discovery
 - View and Collect Inventory
 - Views
- Release Management
- Security
- System Configuration
- System Status and Health
- Task Management
- Settings

System Discovery

Use system discovery to discover manageable single IP address, discover resources of the s you to customize discoveries, including impor resources.

Select a discovery option:
Range of IPv4 addresses

Starting IP address:
9 . 155 . 90 . 181

Ending IP address:
9 . 155 . 90 . 183

Select the resource type to discover:
All

Discover Now
Schedule...

Discovered Manageable Systems:

Name	Discovered
xiv-lab-03.mainz.de.ibm....	Previous

Figure 7-75 Discover Range

After you have set up the Discovery Preferences, the IBM Systems Director will discover the IBM XIV Storage Systems and add them to the inventory as displayed in Figure 7-76.




Figure 7-76 View and Collect Inventory

At this point, the IBM Systems Director is ready to receive SNMP traps from the discovered XIV.

3. To review select **Inventory** → **View and Collect Inventory** → **Refresh View**.

With the IBM Systems Director, you can display general information about your IBM XIVs, monitor the Event Log, and browse more information.

Event Log

To open the Event Log, right-click the corresponding IBM XIV Storage System within the View and Collect Inventory window and select **Event Log within System Status and Health** from the pop-up menu that is shown in Figure 7-77.




Figure 7-77 Select Event Log

7.4 Custom Monitoring

The previous sections of this chapter focused on the predefined methods of monitoring the IBM XIV Storage System. We now take a look at ways to combine the various commands, options and features in a manner that allows you to customize the monitoring of the IBM XIV Storage System based on your specific needs or requirements.

Combining the power of XCLI with the flexibility of scripting languages creates endless possibilities and ways to monitor the system. For instance, you can create a script to monitor a specific component and alert you when it exceeds a custom threshold that you define.

7.4.1 Custom monitoring of a volume

Example 7-29 demonstrates how to create a Windows Scripting Host (WSH) script to monitor the capacity of a specific volume and issue a custom event to alert you when the used capacity reaches a predefined limit. The script can be scheduled to run automatically at specific intervals using the standard Windows Task Scheduler.

Example 7-29 WSH Script to monitor a volume

```
<job id="XIV_Vol_Monitor">
<script language="JScript">

//WSH Example for Custom Monitoring of an XIV Volume
//Modify the following parameters with proper values

var IPAddr = "192.168.1.150";
var Volume = "itso_vol_2";
var Threshold = 700;
```

```

var Severity = "Warning";
var Username = "admin";
var Password = "adminpass";

//Main Script Starts Here
var CmdString = "";
var String1 = "";
var output = "";

var xclicmd = "xcli -m " + IPAddr + " -u " + Username + " -p " + Password + "
";
var Command = "vol_list vol=" + Volume + " -t used_capacity";

CmdString = xclicmd + Command;

var WshShell = WScript.CreateObject("WScript.Shell");
var oExec = WshShell.Exec("%comspec% /d/q");

oExec.StdIn.WriteLine("prompt $_");
oExec.StdIn.WriteLine("cd \\");
oExec.StdIn.WriteLine(CmdString);
oExec.StdIn.WriteLine("exit");

while (!oExec.StdOut.AtEndOfStream)
{
    String1 = oExec.StdOut.ReadLine() + "\n";
    String1 = String1.replace(/^\s+|=|\s+$/g, "");
    if (String1 != "")
        output = String1;
}

if (output >= Threshold)
{
    Command = "custom_event description='Vol " + Volume + " exceeded threshold
of " + Threshold + " GB' severity=" + Severity;

    CmdString = xclicmd + Command;

    WshShell = WScript.CreateObject("WScript.Shell");
    oExec = WshShell.Exec("%comspec% /d/q");

    oExec.StdIn.WriteLine("prompt $_");
    oExec.StdIn.WriteLine("cd \\");
    oExec.StdIn.WriteLine(CmdString);
    oExec.StdIn.WriteLine("exit");
}

while (oExec.Status == 0)
    WScript.Sleep(100);

WshShell.LogEvent(0, "XIV Vol Monitor Script Completed");

</script>
</job>

```

The script in Example 7-29 on page 261 begins with defining variables that will be used throughout the process. It then builds an XCLI command based on these variables and runs the command in a shell.

The output from XCLI commands can be customized using the -t option to limit the display to only the specific fields we are interested in. In this case we only need the used capacity value of the volume so the -t option is used to limit the output to only the used_capacity field. This option makes the output of the XCLI command easy to parse.

Example 7-30 Limiting output of XCLI commands using the -t option

```
var Command = "vol_list vol=" + Volume + " -t used_capacity"
```

You can get a complete list of the field options available for any command by using the XCLI help option with the format=full parameter.

Example 7-31 Using the XCLI help command to list available output fields

```
>> help command=vol_list format=full
```

The output from the XCLI command is parsed and compared with the predefined custom threshold. If the threshold is exceeded, an XCLI command is used to create a custom event with a custom description and severity level. An event rule can be used on the IBM XIV Storage System to notify a user of the condition. See Example 7-32.

Example 7-32 Using XCLI to generate a custom event

```
if (output >= Threshold)
{
    Command = "custom_event description='Vol " + Volume + " exceeded threshold
of " + Threshold + " GB' severity=" + Severity;
```

After the script is created, use the Windows Task Scheduler to schedule the script to run automatically on a daily basis by selecting **Start** → **Programs** → **Accessories** → **System Tools** → **Scheduled Tasks** as shown in Figure 7-78.




Figure 7-78 Using Windows Task Scheduler to run custom monitoring scripts

After running this script, a custom event is generated if the used capacity threshold is exceeded. You can see the custom event in the XIV Event Log as shown in Figure 7-79. An event notification rule can be defined based on severity or the CUSTOM_EVENT event code to send an email or SNMP message as described in section 7.3, “XIV event notification” on page 237.




Figure 7-79 A Custom Event generated by WSH script using XCLI

7.4.2 Custom monitoring of performance attributes

In Example 7-33, we demonstrate how to create a Windows Scripting Host (WSH) script to monitor system performance, specifically the write hit latency, and issue a custom event to alert you when the latency reaches a predefined limit. The script can be scheduled to run automatically at specific intervals using the standard Windows Task Scheduler.

Example 7-33 WSH script to monitor write hit latency

```
<job id="XIV_Latency_Monitor">
<script language="JScript">

//WSH Example for Custom Monitoring of an XIV Latency Performance
//Modify the following parameters with proper values

var IPAddr = "192.168.1.150";
var Threshold = 3000;
var Severity = "Warning";
var Username = "admin";
var Password = "adminpass";

//Main Script Starts Here
var strValidChars = "0123456789";
var strChar = "";
var CmdString = "";
var String1 = "";
var output = "";
var Exceeded = 0;

var xclicmd = "xcli -m " + IPAddr + " -u " + Username + " -p " + Password + "
";
```

```

var now = new Date();
var EndDate = now.getYear() + "-" + (now.getMonth()+1) + "-" +
(now.getDate()-1) + ".00:00:00";

var Command = "statistics_get count=1440 interval=1 resolution_unit=minute -t
write_hit_medium_latency end=" + EndDate;

CmdString = xclicmd + Command;

var WshShell = WScript.CreateObject("WScript.Shell");
var oExec = WshShell.Exec("%comspec% /d/q");

oExec.StdIn.WriteLine("prompt $_");
oExec.StdIn.WriteLine("cd \\");
oExec.StdIn.WriteLine(CmdString);
oExec.StdIn.WriteLine("exit");

while (!oExec.StdOut.AtEndOfStream)
{
    String1 = oExec.StdOut.ReadLine() + "\n";
    String1 = String1.replace(/^\s+|=|\s+$/g,"");
    if (String1.length == 0) continue;

    strChar = String1.charAt(0);
    if (strValidChars.indexOf(strChar) == -1) continue;
    if (String1 >= Threshold) Exceeded = 1;
}

if (Exceeded)
{
    Command = "custom_event description='Medium Latency exceeded threshold of "
+ Threshold + " during 24 hour period ending " + EndDate + "' severity=" +
Severity;
    CmdString = xclicmd + Command;

    WshShell = WScript.CreateObject("WScript.Shell");
    oExec = WshShell.Exec("%comspec% /d/q");
    oExec.StdIn.WriteLine("prompt $_");
    oExec.StdIn.WriteLine("cd \\");
    oExec.StdIn.WriteLine(CmdString);
    oExec.StdIn.WriteLine("exit");
}

while (oExec.Status == 0)
    WScript.Sleep(100);

WshShell.LogEvent(0, "XIV Latency Monitor Script Completed Successfully");
</script>
</job>

```

The script in Example 7-33 on page 264 also begins with defining variables that are used throughout the process. It then builds an XCLI command based on these variables and runs the command in a shell.

The output from XCLI is again customized using the -t option to limit the display to only the specific fields we are interested in. In this case we only need the write hit medium latency field. This option makes the output of the XCLI command easy to parse.

The XCLI command is structured in a way to provide output during a custom window of time. In Example 7-34, we are gathering data from every minute during a 24 hour period on the previous day the script is run.

Example 7-34 Limiting output of XCLI commands using the -t option

```
var Command = "statistics_get count=1440 interval=1 resolution_unit=minute -t  
write_hit_medium_latency end=" + EndDate;
```

In Example 7-35, each resulting value is compared to the threshold and if any one of them exceeds the value, it will set a flag to be used to generate a custom event.

Example 7-35 Looping through XCLI command output

```
while (!oExec.Stdout.AtEndOfStream)  
{  
    String1 = oExec.Stdout.ReadLine() + "\n";  
    String1 = String1.replace(/^\s+|=|\s+$/g, "");  
    if (String1.length == 0) continue;  
  
    strChar = String1.charAt(0);  
    if (strValidChars.indexOf(strChar) == -1) continue;  
    if (String1 >= Threshold) Exceeded = 1;  
}
```

In Example 7-36, we see again how to generate a custom event when our desired threshold is exceeded.

Example 7-36 Using XCLI to generate a custom event

```
if (Exceeded)  
{  
    Command = "custom_event description='Medium Latency exceeded threshold of "  
+ Threshold + " during 24 hour period ending " + EndDate + "' severity=" +  
Severity;
```

After running this script, you can see the custom event in the XIV Event Log as shown in Figure 7-80.

After running this script, a custom event is generated if the write hit medium latency threshold is exceeded. You can see the custom event in the XIV Event Log as shown in Figure 7-80. An event notification rule can be defined based on severity or the CUSTOM_EVENT event code to send an email or SNMP message as described in 7.3, “XIV event notification” on page 237.




Figure 7-80 A Custom Event generated by WSH script using XCLI

7.4.3 Conclusion

The simple examples provided are just a small subset of what can be accomplished by combining the power of XCLI commands with the flexibility of scripting.

By using the output from multiple commands or multiple fields, it is possible to monitor any combination of components or parameters of the system and act upon them based on criteria you define.

Similar results can be accomplished on UNIX based platforms by using shell scripting and crontab for scheduling custom monitoring processes.

7.5 Call Home and Remote support

The Call Home function allows the IBM XIV Storage System to send event notifications to the XIV Remote Support Center. This enables both proactive and failure notifications to be sent directly to IBM for analysis. The XIV Remote Support Center will take appropriate action, up to dispatching an IBM service representative with a replacement part, or engaging level 2 or higher to ensure complete problem determination and solution.

7.5.1 Call Home

Call Home is always configured to use SMTP, and is only configured by qualified IBM service personnel, typically when the IBM XIV Storage System is first installed.

7.5.2 Remote support

The IBM XIV Storage System is repaired by trained IBM service personnel, either remotely with the help of the IBM XIV Remote Support Center, or on-site by an IBM SSR. When problems arise, a remote support specialist can connect to the system to analyze the problem, repair it remotely if possible, or assist the IBM SSR who is on-site.

The remote support center has three ways to connect the system. Depending on the client's choice, the support specialist can either connect by

- ▶ Using a modem dial-up connection, using a analog phone line provided by the client
- ▶ Using a secure, high-speed direct connection through the Internet to the IBM XIV Storage System
- ▶ Using the XIV Remote Support Center (XRSC), which allows the customer to initiate a secure connection from the XIV to IBM. Using XRSC, the IBM XIV Storage System makes a connection to an external XRSC server. Using an internal XRSC server, the IBM XIV Remote Support Center can connect to the XIV, through the connection made to the external server. For details, see Figure 7-81.

XRSC: We encourage all customers to use the secure, high speed remote support solution enabled by the XIV Remote Support Center (XRSC).

These possibilities are depicted in Figure 7-81. In case of problems, the remote specialist is able to analyze problems and also assist an IBM SSR dispatched on-site in repairing the system or in replacing field-replaceable units (FRUs).

To enable remote support, you must allow an external connection, such as either:

- ▶ A telephone line
- ▶ An Internet connection through your firewall that allows IBM to use a secure SSH connection to your IBM XIV Storage System




Figure 7-81 Remote Support connections

XIV Remote Support Center connection

The XIV Remote Support Center (XRSC) uses a high speed Internet connection, but gives the client the ability to initiate an outbound SSH call to a secure IBM server.

The XIV Remote Support Center is comprised of the IBM XIV Storage Servers internal functionality together with a set of globally deployed supporting servers to provide secure IBM support access to the IBM XIV Storage System when necessary and when authorized by the customer personnel.

Objectives and components

The XIV Remote Support Center is designed to meet three major objectives:

- **Security from malicious attack**

The XIV Remote Support Center was designed with security as a major concern, while keeping the system architecture simple and easy to deploy. It relies on standard, proven technologies and minimizes the logic (code) that must reside either on the External XRSC server or on customer machines. The architecture of the Remote Support Center anticipates and handles a variety of potential attack vectors, including encryption protocol and authentication logic. The design of the Remote Support Center identifies, analyzes and defends against these and other possible assault routes.

- **Simplicity of implementation and design**

Simplicity enhances security, facilitates implementation at the customer site, decreases staff training costs and ultimately lowers TCO by allowing more issues to be resolved quickly and with fewer unnecessary visits by support personnel to customer premises.

- **Accessibility by support personnel**

The auditing features of the XIV Remote Support Center both enhance security and enable easy re-assignment of a support case among support personnel. The Remote Support Center can expedite escalation of support cases to higher levels and allows IBM support representatives engineers to pool knowledge to resolve complex customer issues.

Underlying architecture

The XIV Remote Support mechanism has four major components (see Figure 7-82):

- Remote Support Client (machine internal):

The Remote Support Client is a software component inside the IBM XIV Storage System that handles remote support connectivity. It relies only on a single outgoing TCP connection, and has no capability to receive inbound connections of any kind. The Client is controlled using XCLI, and is charged with starting a connection, terminating a connection (due to time-out or customer request) and re-trying the connection in case it terminates unexpectedly.

- Optional Remote Support Proxy

The Remote Support Client can access the Remote Support Center Front Server directly, or through an optional proxy server. The optional Remote Support Proxy can be used when one or more XIV systems do not have direct access to the Internet (for example, due to firewall restrictions), you can use the XIV Remote Support Proxy to facilitate the connection to the XIV Remote Support Center. More information about the Remote Support Proxy can be found in the *Installation and User's Guide*, GA32-0795.

- Remote Support Center Front Server (Internet):

Front Servers are located on an IBM DMZ of the Internet and receive connections from the Remote Support Client and the IBM XIV Remote Support Back Server. Front Servers are security-hardened machines which provide a minimal set of services, namely, maintaining connectivity to connected Clients and to the Back Server. They are strictly inbound, and never initiate anything on their own accord. No sensitive information is ever stored on the Front Server, and all data passing through the Front Server from the Client to the Back server is encrypted so that the Front Server or a malicious entity in control of a Front Server cannot access this data.

► Remote Support Center Back Server (IBM Intranet):

The Back Server manages most of the logic of the system. It is located within IBM's Intranet. The Back Server is access controlled. Only IBM employees authorized to perform remote support of IBM XIV are allowed to use it, and only through specific support interfaces, not with a CLI or a GUI shell. The Back Server is in charge of authenticating a support person, providing the support person with a UI through which to choose a system to support based on the support person's permissions and the list of systems currently connected to the Front Servers and managing the remote support session as it progresses (logging it, allowing additional support persons to join the session, and so on). The Back Server maintains connection to all Front Servers. Support people connect to the Back Server using any SSH client or an HTTPS connection with any browser.

Figure 7-82 provides a representation of the data flow of the XIV to IBM Support.




Figure 7-82 XIV Remote Support Center

To initiate the remote connection process, the following steps are performed:

1. Customer initiates an Internet based SSH connection to XRSC either by the GUI or XCLI
2. XRSC identifies the IBM XIV Storage System and marks it as "connected"
3. Support personnel connects to XRSC using SSH over the IBM Intranet
4. XRSC authenticates the support person against the IBM Intranet
5. XRSC then displays the connected customer system available to the support personnel
6. The IBM Support person then chooses which system to support and connect to
 - Only permitted XIV systems are shown
 - IBM Support personnel log their intended activity
7. A fully recorded support session commences
8. When complete, the support person terminates the session and the XRSC disconnects the IBM XIV Storage System from the remote support system.

Customer usage scenario

After a customer has decided that remote support is needed, possibly due to a telephone request from a remote support professional, communication between the Remote Support Center and the customer site must be established. This process is always initiated by the customer and always remains under the customer's control. It proceeds as follows:

- The customer tells the IBM XIV system to connect to the Remote Support Center using a specific XCLI command. Until this command is issued, an IBM XIV system will not connect to the Remote Support Center for any reason.

- ▶ The customer can set a regular timeout and an 'idle timeout' for the connection, using an optional parameter in the connection command. A timeout will cause the connection to be dropped after a specified length of time whether a support session is in progress or not, while an idle timeout will cause the connection to be dropped after a certain period of inactivity. Specifying the keyword "never" for these two parameters will allow the connection to remain established indefinitely.
- ▶ The customer can also chose to set a unique password for a session when it is initiated. This password must be provided to the IBM support representative in order for them to complete the authentication process and access the IBM XIV Storage Server.
- ▶ In addition, the customer can see the Remote Support Center connectivity status at any time (disconnected, idle, in-session) and can choose to forcibly disconnect the connection regardless of whether a support session is in progress or not.

While forcibly terminating a support session using the timeout or force-disconnect mechanisms is not desirable, it is important to note that these mechanisms do exist and that the customer has full control over whether or not to proceed with a support session.

The initiation or loss of connection to the Remote Support Center causes the system to generate events that can be seen on the machine's event log and that can be forwarded to any destination of the customer's choice (like any other event using the IBM XIV event-rules mechanism). Connection loss events, whether due to a transport error, timeout or due to specific customer action, specify whether the connection was idle or in-use at the time of disconnection. A warning event is issued 15 minutes before the timeout parameter disconnects a busy support session. If the connection is lost due to a network error before the timeout for the session has expired, the system will automatically try to reconnect to any of the configured Remote Support Center servers.

While a support session is in progress, the IBM XIV generates events and displays the machine's status on the GUI screen as usual. So, for example, the customer can see the process of phasing out a module or the restart of customer-visible machine services as they happen.

Starting an XIV Remote Support Center connection using the GUI

The XRSC definitions are configured at the time the system is installed. To start the connection from the GUI, select Tools->Support Center from the drop down menu as shown in Figure 7-83.




Figure 7-83 Starting an XRSC Remote Support Center connection

Select the desired Support Center from the list and click the **Connect Support Center** icon as shown in Figure 7-84.




Figure 7-84 Select the desired Support Center and click Connect

The next pop-up windows, as shown in Figure 7-85, will prompt for the **Session Timeout**, **Idle Timeout** and **Connection Password**. The timeout values are specified in minutes and will disconnect the session when they expire. The password, if set, must be given to the support representative in order for them to establish a connection. Use of a password is not required and can be left blank.

Using **Never** for the timeout values will result in the connection remaining open until explicitly disconnected.




Figure 7-85 Starting an XRSC Connection

After the connection to the Remote Support Center has been established, you can disconnect the session by clicking the Disconnect Support Center icon as shown in Figure 7-86.




Figure 7-86 Disconnecting a Remote Support Center session

Starting an XIV Remote Support Center connection using XCLI

The following XCLI commands are used to manage Remote Support Center connections

- ▶ support_center_list
- ▶ support_center_status
- ▶ support_center_connect
- ▶ support_center_disconnect

To start a Remote Support Center connection using GUI, first display a list of the available Support Centers that were configured during system installation by issuing the XCLI command **support_center_list** as shown in Example 7-37.

Example 7-37 List available support centers

```
>> support_center_list
```

Name	Address	Port	Priority
israel-xiv	195.110.41.141	22	0
israel-xiv-beta	212.143.102.242	22	1

Contact your support representative if the **support_center_list** commands do not return a value and you wish to use this feature.

Next, you can see what the current status of the connection is by using the XCLI command **support_center_status** (Example 7-37).

Example 7-38 Check support center connection status

```
>> support_center_status
```

State	Connected sessions	Timeout (min)	Module	Connected since
no connection	?	?	?	?

In the example before, no connections are currently established. To start a session, use the XCLI command **support_center_connect** followed by **support_center_status**.

Example 7-39 Connect to support center

```
>> support_center_connect
command 0:
administrator:
    command:
        code = "SUCCESS"
        status = "0"
        status_str = "Command completed successfully"
aserver = "DELIVERY_SUCCESSFUL"

>> support_center_status
State   Connected sessions   Timeout (min)   Module           Connected since
-----
idle    0                     no timeout     1:Module:4       2010-10-08 10:45:35
```

If no parameters are specified with the **support_center_connect** command, the timeout values default to never and no password is used. You can specify the parameters for **timeout**, **idle_timeout** and **password** as shown in Example 7-40. The timeout values are in the format hh:mm when using the XCLI commands.

Example 7-40 Specify connection parameters

```
>> support_center_connect timeout=01:00 idle_timeout=00:30 password=4support
command 0:
administrator:
    command:
        code = "SUCCESS"
        status = "0"
        status_str = "Command completed successfully"
aserver = "DELIVERY_SUCCESSFUL"

>> support_center_status
State   Connected sessions   Timeout (min)   Module           Connected since
-----
idle    0                     27.6           1:Module:4       2010-10-08 10:49:40
```

The status will show an idle state until support center representative establishes a connection at which time it will show a state of busy as illustrated in Example 7-41.

Example 7-41 Connection busy

```
>> support_center_status
State   Connected sessions   Timeout (min)   Module           Connected since
-----
busy    1                     40.2           1:Module:4       2010-10-08 10:49:40
```

To terminate the remote support center connection, use the **support_center_disconnect** XCLI command, as shown in Example 7-42.

Example 7-42 Disconnect from support center

```
>> support_center_disconnect
command 0:
administrator:
    command:
        code = "SUCCESS"
        status = "0"
        status_str = "Command completed successfully"
aserver = "DELIVERY_SUCCESSFUL"
```

```
>> support_center_status
```

State	Connected sessions	Timeout (min)	Module	Connected since
-----	-----	-----	-----	-----
no connection	?	?	?	?

7.5.3 Repair flow

In case of system problems, the IBM XIV Support Center will be notified by a hardware or software call generated by a notification from the system or by a user's call.

Flowchart for repair troubleshooting

Based on this call, the remote support center will initiate the necessary steps to repair the problem according to the flow depicted in Figure 7-87.




Figure 7-87 Problem Repair Flow

Either a call from the user or an email notification will generate an IBM internal problem record and alert the IBM XIV Support Center. A Support Center Specialist will remotely connect to the system and evaluate the situation to decide what further actions to take to solve the reported issue:

- ▶ Remote Repair: Depending on the nature of the problem, a specialist will fix the problem while connected.
- ▶ Data Collection: Start to collect data in the system for analysis to develop an action plan to solve the problem.
- ▶ On-site Repair: Provide an action plan, including needed parts, to the call management center (CMC) to initiate an IBM SSR repairing the system on-site.
- ▶ IBM SSR assistance: Support the IBM SSR during on-site repair by remote connection.

The architecture of the IBM XIV Storage System is self-healing. Failing units are logically removed from the system automatically, which greatly reduces the potential impact of the event and results in service actions being performed in a fully redundant state.

For example, if a disk drive fails, it will be automatically removed from the system. The process has a minimal impact on performance, because only a small part of the available resources has been removed. The rebuild time is fast, because most of the remaining drives will participate in redistributing the data.

Due to this self-healing mechanism, with most failures, there is no need for urgent action and service can be performed at a convenient time. The IBM XIV will be in a fully redundant state, which mitigates issues that might otherwise arise if a failure occurs during a service action.

Data collection

It might be necessary to collect data from the system in order to fully diagnose a problem. In most cases, the support center representative can collect this data remotely, but in various cases the customer might be asked to collect the data and send it to the support center. To collect the data from the GUI, select Collect Support Logs from the Tools drop down menu as shown in Figure 7-88. The collected file package is referred to as an Xray.




Figure 7-88 Collecting Support Logs for problem analysis

A dialog box will appear to allow you to specify location for the logs. Chose a location and click **Collect**. The collection process can take a few minutes to complete. A status message with the percentage completed is displayed, as illustrated in Figure 7-89.




Figure 7-89 Specify the location for the log file

After the collection process is complete, click the **Get** button to download the file to the specified location. See Figure 7-90.





Figure 7-90 Downloading system log file to specified location

The support center representative who requested the logs will provide further information about how to send the logs in for analysis.

To view the files contained in an Xray, simply unzip the file into a specified directory using any compression utility that supports the .zip format and open the xray.html file in a standard web browser as shown in Figure 7-91.





A

Additional LDAP information


In this appendix we provide additional LDAP related information:

- ▶ Creating user accounts in Active Directory
- ▶ Creating user accounts in SUN Java Directory
- ▶ Securing LDAP communication with SSL
 - Windows Server SSL configuration
 - SUN Java Directory SSL configuration
- ▶ Certificate Authority setup

Creating user accounts in Microsoft Active Directory

Creating an account in Microsoft Active Directory for use by XIV LDAP authentication is just like creating any regular user account. The only exception is the designated “*description*” attribute (field). This field must be populated with the predefined value in order for the authentication process to work. Proceed as follows:

1. Start Active Directory Users and Computer by selecting **Start** → **Administrative Tools** → **Active Directory Users and Computers**
2. Right-click “Users” container, select **New** → **User**. The New Object User dialog window opens as seen in Figure A-1.



The screenshot shows the 'New Object - User' dialog box. At the top, it says 'Create in: xivhost1ldap.storage.tucson.ibm.com/Users'. Below this are several input fields: 'First name:', 'Last name:', 'Full name:' (containing 'xivtestuser1'), 'User logon name:' (containing 'xivtestuser1'), and a dropdown menu for the domain (showing '@xivhost1ldap.storage.tucson.ibm.com'). There is also a section for 'User logon name (pre-Windows 2000):' with two sub-fields containing 'XIVHOST1LDAP\' and 'xivtestuser1'. At the bottom are three buttons: '< Back', 'Next >', and 'Cancel'.

Figure A-1 Creating Active Directory user account

The value entered in “Full name” is what XIV will use as the User name. The only other mandatory field in this form is “User logon name”. For simplicity the same *xivtestuser1* value is entered into both fields. Other fields can also be populated but it is not required.

3. Proceed with creating the account by clicking **Next**. A new dialog window, shown in Figure A-2 is displayed.


A screenshot of a Windows-style dialog box titled "New Object - User". The dialog has a blue title bar with a close button. Below the title bar, there is a small user icon and the text "Create in: xivhost1ldap.storage.tucson.ibm.com/Users". The main area contains two password input fields: "Password:" and "Confirm password:", both filled with dots. Below these fields are four checkboxes: "User must change password at next login" (checked), "User cannot change password" (unchecked), "Password never expires" (unchecked), and "Account is disabled" (unchecked). At the bottom, there are three buttons: "< Back", "Next >", and "Cancel".

Figure A-2 Assigning password

Note that by default, the password is set to “User must change password at next login”. After the account is created, the user must logon to a server that is part of the Active Directory managed domain to change the password. After the password is changed, all the security rules and policies related to password management are in effect, such as password expiration, maintaining password change history, verifying password complexity, and so on.

Password: If the password initially assigned to an Active directory user is not changed - XIV will not authenticate that user.

4. Complete the account creation by pressing **Next** → **Finish**.
5. Proceed with populating the “*description*” field with predefined value for XIV category (role) mapping by selecting the *xivtestuser1* user name followed by a right-click and selecting **Properties**, as illustrated in Figure A-3.

The screenshot shows the 'xivtestuser1 Properties' dialog box. The 'General' tab is active. The 'Description' field is populated with 'Storage Administrator'. The 'Display name' field shows 'xivtestuser1'. Other fields like 'First name', 'Last name', 'Office', 'Telephone number', 'E-mail', and 'Web page' are empty. The 'OK', 'Cancel', and 'Apply' buttons are visible at the bottom.

Figure A-3 Entering predefined value into Description field

Complete the account information update by clicking **OK**.

After the user account is created in Active Directory, its accessibility can be verified from any of the available LDAP clients. In our case we used the OpenLDAP client, as shown in Example A-1.

Example: A-1 Active Directory account verification using OpenLDAP client

```
$ ldapsearch -x -H "ldap://xivhost1.xivhost1ldap.storage.tucson.ibm.com:389" -D
'CN=xivtestuser1,CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com' -w
pass2remember -b 'CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com'
cn=xivtestuser1
```

```
dn: CN=xivtestuser1,CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: user
cn: xivtestuser1
description: Storage Administrator
distinguishedName: CN=xivtestuser1,CN=Users,DC=xivhost1ldap,DC=storage,DC=tuc
on,DC=ibm,DC=com
instanceType: 4
whenCreated: 20090622172440.0Z
whenChanged: 20090622180134.0Z
displayName: xivtestuser1
uSNCreated: 98467
uSNChanged: 98496
name: xivtestuser1
```



```
objectGUID:: apHajqyazEyALYHDAJrjNA==
userAccountControl: 512
badPwdCount: 0
codePage: 0
countryCode: 0
badPasswordTime: 128901682350000000
lastLogoff: 0
lastLogon: 128901682415312500
pwdLastSet: 128901672940468750
primaryGroupID: 513
objectSid:: AQUAAAAAAAAUVAAN59TxdIlskwvBQmdAQAAA==
accountExpires: 9223372036854775807
logonCount: 3
sAMAccountName: xivtestuser1
sAMAccountType: 805306368
userPrincipalName: xivtestuser1@xivhost1ldap.storage.tucson.ibm.com
objectCategory: CN=Person,CN=Schema,CN=Configuration,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com
```

The **ldapsearch** command syntax might appear overly complex and its output difficult for interpretation. However this might be the easiest way to verify that the account was created as expected. The **ldapsearch** command can also be very useful for troubleshooting purposes when you are unable to communicate with Active directory LDAP server.

Here is a brief explanation of the **ldapsearch** command line parameters:

-H 'ldap://xivhost1.xivhost1ldap.storage.tucson.ibm.com:389' - specifies that the LDAP search query must be sent to "xivhost1.xivhost1ldap.storage.tucson.ibm.com" server using port number 389.

-D 'CN=xivtestuser1,CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com' - the query is issued on behalf of user "xivtestuser1" registered in "Users" container in "xivhost1ldap.storage.tucson.ibm.com" Active Directory domain.

-w pass2remember - is the current password of the user "xivtestuser1" (after the initially assigned password was changed to this new password).

-b 'CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com' - Base_DN, the location in the directory where to perform the search, the "Users" container in the "xivhost1ldap.storage.tucson.ibm.com" Active Directory domain.

cn=xivtestuser1 - specifies what object to search for

The output of the **ldapsearch** command shows the structure of the LDAP object retrieved from the LDAP repository. We do not need to describe every attribute of the retrieved object, however, at least two attributes must be checked to validate the response:

```
name: xivtestuser1
description: Storage Administrator
```

The fact that **ldapsearch** returns the expected results in our example indicates that:

1. The account is indeed registered in Active Directory
2. The distinguished name (DN) of the LDAP object is known and valid
3. The password is valid
4. The designated attribute "description" has a predefined value assigned "Storage Administrator"

When the Active Directory account verification is completed, we can proceed with configuring the XIV System for LDAP authentication mode. At this point we still have a few unassigned LDAP related configuration parameters in our XIV System as can be observed in Example A-2.

Example: A-2 Remaining XIV LDAP configuration parameters

>> ldap_config_get	
Name	Value
base_dn	
xiv_group_attr	description
third_expiration_event	7
version	3
user_id_attr	objectSid
current_server	
use_ssl	no
session_cache_period	
second_expiration_event	14
read_only_role	Read Only
storage_admin_role	Storage Administrator
first_expiration_event	30
bind_time_limit	0

base_dn - base DN (distinguished name), the parameter which specifies where in the Active Directory LDAP repository that a user can be located. In our example we use "CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com" as base DN, see Example A-1 on page 282.

current_server - is read-only parameter and can not be populated manually.. It will get updated by the IBM XIV Storage System after the initial contact with LDAP server is established.

session_cache_period - duration in minutes the IBM XIV Storage System keeps user credentials in its cache before discarding the cache contents. If a user repeats the login attempt within **session_cache_period** minutes from the first attempt, authentication will be done from the cache content without contacting LDAP server for user credentials.

bind_time_limit - the timeout value in seconds after which the next LDAP server on the **ldap_list_servers** is called. The default value for this parameter is 0. It must be set to a non-zero value in order for bind (establishing LDAP connection) to work. The rule also applies to configurations where the XIV System is configured with only a single server on the **ldap_list_servers** list.

The populated values are shown in Example A-3.

Example: A-3 Completing and verifying LDAP configuration on XIV

```
>> ldap_config_set
base_dn="CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com"
session_cache_period=10 bind_time_limit=30
Command executed successfully.

$ xcli -c "XIV MN00019" -u ITS0 -p redb00k ldap_config_get
Name                               Value
base_dn CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com
xiv_group_attrib                    description
third_expiration_event              7
version                             3
user_id_attrib                      objectSid
current_server
use_ssl                             no
session_cache_period                10
second_expiration_event             14
read_only_role                      Read Only
storage_admin_role                  Storage Administrator
first_expiration_event              30
bind_time_limit                     30
```

To complete our description of the LDAP related configuration parameters (at the IBM XIV Storage System), we have to discuss the parameters that had default values assigned and did not have to be set explicitly:

- ▶ **version** - version of LDAP protocol used, default “3”. This parameter must never be changed. Both LDAP products - Active Directory and Sun Java Services Directory Server Enterprise Edition support LDAP protocol version 3
- ▶ **user_id_attrib** - LDAP attribute set to identify the user (in addition to user name) when recording user operations in the XIV event log, Default objectSid value corresponds to the existing attribute name in Active Directory LDAP object class.
- ▶ **use_ssl** - indicates if secure (SSL encrypted) LDAP communication is mandated. Default value is “no”. If set to “yes” without configuring both sides for SSL encrypted communication, will result in failing LDAP authentication at the IBM XIV Storage System
- ▶ **first_expiration_event** - number of days before expiration of certificate to set first alert (severity “warning”). Set this parameter to a number of days that will give you enough time to generate and deploy new security certificate.
- ▶ **second_expiration_event** - number of days before expiration of certificate to set second alert (severity “warning”)
- ▶ **third_expiration_event** - number of days before expiration of certificate to set third alert (severity “warning”)

Now that all configuration and verification steps are completed, the XIV System is ready for the LDAP mode to be activated.

Creating user accounts in SUN Java Directory

Creating an account in SUN Java Directory can be done in many ways using various LDAP clients. For illustration purposes we used the LDAP GUI client that is part of Java System Directory Service Control Center - web tool that is part of the SUN Java Directory Server product suit.

The designated *description* attribute must be populated with the predefined value in order for the authentication process to work. From the SUN Java Directory LDAP Server perspective, assigning value to the *description* attribute is not mandatory and will not be enforced by the server itself. LDAP server will allow creating an account with no value assigned to the attribute. However this attribute value is required by the IBM XIV Storage System for establishing LDAP role mapping. This field must be populated with the predefined value in order for the authentication process to work.

To launch Java System Directory Service Control Center point your browser to the ip address of your SUN Java Directory LDAP Server, for a secure connection on port 6789:

In our example we use the following URL for accessing Java System Directory Service Control Center: "https://xivhost2.storage.tucson.ibm.com:6789"

Before the first user account can be created, the LDAP administrator must create a suffix. A suffix (also known as a naming context) is a DN that identifies the top entry in the directory hierarchy. SUN Java Directory LDAP server can have multiple suffixes, each identifying a locally held directory hierarchy. For example, o=ibm or in our specific example dc=xivauth

To create a suffix, login to the Java Console using your own userid and password and select "Directory Service Control Center (DSCC)" link in "Services" section. Authenticate to "Directory Service Manager" application. In "Common Tasks" tab select **Directory Entry Management** → **Create New Suffix or Replication Topology**. Continue as follows:

1. Enter Suffix Name. Specify the new suffix DN; In our example dc=xivauth. Click **Next**.
2. Choose Replication Options. Accept the default "Do Not Replicate Suffix" (LDAP replication is beyond the scope for this book).
3. Choose Server(s). Select "xivhost2.storage.tucson.ibm.com:389" in Available Servers list and click **Add**. The server name displays in the Chosen Servers list. Click **Next**.
4. Choose Settings. Accept the default "Use Default Settings".
5. Choose Database Location Options. Accept the default database location. Click **Next**.
6. Choose Data Options. Select "Create Top Entry for the Suffix. Click **Next**.
7. Review settings for the suffix about to be created and click **Finish** if they are correct.

After the new suffix creation is confirmed, you can proceed with LDAP entry creation.

To create a new LDAP entry login to the Java Console using your own and password, select "Directory Service Control Center (DSCC)" link in "Services" section. Authenticate to "Directory Service Manager" application. In "Common Tasks" tab select **Directory Entry Management** → **Create New Entry**.

Now you can launch the New Entry configuration wizard, as follows:

1. The first step of the process is selecting a server instance. SUN Java Directory allows you to create multiple instances of an LDAP server. However the only instance that uses port 389 for non-SSL LDAP and port 636 for SSL LDAP communication can be used for XIV LDAP authentication services. In step one select an instance configured on port 384, as illustrated in Figure A-4.




Figure A-4 Selecting Directory Server Instance

2. The second step (see Figure A-5) is selecting the new entry location. The LDAP administrator determines the location of a new entry. Unlike the Active Directory LDAP repository, where location is directly linked to the domain name, SUN Java Directory LDAP server provides greater flexibility in terms of placement for the new entry.

The location of all entries for XIV accounts must be the same because the IBM XIV Storage System has only one LDAP configuration parameter that specifies the location. In our example we use `dc=xivauth` as the entry location for XIV user accounts. For simplicity in this example the location name is the same as the suffix name. There are certain similarities between Windows file system and LDAP directory structures. You can think of LDAP suffixes as drive letters. A drive letter can contain directories but you can also put files onto the root directory on a drive letter. In our example we put new account at the level (by analogy) of a root directory, the `dc=xivauth` location.

As your LDAP repository grows, it might no longer be practical to put types of entries into the same location. In this case, just like with Windows file system, you can create subdirectories and place new entries there and so on. And the LDAP equivalent of what has become a directory hierarchy in your file system is called the Directory Information Tree (DIT). After the entry location is selected and the XIV System is configured to point to that location, all the new account entries can only be created in that location.

Step 2: Specify Entry Location

Provide the location DN where the entry will be written in the Directory Server.

* Indicates required field

* Entry Browse...

Parent For example: dc=example,dc=com

DN:

Previous Next Cancel

Figure A-5 Selecting entry location

3. Step 3 of the process is to select an object class for the new entry. And again, unlike a predefined object class for a user account in Active Directory LDAP, SUN Java Directory LDAP presents you with a choice of object class types.

The LDAP object class describes the content and purpose of the object. It also contains a list of attributes, such as a name, surname or telephone number. Traditionally `inetOrgPerson` object class type is used for LDAP objects describing personal information hence its name - Internet Organizational Person. To be compatible with XIV System an object class must include a minimal set of attributes. These attributes are:

- uid - user identifier (user name)
- userPassword - user password
- description (configurable) - LDAP role mapping attribute

You can select another object class type as long as it contains the same minimal set of attributes. Note that object class type can enforce certain rules, for instance, various attributes can be designated as mandatory in which case a new LDAP object can not be created without assigning a value to that attribute. In case of `inetOrgPerson` object there are two mandatory attributes - `cn` "Full Name" also called "Common Name" and `sn` - "Full Name" also called "Surname". Although it is possible to populate these two objects with various values for simplicity reasons we will be using the `uid` value to populate both `cn` and `sn` attributes. See Figure A-6.

4. Step 4 of the process is entering the attribute values. The first field "Naming Attribute" must remain "uid"; XIV is using that attribute name for account identification. Then, we populate the mandatory attributes with values as described in the previous step. You can also choose to populate other optional attributes and store their values in the LDAP repository; However XIV will not use those attributes. See Figure A-7.

Step 3:Choose Object Class

Choose the object type of the entry you want to create.

Entry Type:

Common Objects

User - (inetOrgPerson)

Static Group - (groupOfUniqueNames)

Dynamic Group - (groupOfUrls)

Certificates Group - (groupOfCertificates)

Managed Role - (nsManagedRoleDefinition)

Filtered Role - (nsFilteredRoleDefinition)

Domain Component - (domain)

Organizational Unit - (organizationalUnit)

Referral - (referral)

Class of Service - (cosPointerDefinition)

User-defined Object Classes

-No user-defined objectclass-

Previous

Next

Cancel

Figure A-6 Object class selection

Step 4:Configure Attributes

Enter the attribute values for the new entry. For multi-valued attributes, press the Enter key in the field to make the field taller and enter values on separate lines.

Required Attributes

Naming Attribute:

User ID (uid)

* Full Name (cn):

xivtestuser2

* Last Name (sn):

xivtestuser2

Allowed Attributes

First Name (givenname):

User ID (uid):

xivtestuser2

Password (userPassword):

.....

Confirm Password:

.....

E-mail (mail):

Telephone Number (telephoneNumber):

Fax Number (facsimileTelephoneNumber):

Locality (l):

Organization (o):

Organizational Unit (ou):

audio:

businessCategory:

carLicense:

departmentNumber:

description:

Storage Administrator

destinationIndicator:

Figure A-7 Entering object attribute values

- Step 5 is the last step of the process. A Summary panel (Figure A-8) shows what you have selected and entered in the previous steps. You can go back and change parameters if you choose to do so. Otherwise you proceed with the entry creation.

Step 5: Summary

Review your settings and click finish if they are correct.

Entry DN:	uid=xivtestuser2,dc=xivauth
Object Class:	User - (inetOrgPerson)
Full Name (cn):	xivtestuser2
Last Name (sn):	xivtestuser2
description:	Storage Administrator
User ID (uid):	xivtestuser2
Password (userPassword):	*****

Figure A-8 Reviewing entry settings

If all the information was entered correctly, you get an “Operation completed Successfully” message on the next panel, shown in Figure A-9. If the operation failed for one reason or another you need to go back and make necessary changes before resubmitting your request.

 **Operation Completed Successfully**
The Entry was successfully created.

Creating New Entry on xivhost2.storage.tucson.ibm.com:389

Creating entry uid=xivtestuser2,dc=xivauth on Directory Server xivhost2.storage.tucson.ibm.com:389
Done.
Operation Completed Successfully

☐ Close window upon completion

Figure A-9 Entry creation confirmation

After the user account was created in SUN Java Directory LDAP its accessibility can be verified using any of the available LDAP clients. In our example (Example A-4), we use the SUN Java Directory LDAP client.

Example: A-4 SUN Java Directory account verification using SUN Java Directory LDAP client

```
$ /opt/sun/dsee6/bin/ldapsearch -b dc=xivauth -h xivhost2.storage.tucson.ibm.com
-D uid=xivtestuser2,dc=xivauth -w pwd2remember uid=xivtestuser2
```

```
dn: uid=xivtestuser2,dc=xivauth
uid: xivtestuser2
description: Storage Administrator
objectClass: inetOrgPerson
objectClass: organizationalPerson
objectClass: person
objectClass: top
sn: xivtestuser2
cn: xivtestuser2
```

The **ldapsearch** command syntax might appear overly complex and its output difficult for interpretation. However, this might be the easiest way to verify that the account was created as expected. The **ldapsearch** command can also be very useful for troubleshooting purposes when you are unable to communicate with Active directory LDAP server.

Here is a brief explanation of the **ldapsearch** command line parameters:

- ▶ **-h xivhost2.storage.tucson.ibm.com** - specifies that the LDAP search query must be sent to “xivhost2.xivhost1ldap.storage.tucson.ibm.com” server using default port 389
- ▶ **-b dc=xivauth** - Base_DN, the location in the Directory Information Tree (DIT).
- ▶ **-D uid=xivtestuser2,dc=xivauth** - the query is issued on behalf of user xivtestuser2 located in dc=xivauth SUN Directory repository.
- ▶ **-w pwd2remember** - is the current password of the user xivtestuser2
- ▶ **uid=xivtestuser2** - specifies what object to search

The output of the “ldapsearch” command shows the structure of the object found. We do not need to describe every attribute of the returned object, however, at least two attributes must be checked to validate the response:

```
uid: xivtestuser2
description: Storage Administrator
```

The fact that **ldapsearch** returns the expected results in our example indicates that:

1. The account is registered in SUN Java Directory.
2. We know where in the SUN Java Directory repository the account is located.
3. We know the valid password.
4. Designated attribute “description” has predefined value assigned “Storage Administrator”.

When the SUN Java Directory account verification is completed we can proceed with configuring XIV System for LDAP authentication mode. At this point we still have a few unassigned LDAP related configuration parameters in our XIV System as can be observed in Example A-5.

Example: A-5 Remaining XIV LDAP configuration parameters

>> ldap_config_get	
Name	Value
base_dn	
xiv_group_attr	description
third_expiration_event	7
version	3
user_id_attr	objectSid
current_server	
use_ssl	no
session_cache_period	
second_expiration_event	14
read_only_role	Read Only
storage_admin_role	Storage Administrator
first_expiration_event	30
bind_time_limit	0

- ▶ **base_dn** - base DN (distinguished name) ,the parameter which specifies where in SUN Java Directory DIT a user can be located. In our example we use “dc=xivauth” as base DN
- ▶ **user_id_attrib** - LDAP attribute set to identify the user (in addition to user name) when recording user operations in the XIV event log. The default value for the attribute is objectSid which is suitable for Active Directory but not for SUN Java Directory LDAP. objectSid attribute is not defined in inetOrgPerson object class used by SUN Java Directory. In our example we set it to uid.
- ▶ **current_server** - is read-only parameter and can not be populated manually.. It will get updated by XIV System after the initial contact with LDAP server is established.
- ▶ **session_cache_period** - duration in minutes XIV System keeps user credentials in its cache before discarding the cache contents. If a user repeats login attempt within session_cache_period minutes from the first attempt - authentication will be done based on the cache content without contacting LDAP server for user credentials.
- ▶ **bind_time_limit** - the timeout value in seconds after which the next LDAP server on the ldap_list_servers is called. Default value for this parameter is 0. It must be set to a non-zero value in order for bind (establishing LDAP connection) to work. The rule also applies to configurations where XIV System is configured with only a single server on the ldap_list_servers list.

The populated values are shown in Example A-6.

Example: A-6 Completing and verifying LDAP configuration on XIV

```
$ xcli -c "ARCXIVJEMT1" -u admin -p s8cur8pwd ldap_config_set base_dn="dc=xivauth"
user_id_attrib=uid session_cache_period=10 bind_time_limit=30
Command executed successfully.
```

```
$ xcli -c "XIV MN00019" -u admin -p s8cur8pwd ldap_config_get
```

Name	Value
base_dn	dc=xivauth
xiv_group_attrib	description
third_expiration_event	7
version	3
user_id_attrib	sid
current_server	
use_ssl	no
session_cache_period	10
second_expiration_event	14
read_only_role	Read Only
storage_admin_role	Storage Administrator
first_expiration_event	30
bind_time_limit	30

To complete our description e of the LDAP related configuration parameters (at the IBM XIV Storage System), let us discuss the parameters that had default values assigned and did not have to be set explicitly. Those are:

- ▶ **version** - version of LDAP protocol used, default “3”. This parameter must never be changed. Both LDAP products - Active Directory and Sun Java Services Directory Server Enterprise Edition support LDAP protocol version 3
- ▶ **user_id_attrib** - LDAP attribute set to identify the user (in addition to user name) when recording user operations in the XIV event log

- ▶ `use_ssl` - indicates if secure (SSL encrypted) LDAP communication is mandated. If set to “yes” without configuring both sides for SSL encrypted communication will result in failing LDAP authentication on XIV System
- ▶ `first_expiration_event` - number of days before expiration of certificate to set first alert (severity “warning”). This parameter must be set to a number of days so it will give you enough time to generate and deploy new security certificate.
- ▶ `second_expiration_event` - number of days before expiration of certificate to set second alert (severity “warning”)
- ▶ `third_expiration_event` - number of days before expiration of certificate to set third alert (severity “warning”)

Securing LDAP communication with SSL

In any authentication scenario, information is exchanged between the LDAP server and XIV system where access is being sought. Security Socket Layer (SSL) can be used to implement secure communications between the LDAP client and server. LDAPS (LDAP over SSL, the secure version of LDAP protocol) allows secure communications between the IBM XIV Storage System and LDAP server with encrypted SSL connections. This allows a setup where user passwords never appear on the wire in clear text.

SSL provides methods for establishing identity using X.509 certificates and ensuring message privacy and integrity using encryption. In order to create an SSL connection the LDAP server must have a digital certificate signed by a trusted certificate authority (CA). Companies have the choice of using a trusted third-party CA or creating their own certificate authority. In this scenario, the `xivauth.org` CA will be used for demonstration purposes.

To be operational SSL, has to be configured on both the client and the server. Server configuration includes generating a certificate request, obtaining a server certificate from a certificate authority (CA), and installing the server and CA certificates.

Windows Server SSL configuration

To configure SSL for LDAP on a Windows Server, you must install the MMC snap-in to manage local certificates, create a certificate request (CER), have CER signed by a CA, import the signed certificate into the local keystore, import a CA certificate as a trusted root CA, and then reboot the server for the new configuration to take effect.

Installation of the local certificates MMC snap-in

Install the certificate snap-in for MMC to allow you to manage the certificates in your local machine keystore. The procedure to install the MMC snap-in is as follows:

1. Start the Management Console (MMC) by selecting **Start**, then **Run**. Type `mmc /a` and select **OK**.
2. Select the **File** → **Add/Remove Snap-in** menu to open the Add/Remove Snap-in dialog.
3. Select the **Add** button to open the Add Standalone Snap-In dialog. Select the **Certificates** snap-in and then click **Add**.
4. Select the **Computer Account** option to manage system-wide certificates. Click **Next** to continue.
5. Select the **Local Computer** option to manage certificates on the local computer only. Select the **Finish**, **Close**, and then click **OK** to complete the snap-in installation.

6. Select **File** → **Save as** and save the console configuration in the %SYSTEMROOT%\system32 directory with a file name of localcert.msc.
7. Create a shortcut in the Administrative Tools folder in your Start menu by right-clicking the **Start** menu and then **Open All Users**. Select the **Program** folder and then the **Administrative Tools** folder.
8. Select the **File** → **New** → **Shortcut** menu. Then enter the location of the saved console, %SYSTEMROOT%\system32\localcert.msc, in the Type the location of the item field. Click **Next** to continue.
9. Enter the name of the new shortcut, Certificates (Local Computer), in the “Type a name for this shortcut” field.
10. To start the local certificate management tool, select **Start** → **Administrative tools** → **Certificates (Local Computer)**.

When the local certificate management tool starts, it will appear as shown in Figure A-10. The certificates used by Active Directory are located in the **Console Root** → **Certificates (Local Computer)** → **Personal** → **Certificates** folder. The list of trusted root certificates authorities is located in the **Console Root** → **Certificates (Local Computer)** → **Trusted Certification Authorities** → **Certificates** folder. See Figure A-10.




Figure A-10 Windows certificate MMC snap-in (local computer)

Generating a Windows Server certificate request

You must use the **certreq** command to generate a certificate request. The **certreq** command uses a text instruction file, which specifies the attributes needed to generate a certificate. It contains attributes such as the subject's common name, certificate key length, and additional key usage extensions. The Active Directory requires that the certificate meet the following requirements:

- ▶ The private key and certificate for the local machine must be imported into the local computer's personal keystore.

- ▶ The fully qualified domain name (FQDN) for the Active Directory must appear
- ▶ in the common name (CN) in the subject field or DNS entry in the subject alternative name extension.
- ▶ The certificate must be issued by a CA that the Active Directory server and the IBM XIV Storage System trust.
- ▶ The certificate must contain the enhanced key usage extension that specifies the server authentication object identifier (OID) 1.3.6.1.5.5.7.3.1. This OID indicates that the certificate will be used as a SSL server certificate.

Example A-7 shows the text instruction file used to generate the certificate for the xivhost1ldap.storage.tucson.ibm.com domain controller. The subject field is set to CN=xivhost1.xivhost1ldap.storage.tucson.ibm.com, which is the FQDN of the domain controller. You then use the **certreq** command to generate the certificate request file.

Example: A-7 Text instruction file for certificate request generation

[Version]

Signature="\$Windows NT\$

[NewRequest]

Subject = "CN=xivhost1.xivhost1ldap.storage.tucson.ibm.com"

KeySpec = 1

KeyLength = 1024

; Can be 1024, 2048, 4096, 8192, or 16384.

; Larger key sizes are more secure, but have

; a greater impact on performance.

Exportable = TRUE

MachineKeySet = TRUE

SMIME = False

PrivateKeyArchive = FALSE

UserProtected = FALSE

UseExistingKeySet = FALSE

ProviderName = "Microsoft RSA SChannel Cryptographic Provider"

ProviderType = 12

RequestType = PKCS10

KeyUsage = 0xa0

[EnhancedKeyUsageExtension]

OID=1.3.6.1.5.5.7.3.1 ; this is for Server Authentication

C:\SSL> certreq -new xivhost1_cert_req.inf xivhost1_cert_req.pem

C:\SSL>

Signing and importing Windows server certificate

After the CER is generated (xivhost1_cert_req.pem), you must send the request to the certificate authority to be signed. For more information about signing this certificate see "Signing a certificate for xivhost1 server" on page 308. After the signed certificate is returned, you must import the certificate into the local machine's personal keystore.

Example A-8 shows how to import the signed certificate using the **certreq** command. Confirm that the certificate is imported correctly by using the **certutil** command.

Example: A-8 Accepting the signed certificate into local certificate keystore

```
C:\>certreq -accept xivhost1_cert.pem
C:\SSL>certutil -store my
===== Certificate 0 =====
Serial Number: 01
Issuer: E=ca@xivstorage.org, CN=xivstorage, O=xivstorage, L=Tucson, S=Arizona,
C=US
Subject: CN=xivhost1.xivhost1ldap.storage.tucson.ibm.com
Non-root Certificate
Cert Hash(sha1): e2 8a dd cc 84 47 bc 49 85 e2 31 cc e3 23 32 c0 ec d2 65 3a
Key Container =
227151f702e7d7b2105f4d2ce0f6f38e_8aa08b0a-e9a6-4a73-9dce-c84e45aec165
Provider = Microsoft RSA SChannel Cryptographic Provider
Encryption test passed
CertUtil: -store command completed successfully.
```

Importing a Certificate Authority certificate

Until the xivstorage.org CA is designated as a trusted root, any certificate signed by that CA will be untrusted. You must import the CA's certificate, using the local certificate management tool, into the Trusted Certification Authorities folder in the local keystore.

To start the local certificate management tool, select **Start** → **Administrative tools** → **Certificates (Local Computer)**:

1. After the certificate tool opens, select the **/Console Root/Certificates (Local Computer)/Trusted Certification Authorities** folder.
2. Start the certificate import wizard by selecting the **Action** → **All Tasks** → **Import** menu. Click **Next** to continue.
3. Select the file you want to import. The xivstorage.org CA certificate is located in the cacert.pem file. Click **Next** to continue.
4. Select the **Place all certificates in the following store** option and make sure the certificate store field is set to **Trusted Root Certification Authorities**. Click **Next** to continue.
5. The CA certificate is now imported. Click **Finish** to close the wizard.

After the CA and server certificates are imported into the local keystore, you can then use the local certificate management tool to check whether the certificates are correctly imported. Open the **Console Root** → **Certificates (Local Computer)** → **Personal** → **Certificates** folder and select the certificate issued to xivhost1.xivhost1ldap.storage.tucson.ibm.com.

Figure A-11 shows that the certificate which was issued to xivhost1.xivhost1ldap.storage.tucson.ibm.com is valid and was issued by the xivstorage CA. The certificate has a corresponding private key in the keystore. The “Ensures the identity of the remote computer” text indicates that the certificate has the required server authentication key usage defined.

To check the xivstorage certificate, open the **Console Root** → **Certificates (Local Computer)** → **Trusted Certification Authorities** → **Certificates** folder and select the certificate issued by xivstorage. Figure A-12 shows that the certificate issued to and by the xivstorage CA is valid.




Figure A-11 Certificate information dialog




Figure A-12 Certificate information dialog for xivstorage certificate authority

Low-level SSL validation using the openssl command

The easiest way to test the low-level SSL connection to the LDAP server is by using the `openssl s_client` command with the `-showcerts` option. This command will connect to the specified host and list the server certificate, the certificate authority chain, supported ciphers, SSL session information, and verify return code. If the SSL connection worked, the `openssl s_client` command result in the verify return code will be 0 (Ok).

Example A-9 shows the output of the `openssl s_client` command connecting Linux server (xivstorage.org) to the Active Directory server `xivhost1.xivhost1ldap.storage.tucson.ibm.com`. This command connects to the Active Directory server using the secure LDAP port (636).

Example: A-9 Low-level SSL validation using the openssl s_client

```
openssl s_client -host xivhost1.xivhost1ldap.storage.tucson.ibm.com -port 636
-Cafile cacert.pem -showcerts
...
---
Server certificate
subject=/CN=xivhost1.xivhost1ldap.storage.tucson.ibm.com
issuer=/C=US/ST=Arizona/L=Tucson/O=xivstorage/CN=xivstorage/emailAddress=ca@xivsto
rage.org
---
...
New, TLSv1/SSLv3, Cipher is RC4-MD5
Server public key is 1024 bit
SSL-Session:
    Protocol  : TLSv1
    Cipher    : RC4-MD5
    Session-ID: 9E240000CE9499A4641F421F523ACC347ADB91B3F6D3ADD5F91E271B933B3F4F
    Session-ID-ctx:
    Master-Key:
F05884E22B42FC4957682772E8FB1CA7772B8E4212104C28FA234F10135D88AE496187447313149F2E
89220E6F4DADF3
    Key-Arg   : None
    Krb5 Principal: None
    Start Time: 1246314540
    Timeout   : 300 (sec)
    Verify return code: 0 (ok)
```

Attention: In order to complete the configuration of SSL for the Active Directory, you must reboot the Windows server.

Basic secure LDAP validation using the ldapsearch command

After you have confirmed that the SSL connection is working properly, you must confirm that you are able to search your LDAP directory using LDAP on a secure port. This will confirm that the LDAP server can communicate using an SSL connection.

In our example we use OpenLDAP client for SSL connection validation. CA certificate needs to be added to the key ring file used by OpenLDAP client. `TLS_CERTS` option in OpenLDAP configuration file (typically `/etc/openldap/ldap.conf`) specifies the file that contains certificates for all of the Certificate Authorities the client will recognize. See Example A-10.

Example: A-10 Testing LDAP over SSL using ldapsearch command

```
# /usr/bin/ldapsearch -x -H
"ldaps://xivhost1.xivhost1ldap.storage.tucson.ibm.com:636" -D
'CN=xivtestuser1,CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com' -w
pass2remember -b 'CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com'

dn: CN=xivtestuser1,CN=Users,DC=xivhost1ldap,DC=storage,DC=tucson,DC=ibm,DC=com
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: user
cn: xivtestuser1
description: Storage Administrator
distinguishedName: CN=xivtestuser1,CN=Users,DC=xivhost1ldap,DC=storage,DC=tuc
on,DC=ibm,DC=com
...
# search result
search: 2
result: 0 Success
```

The URI format used with “-H” option specifies that LDAPS (LDAP over SSL) must be used on port 636 (LDAP secure port).

SUN Java Directory SSL configuration

This section illustrates the use of an SSL protocol for communication with the SUN Java Directory.

Creating SUN Java Directory certificate request

To configure SSL for SUN Java LDAP Directory, you must create a certificate request (CER), have CER signed by a CA, import the signed certificate into the local keystore, import a CA certificate as a trusted root CA, and then restart the LDAP server for the new configuration to take effect.

Generating a SUN Java Directory server certificate request

To generate a certificate request using SUN Java Web Console tool:

- ▶ Point your web browser HTTPS port 6789, in our example “https://xivhost2.storage.tucson.ibm.com:6789”.
- ▶ Login to the system and select “Directory Service Control Center (DSCC)” application and Authenticate to Directory Service Manager.
- ▶ Select **Directory Servers** → **xivhost2.storage.tucson.ibm.com:389** → **Security** → **Certificates** → **Request CA-Signed Certificate**. Fill out the certificate request form. Sample of certificate request form is shown in Figure A-13.

Request CA-Signed Certificate

This will generate a certificate request. The text of the request will appear on the progress dialog. This text should be submitted to a Certificate Authority who will process it and issue a certificate. You can submit the request either by sending the text in an e-mail message to the CA or by submitting it through the CA's web site.

* Indicates required field

Server: xivhost2.storage.tucson.ibm.com:389

Certificate Details: ☒ Specify Values Separately:

* Common Name (cn): xivhost2.storage.tucson.ibm.com

Organization (o): xivstorage

Organizational Unit (ou): ITSO

City/Locality (l): Tucson

State/Province (st): Arizona

Country (c): US

☐ Specify as Subject DN:

* Subject DN:

Figure A-13 Certificate request

Copy the generated certificate shown in Figure A-14 request into xivhost2_cert_req.pem file.

Operation Completed Successfully
 The Request was successfully generated.

Generating the CA-Signed Certificate Request

Close

Generating the CA-Signed Certificate Request for xivhost2.storage.tucson.ibm.com:389 Done.

The request is:
 Certificate request generated by Sun-Java(tm)-System-Directory/6.0

Common Name: xivhost2.storage.tucson.ibm.com
 Email: (not specified)
 Phone: (not specified)
 Organization: xivstorage
 State: Arizona
 Country: US

```

-----BEGIN NEW CERTIFICATE REQUEST-----
MIIbVjCCAScCAQAwFjELMAkGA1UEBhMCVVMxEADA0BgNVBAgTB0FyaXpvcmlExDzANBgNVBAcTB1R1Y3NybWJEMAsGA1UECjMESEVVRTTzETMBEGA1UEChMKc2c3RvcmlFZTEoMCYGA1UEAxMFeG12aG9zdDIuc3RvcmlFZS50dWZb24uaWJtLmNvbCCEBgkqhkiG9w0BAQEEFAAOBjQAwgYKCGYEAoLxzk/ay16NDT1LQWdMmQhREJTfXBL+BlvvtFItMlqv8ub+haouH/gZc8rEPIyHAHtvvrthwu+LWQJ5ZrYLP6ESEMU31KwAsMEQLZF+dyUQ3gtATCisZZ1/IdPgOXW25wYQ1nfjJWsdDN2sbkMYXPK6P53R/MCAwEAAaAAMAOGCSqGSIb3DQEBBAAUAA4GBAIKxs13kKwIutw+rIMgmPUo8pvxF1fXpTyQM0I+ooKscT8RksAwzcWFOXQg3beMH+FWcuF2Qg18Gkwc69n/694nRXc81Rca999G6EgTZOM7e5cFIn+DKyExvemKL5YTfSvm5SK0F6n5XZp2xnjd/UaTXy1COBUu
-----END NEW CERTIFICATE REQUEST-----

```

You must copy the request and send it to the Certificate Authority to obtain your certificate.

Operation Completed Successfully


Figure A-14 Generated certificate

Signing and importing a server certificate

After the CER is generated (xivhost2_cert_req.pem), you must send the request to the certificate authority to be signed. For more information about signing this certificate, see “” on page 308. After the signed certificate xivhost2_cert.pem file is returned, you must import the certificate into the local machine’s personal keystore.

To add signed certificate in “Directory Service Manager” manager application, select **Directory Servers** → **xivhost2.storage.tucson.ibm.com:389** → **Security** → **Certificates** → **Add**.

Copy and paste certificate stored in xivhost2_cert.pem file as shown in Figure A-15.



Add Certificate

Paste the certificate into the text field below. The certificate text must be in ASCII format.

Server: xivhost2.storage.tucson.ibm.com:389

* Certificate Name: xivstorage.org sample CA certificate

* Certificate:

```
Certificate:
Data:
  Version: 3 (0x2)
  Serial Number: 2 (0x2)
  Signature Algorithm: md5WithRSAEncryption
  Issuer: C=US, ST=Arizona, L=Tucson, O=xivstorage, CN=xivstorage/emailAddress=ca@xivstorage.org
  Validity
    Not Before: Jul  2 22:48:43 2009 GMT
    Not After : Jul  2 22:48:43 2010 GMT
  Subject: C=US, ST=Arizona, L=Tucson, O=xivstorage, OU=ITSO, CN=xivhost2.storage.tucson.ibm.com
  Subject Public Key Info:
```

Figure A-15 Adding signed certificate

Importing a Certificate Authority certificate

Until the xivstorage.org CA is designated as a trusted root, any certificate signed by that CA will be untrusted. You must import the CA’s certificate using “Directory Service Manager” manager application by selecting **Directory Servers** → **xivhost2.storage.tucson.ibm.com:389** → **Security** → **CA Certificates** → **Add**.

Add Certificate

Paste the certificate into the text field below. The certificate text must be in ASCII format.

* Indicates required field

Server:

xivhost2.storage.tucson.ibm.com:389

*

Certificate Name:

xivstorage.org sample CA Certificate Authority certificate

*

Certificate:

—BEGIN CERTIFICATE—
MIIDSTCCArKgAwIBAgIBADANBgkqhkiG9w0BAQQFADB8MQswCQYDVQQGEwJVUzEQ
MA4GA1UECBMHQXJpem9uYTEPMAOGA1UEBxMGVHVjc29uMRMwEQYDVQQKEwp4aXZz
dG9yYWdlMRMwEQYDVQQDEWp4aXZzdG9yYWdlMSAwHgYJKoZIhvcNAQkBFhFjYUB4
aXZzdG9yYWdlLm9yZzAeFw0wOTA2MjkvMTI0MjNaFw0xMDA2MjkvMTI0MjNaMHwx
CzajBgNVBAYTAiVTMTRAwDgYDVOQIEwdBcmI6b25hMQ8wDOYDVQQHEwZUdWNzb24x
EzARBgNVBAoTChdpdnN0b3JhZ2UxExARBgNVBAMTCnhpdnN0b3JhZ2UxIDAeBgkq
hkiG9w0BCQEWEWNhQHhpdnN0b3JhZ2Uub3JnMIGfMA0GC SqSIlb3DQEBQUAA4GN
ADCBIQKbgODWXxVtuYtPlafRLlGgq24KdECZYtt+dLSqJPBhcCYcvAtzTnzanzRZj
jl1Z13dtOEYnj8NvpSpVNF62eG9qccASo/N+JgXC9+idSPHWKpKYmm9kmGdLFW6
TA1BEJ/EffkbortCNxeStBeWmmioSAahDyPDxZqNo2H2+M1/V3/PCwIDAQABo4Ha

Figure A-17 shows that the certificate issued to `xivhost2.storage.tucson.ibm.com` is valid and was issued by the `xivstorage` Certificate Authority.

General

Issued To:	CN=xivhost2.storage.tucson.ibm.com, OU=ITSO, O=xivstorage, L=Tucson, ST=Arizona, C=US
Issued By:	EMAILADDRESS=ca@xivstorage.org, CN=xivstorage, O=xivstorage, L=Tucson, ST=Arizona, C=US
Valid From:	7/2/09 3:48 PM
Expires On:	7/2/10 3:48 PM
Type:	X.509
Serial Number:	2
Signature:	[B@3126d4
Signature Algorithm:	MD5withRSA
Public Key:	

```
Sun RSA public key, 1024 bits
modulus: 114013868825707066374043683934473885929025216935889507771736354172179899887955529768116176
418109480816200394446098553506861611177281191517927909278808722357447377156239184381783246525075007
218852949411716333024299916525339877472184384252011146708924534136193250322949095272090889371573877
142807791146638985203
public exponent: 65537
```

Version: 3

To check the xivstorage certificate, open the **Directory Servers** → **xivhost2.storage.tucson.ibm.com:389** → **Security** → **CA Certificates** and click the link, **xivstorage.org sample CA Certificate Authority certificate**. Figure A-18 shows that the certificate issued to and by the xivstorage CA is valid.

General	
Issued To:	EMAILADDRESS=ca@xivstorage.org, CN=xivstorage, O=xivstorage, L=Tucson, ST=Arizona, C=US
Issued By:	EMAILADDRESS=ca@xivstorage.org, CN=xivstorage, O=xivstorage, L=Tucson, ST=Arizona, C=US
Valid From:	6/29/09 2:24 PM
Expires On:	6/29/10 2:24 PM
Type:	X.509
Serial Number:	0
Signature:	[B@15b0bc6
Signature Algorithm:	MD5withRSA
Public Key:	<pre> Sun RSA public key, 1024 bits modulus: 150536730995091067593670568490000314918699882632837962817648888644808296689992063936025114 118715826260207280255942458541747085068940168153850225658409003397463038366659635042876751283164927 321393827597558264885346828836943619731303484929767031031717008539702750869126516194155733126618341 652906240398868860683 public exponent: 65537 </pre>
Version:	3

Figure A-18 Certificate information for xivstorage certificate authority

To activate imported certificate open the **Directory Servers** → **xivhost2.storage.tucson.ibm.com:389** → **Security**. In General tab in “Certificate” field expand scroll down list and select xivstorage.org sample CA certificate as shown in Figure A-19.

Directory Servers > xivhost2.storage.tucson.ibm.com:389

Server Operation	Suffixes	Entry Management	Schema	Security	Server Configuration
General	Certificates	CA Certificates			

xivhost2.storage.tucson.ibm.com:389 - Security

SSL Settings

SSL Encryption: ☒ Enabled – [Configure LDAP Secure Port](#) | [Configure DSML Secure Port](#)

Certificate: Default Certificate

Cipher Family: Default Certificate
xivstorage.org sample CA certificate

Security Device: Internal (software)

Client Authentication

LDAP Settings: Allow Certificate-Based Client authentication

DSML Settings: Try to use client certificate first

Figure A-19 New signed certificate activation

As depicted in Figure A-20, you will be prompted to restart the LDAP server in order for the new certificate to take effect


Restart Required

You have modified the Certificate used by the Directory Server. You must restart the Directory Server in order this modification to be taken into account.

Note: Once the Java(TM) Web Console session has ended this message will disappear.

Figure A-20 Manual restart request after activating new certificate

Low-level SSL validation using the openssl command

The easiest way to test the low-level SSL connection to the LDAP server is by using the `openssl s_client` command with the `-showcerts` option. This command will connect to the specified host and list the server certificate, the certificate authority chain, supported ciphers, SSL session information, and verify return code. If the SSL connection worked, the `openssl s_client` command result in the verify return code will be 0 (Ok).

Example A-11 shows the output of the **openssl s_client** command connecting Linux server (xivstorage.org) to the SUN Java Directory server `xivhost2.storage.tucson.ibm.com`. This command connects to the SUN Java Directory server using the secure LDAP port (636).

Example: A-11 Low-level SSL validation using the openssl s_client

```
# openssl s_client -host xivhost2.storage.tucson.ibm.com -port 636 -CAfile
cacert.pem -showcerts
...
---
Server certificate
subject=/C=US/ST=Arizona/L=Tucson/O=xivstorage/OU=ITS0/CN=xivhost2.storage.tucson.
ibm.com
issuer=/C=US/ST=Arizona/L=Tucson/O=xivstorage/CN=xivstorage/emailAddress=ca@xivsto
rage.org
---
Acceptable client certificate CA names
/O=Sun Microsystems/CN=Directory Server/CN=636/CN=xivhost2.storage.tucson.ibm.com
/C=US/ST=Arizona/L=Tucson/O=xivstorage/CN=xivstorage/emailAddress=ca@xivstorage.or
g
---
SSL handshake has read 2144 bytes and written 328 bytes
---
New, TLSv1/SSLv3, Cipher is AES256-SHA
Server public key is 1024 bit
Compression: NONE
Expansion: NONE
SSL-Session:
    Protocol : TLSv1
    Cipher    : AES256-SHA
    Session-ID: 48B43B5C985FE1F6BE3F455F8350A4155DD3330E6BD09070DDCB80DCCB570A2E
    Session-ID-ctx:
    Master-Key:
1074DC7ECDD9FC302781C876B3101C9C618BB07402DD7062E7EA3AB794CA9C5D1A33447EE254288CEC
86BBB6CD264DCA
    Key-Arg   : None
    Krb5 Principal: None
    Start Time: 1246579854
    Timeout   : 300 (sec)
    Verify return code: 0 (ok)
---
```

Basic secure LDAP validation using the ldapsearch command

After you have confirmed that the SSL connection is working properly, you must verify that you are able to search your LDAP directory using LDAPS on port 636. This will confirm that the LDAP server can communicate using SSL connection.

In Example A-12, we use OpenLDAP client for SSL connection validation. CA certificate needs to be added to key ring file used by OpenLDAP client. TLS_CERTS option in OpenLDAP configuration file (typically /etc/openldap/ldap.conf) specifies the file that contains certificates for all of the Certificate Authorities the client will recognize.

Example: A-12 Testing LDAP over SSL using ldapsearch command

```
# /usr/bin/ldapsearch -x -H "ldaps://xivhost2.storage.tucson.ibm.com:636" -D
'uid=xivtestuser2,dc=xivauth' -w pwd2remember -b 'dc=xivauth'

# extended LDIF
#
# LDAPv3
# base <dc=xivauth> with scope subtree
# filter: uid=xivtestuser2
# requesting: ALL
#

# xivtestuser2, xivauth
dn: uid=xivtestuser2,dc=xivauth
uid: xivtestuser2
objectClass: inetOrgPerson
objectClass: organizationalPerson
objectClass: person
objectClass: top
sn: xivtestuser2
cn: xivtestuser2
description: custom_role_01

# search result
search: 2
result: 0 Success
```

The URI format used with “-H” option specifies that LDAPS (LDAP over SSL) must be used on port 636 (LDAP secure port).

Certificate Authority setup

This section describes the setup and use of the certificate authority that was used with all example scenarios in this book to issue certificates.

OpenSSL comes with most Linux distributions by default. Information about OpenSSL can be found at the OpenSSL website:

<http://www.openssl.org>

Creating the CA certificate

To set up the CA for the xivstorage.org domain we need to make certain assumptions. We modify the openssl.cnf to reflect these assumptions to the CA. The file can be found at /usr/share/ssl/openssl.cnf and the interesting sections are shown in Example A-13.

```
[ CA_default ]

dir           = /root/xivstorage.orgCA# Where everything is kept
certs        = $dir/certs             # Where the issued certs are kept
crl_dir       = $dir/crl               # Where the issued crl are kept
database      = $dir/index.txt        # database index file.
new_certs_dir = $dir/newcerts         # default place for new certs.
certificate    = $dir/cacert.pem      # The CA certificate
serial        = $dir/serial           # The current serial number
crl           = $dir/crl.pem          # The current CRL
private_key    = $dir/private/cakey.pem# The private key
RANDFILE      = $dir/private/.rand    # private random number file
x509_extensions = usr_cert           # The extensions to add to the cert
name_opt       = ca_default           # Subject Name options
cert_opt       = ca_default           # Certificate field options
default_days   = 365                  # how long to certify for
default_crl_days= 30                  # how long before next CRL
default_md     = md5                  # which md to use.
preserve      = no                    # keep passed DN ordering
.
.
.

[ req_distinguished_name ]
countryName          = Country Name (2 letter code)
countryName_default  = US
countryName_min      = 2
countryName_max      = 2

stateOrProvinceName  = State or Province Name (full name)
stateOrProvinceName_default = TX
localityName         = Locality Name (eg, city)
localityName_default = Tucson
0.organizationName   = Organization Name (eg, company)
0.organizationName_default = xivstorage
organizationalUnitName = Organizational Unit Name (eg, section)
commonName           = xivstorage.org (eg, your server's hostname)
commonName_max       = 64
emailAddress         = ca@xivstorage.org
emailAddress_max     = 64
```

Also, the directories to store the certificates and keys must be created:

```
mkdir /root/xivstorage.orgCA /root/xivstorage.orgCA/certs
/root/xivstorage.orgCA/crl /root/xivstorage.orgCA/newcerts
/root/xivstorage.orgCA/private
```


OpenSSL is using a couple of files, which it uses to maintain the CA. These files must be created:

```
touch /root/xivstorage.orgCA/index.txt
```

```
echo "01" >> /root/xivstorage.orgCA/serial
```

The access rights on the directories and files need to be reviewed to restrict access to the CA and, most importantly, to the private key as far as possible.

To create the CA certificate certified for 365 days, the OpenSSL command is issued directly, as shown in Example A-14.

Example: A-14 Generating the CA certificate

```
openssl req -new -x509 -days 365 -keyout /root/xivstorage.orgCA/private/cakey.pem
-out /root/xivstorage.orgCA/cacert.pem
Generating a 1024 bit RSA private key
....+++++
.....+++++
writing new private key to '/root/xivstorage.orgCA/private/cakey.pem'
Enter PEM pass phrase:
Verifying - Enter PEM pass phrase:
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [US]:
State or Province Name (full name) [Arizona]:
Locality Name (eg, city) [Tucson]:
Organization Name (eg, company) [xivstorage]:
Organizational Unit Name (eg, section) []:
Common Name (eg, YOUR name) []:xivstorage
Email Address []:ca@xivstorage.org
```

During the creation of the certificate missing information must be provided. Also, the information that has been defined as default in the openssl.cnf file must be confirmed. The password for the CA private key must be given during the creation process. This password is needed whenever the CA's private key is used. The following command can be used to view CA certificate: "openssl x509 -in cacert.pem -text"

Signing a certificate

The client or server that needs to get a certificate must create a certificate signing request and send this request to the CA.

Certificate request details can be viewed using the following command:

```
openssl req -in xivhost1_cert_req.pem -text
```

xivhost1_cert_req.pem is the certificate signing request, the file generated on xivhost1.xivhost1dap.storage.tucson.ibm.com server

Signing a certificate for xivhost1 server

To sign the certificate, the **openssl** command is used with a specified policy, as shown in Example A-15.

Example: A-15 Signing certificate for xivhost1 server

```
# openssl ca -policy policy_anything -cert cacert.pem -keyfile private/cakey.pem
-out xivhost1_cert.pem -in xivhost1_cert_req.pem
Using configuration from /usr/share/ssl/openssl.cnf
Enter pass phrase for private/cakey.pem:
Check that the request matches the signature
Signature ok
Certificate Details:
    Serial Number: 1 (0x1)
    Validity
        Not Before: Jun 29 21:35:33 2009 GMT
        Not After : Jun 29 21:35:33 2010 GMT
    Subject:
        commonName =
xivhost1.xivhost1ldap.storage.tucson.ibm.com
    X509v3 extensions:
        X509v3 Basic Constraints:
            CA:FALSE
        Netscape Comment:
            OpenSSL Generated Certificate
        X509v3 Subject Key Identifier:
            C8:EB:8D:84:AB:86:BB:AF:5B:74:4D:35:34:0E:C5:84:30:A1:61:84
        X509v3 Authority Key Identifier:
            keyid:A8:0B:D1:B5:D6:BE:9E:61:62:E3:60:FF:3E:F2:BC:4D:79:FC:E3:5A

DirName:/C=US/ST=Arizona/L=Tucson/O=xivstorage/CN=xivstorage/emailAddress=ca@xivstorage.org
serial:00

X509v3 Extended Key Usage:
    TLS Web Server Authentication
X509v3 Key Usage:
    Digital Signature, Key Encipherment
Certificate is to be certified until Jun 29 21:35:33 2010 GMT (365 days)
Sign the certificate? [y/n]:y

1 out of 1 certificate requests certified, commit? [y/n]y
Write out database with 1 new entries
Data Base Updated
```

Related publications

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

IBM Redbooks publications

For information about ordering this publication, see “How to get IBM Redbooks publications” on page 310. The following documents might be available in softcopy only:

- ▶ *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904
- ▶ *IBM TXIV Storage System: Copy Services and Data Migration*, SG24-7759

Other publications

These publications are also relevant as further information sources:

- ▶ *IBM XIV Storage System Application Programming Interface*, GA32-0788
- ▶ *IBM XIV Storage System User Manual*, GC27-2213
- ▶ *IBM XIV Storage System: Product Overview*, GA32-0791
- ▶ *IBM XIV Storage System Planning Guide*, GA32-0770
- ▶ *IBM XIV Storage System Pre-Installation Network Planning Guide for Customer Configuration*, GC52-1328-01
- ▶ *The iSCSI User Guide*:
<http://download.microsoft.com/download/a/e/9/ae91dea1-66d9-417c-ade4-92d824b871af/uguide.doc>

Online resources

These websites are also relevant as further information sources:

- ▶ IBM XIV Storage System Information Center:
<http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp>
- ▶ IBM XIV Storage 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>
- ▶ SNIA (Storage Networking Industry Association) website:
<http://www.snia.org/>
- ▶ IBM Director Software Download Matrix page:
<http://www.ibm.com/systems/management/director/downloads.html>
- ▶ IBM Director documentation:
<http://www.ibm.com/systems/management/director/>

How to get IBM Redbooks publications

You can search for, view, or download IBM Redbooks publications, Redpapers publications, Technotes, draft publications, and Additional materials, as well as order hardcopy IBM Redbooks publications, at this website:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

A

Access Control List (ACL) 7
ACL 7
Active Directory (AD) 6, 149, 151, 166, 168, 174, 179, 181–182, 188–189, 279–281, 283, 285, 287–288, 291–292, 294, 298
 LDAP server 180
 user 280–281
admin 91, 101, 133, 135–136, 138–139, 146–147, 152, 167, 175–176, 181–182, 184–185, 194, 292
alerting event 213
Application Administrator 134, 137, 139, 141, 169, 176
applicationadmin 134–135, 142, 169–171
applicationadmin role 137, 139, 152, 154
Automatic Transfer Switch (ATS) 3, 51
autonomic features 36
availability 14

B

bandwidth 17–18, 48, 57, 64
basic configuration 67
battery 52, 211
block 20–21, 113, 118, 124
block-designated capacity 21
boot device 58
buffer 61

C

CA certificate 187, 293, 301–302
cache 14–15, 48, 57
 buffer 61
 growth 17
 page 199
caching
 management 198
call home 203, 207
capacity 83, 85, 89
 depleted 104
 unallocated 26, 28, 31
category 133–134, 169
CE/SSR 70, 78–79, 276
Certificate Request Copy 300–301, 306
cg_move 111–112
CIM agent 215, 221
 directory look-up 217
CIMOM discovery 217
Command Line Interface 84
Common Information Model Object Manager (CIMOM) 215
Common Name (CN) 174, 180, 182, 288, 294, 307
Compact Flash Card 58
component_list 210
computing resource 16

configuration flow 89
Consistency Group 25, 43, 109, 111–112, 114, 116, 122–123, 192
 application volumes 25
 special snapshot command 25
context menu 126–127, 138–141, 145, 176–177, 191
 Change Password 138, 145
 select events 191
cooling 57
copy on write 19
created Storage Pool
 actual size 105
Cyclic Redundancy Check (CRC) 59

D

daemon 247
Data Collection 276
data distribution algorithm 49
data integrity 25, 44
data migration 6, 20, 44, 60, 192
Data Module 15–16, 37, 39, 47, 49, 55–56, 59, 198
 separate level 37
data redundancy 19–20, 35, 39–40
 individual components 35
data stripe 199
default IP
 address 92
 gateway 75
default password 133, 136, 138
default value 106, 136, 284, 292, 307
deletion priority 26, 34
demo 8
demo mode 8, 85
depleted capacity 106, 112
depletion 25
description attribute 152, 170
destage 25, 38
destination 242, 249–250
destination type 242
detailed information 68, 84, 103, 205–206, 209, 215, 218, 251, 256
DIMM 57
Director Console 253
Directory Information Tree (DIT) 148–149, 174, 179, 183, 287
Directory Service Control Center (DSCC) 183, 286, 299
dirty data 36, 38, 45
disaster recovery 36, 43
disk drive 37, 48, 55, 57, 61, 103, 198–199, 276
 reliability standards 44
 transient, anomalous service time 45
disk scrubbing 44
disk_list 210
distinguished name (DN) 148, 283, 292, 307

- distribution 39–40
- distribution table 22
- distribution 39
- DNS 240, 250
- Domain Name
 - Server 155
 - System 69, 250
- Dynamic Host Configuration Protocol (DHCP) 77

E

- email address 69, 136, 148, 241, 256
- email notification 90
- emergency power off (EPO) 82
- enclosure management card 57
- encrypted SSL connection
 - LDAP server 293
- EPO 82
- Ethernet fabric 14–15
- Ethernet network 15
- Ethernet port 56, 59, 76
 - following configuration information 75
- Ethernet switch 3, 15, 48, 51, 64, 75
- event 206, 208, 213
 - alerting 213
 - severity 204, 206–207, 213
- event_list 213
- event_list command 194, 213

F

- fan 57
- FBDIMM 57
- Fibre Channel
 - cabling 75
 - configuration 71–72, 75
 - connection 70
 - connectivity 59, 70
 - host access 70
 - host I/O 70
 - network 75
 - parameter 70
 - ports 51, 59, 70
 - switch 72, 75
- Fibre Channel (FC) 6, 15, 17, 59, 70–71, 74–76, 78
- Field-Replaceable Unit (FRU) 58, 268
- filter 227–229
- Fluid Dynamic Bearing (FDB) 62
- free capacity 107–108
- Front Server 269–270
- full installation 87
- Fully Buffered DIMM (FBDIMM) 57
- fully qualified domain name (FQDN) 217, 295
- Function icons 93

G

- Gateway 60, 91, 238–239, 245
- GHz clock 57
- Gigabit Ethernet 3
- GigE adapter 56

- goal distribution 23–24, 39–42
 - Full Redundancy 24
 - priority 39
- graceful shutdown 38
- Graphical User Interface (GUI) 1, 8–9, 84, 89–95, 97, 99–100, 103–104, 110, 114, 119, 121–122, 130, 134, 136, 139, 145–146, 150, 152, 169, 173, 175–176, 178, 181, 185–186, 188–189, 191, 195, 197, 204, 213, 222, 224, 226, 228–229, 237, 245, 247, 249, 270
- grid architecture 14, 16
- grid topology 17, 36, 38
- GUI 84–85
 - demo mode 8

H

- hard capacity 83, 85, 89
 - depletion 35
- hard pool size 28–29
- hard size 103, 108, 117
- hard space 25, 27
- hard storage capacity 94, 205
- hard system size 30, 34
- hard_size 112
- hardware 47–48, 50, 67
- high availability 35, 43
- host bus adapter (HBA) 72
- host connectivity 60, 71
- Host Rate Limiting 7
- host server
 - example power 128
- host transfer size 200
- hot-spot 22
- hotspots 1

I

- IBM development and service 133
 - internal IDs 133
- IBM Director 247, 249–254, 256, 258–260
 - components 251
 - MIB file 251–252, 257–258
- IBM Intranet 77, 270
 - support person 77
- IBM SSR 69–70, 76, 78–79, 267, 276
- IBM XIV 48–49, 51–53, 61–62, 65–71, 74, 76–79, 205, 207–208, 212, 215, 224, 247, 252, 258, 267, 269–270, 275–276
 - development team 133
 - final checks 79
 - hardware 65
 - hardware component 51
 - installation 78
 - internal network 51
 - line cord 79
 - maintenance team 133
 - personnel 77
 - power components 51
 - power connector 79
 - remote support 69
 - remote support center 267

- repair 65
- SATA disks 62
- software 5, 247, 267
- Storage Manager 8, 99
- Storage Manager Software 99
- Storage System 1, 3, 5, 10, 66, 68–71, 133, 139, 150, 224
- Storage System Information Center 99
- Storage System open architecture 5
- Storage System software 45, 70
- Support 65, 69, 78, 267, 276
- Support Center 65
- technician 39
- XCLI User Manual 212
- IBM XIV Storage
 - Manager 7, 84, 130
 - Manager GUI 8
 - Manager window 139
 - System 76
- Intel Xeon 56
- interface 17
- Interface Module 15–18, 37, 39, 47, 49, 55–56, 59–60, 63, 67, 69–70, 72, 74, 84, 90, 200
 - GigE adapters 55
 - host requests 37
 - Maximum number 71
- IOPS 226–228
- IP address 60, 69, 75–76, 90, 92, 100, 217, 240, 247, 286
- IP host 247–248
- IQN 126
- iSCSI 59–60
 - initiator 60
 - ports 59–60
 - target 60
- iSCSI connection 75, 126, 129
- iSCSI host
 - port 71
- iSCSI Port 51, 59, 71, 210
 - Interface Module 59
 - Maximum number 71
- ITSO Pool 111–112, 123–124

J

- just-in-time 103

K

- KB 20

L

- LAN subnet 92
- latency 227
- LDAP 6
- LDAP administrator 151, 153–154, 174, 286–287
- LDAP client 148, 187, 286
 - secure communications 293
- LDAP communication 131, 187, 190, 279, 285, 287, 293
- LDAP directory

- Information Tree 179, 183
- server 148
- structure 287
- LDAP entry
 - creation 286
 - login 286
- LDAP object 151, 170, 180–181, 184, 283, 285, 288
 - class 151, 288
- LDAP role 140, 150, 154, 171–172, 176–178, 181–182, 185
 - mapping 150, 166, 286, 288
- LDAP role mapping
 - process 150, 169–170
- LDAP server 147–151, 166–167, 169, 172–173, 175, 180, 184, 187–188, 190, 283–284, 286–287, 291–293, 298–299, 303–304
 - multiple instances 287
- LDAP user 151, 167, 169–174, 180, 182, 184, 186
 - User group membership 170
- ldap_config_set 151
- ldap_role parameter 170–171
- ldapsrch command 181, 184, 283, 291, 298–299, 304
 - line parameter 283, 291
 - syntax 283
- Least Recently Used (LRU) 198
- left pane 128
- Lightweight Directory Access Protocol (LDAP) 6, 147
- load balancing 16
- local computer 293–294, 296
- local keystore 293, 296, 299, 302
- locking 34
- logical structure 19
- logical unit number (LUN) 78, 122, 173, 194
- logical volume 4, 18, 21
 - associated group 24
 - hard space 29
 - layout 18
 - placement 18
 - related group 102
 - size 27
- logical volume (LV) 18, 24, 102, 113
- Logical Volume Manager (LVM) 200
- Logical volume size 27
- logical volume size 27

M

- machine type 3, 48, 67
- MacOS 84, 90
- Main display 93
- main GUI
 - management window 110
 - window 191, 204
- Maintenance Module 3, 48, 53, 65
- Management Information Base (MIB) 247–248
- Management Tools 7
- management workstation 84, 90, 92
- mapping 19–20, 89, 99, 122
- master volume 19, 25, 109, 199
- maximum number 22, 70
- Maximum Transmission Unit (MTU) 60, 75

- MB partition 22, 27, 198
- memory 57
- Menu bar 94
- metadata 15, 56
- metrics 226, 228, 236
- MIB 248
- MIB extensions 249
- MIB file 250–252, 257–258
- Microsoft Active Directory (AD) 6
- migration 118
- modem 3, 65, 268
- module_list 210
- modules 14
- monitor
 - statistics 209
- monitoring 203–205
- MTU 60, 91
- multiple system 136, 145–147, 186
- multivalued attribute 181, 184

N

- NDCL 6, 45
- Network mask 91
- Network Time Protocol (NTP) 69, 79
- Non-Disruptive Code Load (NDCL) 6, 45
- Non-disruptive code load (NDCL) 45

O

- Object class
 - ePerson 148
- On-site Repair 276
- OpenLDAP client 282, 298, 305
- openssl s_client
 - command 298, 304
 - command result 298, 304
- orphaned space 23

P

- parallelism 14, 16–17, 49, 197, 200
- partition 20
- pass2remember ldap_user_test 182, 185
- patch panel 59–60, 63, 70–72, 76, 79, 204
 - IP network 60
- PCI Express 57
- PCI-e 198
- performance 19, 197–198
 - metrics 226, 228, 236
- phase-out 39, 42
- phases-out 35
- phone line 65
- physical capacity 5–6, 19, 21, 24, 26, 28, 30, 34, 112
- physical disc 20–22, 40, 113
 - logical volumes 22
- physical disk 20
- pool size 28–29, 103, 105–106, 108
 - hard 108
 - soft 108
- pool soft size 28

- pool_change_config 111
- pool_delete 111
- pool_rename 111
- pool_resize 111
- power outage 53
- power supply 49, 51–53, 58, 76
- Power Supply Unit (PSU) 58, 212
- Power-off 81
- Power-on 80
- predefined user
 - role 136
- prefetch 17
- primary partition 20, 199
- Proactive phase-out 40
- probe job 218
- problem record 276
- pseudo random distribution
 - MB partitions 198
- pseudo-random 18
- pseudo-random distribution algorithm 20
- pseudo-random distribution function 22

Q

- Quality of Service (QoS) 7

R

- rack 47–48
- rack door 52
- RAID 19, 22–23, 113
- RAID striping 43
- raw capacity 49
- raw read error count 45
- RBAC 135, 170
- read command 122
- readonly 143
- rebuild 19, 39
- redirect on write 6, 19
- redistribution 23, 36, 39, 199
- redundancy 13, 18, 20, 35
- redundancy-supported reaction 36, 45
- Redundant Power Supply (RPS) 3, 48, 51, 56, 58, 64
- regular pool 26–29, 32, 103, 108
- regular Storage Pool 28, 31–32, 117
 - final reserved space 32
 - snapshot reserve space 33
- remote connection 65
- remote mirroring 15, 36, 60, 78
- remote repair 276
- remote support 203–204, 267–268
- report 218
- reserve capacity 27, 29
- resiliency 13, 17, 35
- resize volume 120
- resume 25, 38
- Role Based Access Control 135, 170
- Role Based Access Control (RBAC) 135, 170
- role mapping 151
- roles 132–133, 167
- Rotation Vibration Safeguard (RVS) 61

rule 243–244, 246
RVS 61

S

SAN Volume controller (SVC) 224
SAS adapter 56
SATA disk 48, 61–62
scalability 17, 19
script 85, 99, 130
scrubbing 23, 44
sector count 45
Secure Sockets Layer (SSL) 91
security 131–132
Security Socket Layer (SSL) 131, 187, 293
security-hardened machine 269
self-healing 5, 35–36, 44
Self-Monitoring, Analysis and Reporting Technology (SMART) 44, 62
Self-Protection Throttling (SPT) 62
Serial Advanced Technology Attachment (SATA) 2, 48, 61
Serial-Attached SCSI (SAS) 56–57
Service Agent (SA) 217
service support representative (SSR) 65, 69, 76, 78–79, 133
serviceability 35, 44
severity 204, 206–207
shell 111, 123
Short Message Service (SMS) 193, 204
shutdown 38
shutdown sequence 38
Simple Mail Transfer Protocol
 DNS names 76
Simple Mail Transfer Protocol (SMTP) 69, 76, 192, 239, 246, 267
Simple Network Management Protocol (SNMP) 90, 247
sizing 202
small form-factor pluggable (SFP) 64
SMS 204, 239
 message tokens 245
SMS Gateway 240
SMS gateway 192–193, 245
SMS message 136, 193, 237, 245–246
smsgw_define smsgw 245
SMTP 239, 245
SMTP Gateway 69, 192–193, 240, 245
SMTP gateway 193
Snapshot
 performance 199
snapshot 6, 17, 19, 22, 25, 34
 permission 135
 reserve capacity 27, 29
SNMP 215, 248–249
 destination 242, 249–250
SNMP agent 247–248
SNMP communication 248
SNMP manager 76, 247–250, 256
SNMP trap 247–250, 254, 260
soft capacity 83, 85, 89
soft pool size 29–30
soft size 103, 107–108
soft system size 30, 34
soft_size 112
software services 16
space depletion 25
spare capacity 23–24, 42
spare disk 199
SSIC 10
SSL 91
SSL certificate 188, 190
SSL connection 187, 293, 298, 304
state 120, 122
state_list 209
static allocation 117, 121
statistics 197, 226–227
 monitor 209
statistics_get 235–236
statistics_get command 214, 235–236
Status bar 94, 205
Storage Administrator 18, 20, 26–27, 29–30, 35, 39, 69, 78, 83–84, 91, 133–134, 136, 138–139, 146–147, 152–153, 169, 174–176, 186, 195, 237, 282–283, 285, 290–292, 299
storage capacity 15
Storage Management software 84–85, 94
 installation 85
Storage Networking
 Industry Association 132
Storage Pool 6, 18, 22, 24–35, 42–43, 69, 78, 83, 94, 102–108, 110–114, 116–117, 119–122, 205, 215, 222–224
 and hard capacity limits 31
 available capacity 35
 capacity 26–27
 delete 108
 future activity 105
 logical volumes 26
 overall information 104
 over-provision storage 6
 required size 106
 resize 108
 resource levels 121
 snapshot area 112
 snapshot capacity 107
 snapshot sets 111
 space allocation 35
 system capacity 105, 107
 unused hard capacity 32
 user-created 105
 XCLI 111
Storage Pools 22
storage space 24, 102, 106, 112–113, 117–118, 121–122, 124
 Improved regulation 102
storage system 89, 113, 118, 132, 137, 141, 177
Storage System software 84
storage virtualization 5, 18–20
storageadmin 133–135, 151, 153–154, 169–171, 280–281, 286
striping 200

- SUN Java Directory 151, 168, 174–175, 183–185, 187–188, 279, 287–288, 290–292, 299, 302, 304
 - group 184
 - group creation 184
 - group membership 183, 186
 - group XIVapp01_group 184
 - groups XIVStorageadmin 186
 - new group 183
 - product suit 175
 - such LDAP front end 175
 - user account 286
- Sun Java Systems Directory Server 6
- suspend 25
- switch_list 212
- switching 17
- System level thin provisioning 30–31
- System Planar 55–57
- system quiesce 38
- system services 15
- system size 30, 34, 42
 - hard 30
 - soft 30
- System Storage Interoperability Center (SSIC) 10, 84
- system time 210
- system_capacity_list 209

T

- target volume 122
 - source volume 122
- TCO 44
- technician 133, 135
- telephone line 268
- Thermal Fly-height Control (TFC) 62
- thick-to-thin provisioning 20
- thin provisioning 5–6, 20, 26, 28, 103, 112, 118
 - system level 30
- three-step process 245
- time 210
- time_list 210, 235
- TimeStamp 235
- Tivoli Storage Productivity Center (TPC) 7, 132, 204, 215, 218, 222
- token 245
- toolbar 94, 237, 249
- total cost of ownership (TCO) 44
- TPC 7
- transfer size 200
- Transient system 42
- trap 247–248

U

- UDP 248
- unallocated capacity 26, 28, 31
- uninterruptible power supply (UPS) 3, 38, 52, 93, 132, 192, 211
 - battery charge levels 38
 - current status 211
- unlocked 116, 120, 122
- UPS 3

- UPS module 49, 52
- ups_list 211–212
- usable capacity 50
- USB to Serial 64
- user account 133–136, 143, 146–147, 151, 153–154, 166–167, 170, 173–175, 182, 186, 195, 279–280, 287
- User group 132, 134–137, 139–145, 154–155, 167, 169–173, 176–181, 184, 192
 - Access Control 134, 177
 - Detailed description 170
 - echo Member 172
 - Unauthorized Hosts/Clusters 141, 177
 - user group 136
 - user membership 141
- user group
 - Access Control 140, 177
- User name 91, 101, 132, 134, 138–139, 141, 146, 150, 167, 180, 194, 280–281, 285, 288, 292
 - XIV system limitations 167
- User password 132
- user role 84, 132, 134–135, 143, 151, 167, 169–170
- users
 - maximum 138
- users location 210

V

- VAAI 7
- version_get 210
- virtualization 18–19, 113
- virtualization algorithm 19
- vol_move 111
- Volume 6, 109, 114
 - resize 120
 - state 120, 122
- Volume Shadow Copy Services (VSS) 25
- volume size 27, 103, 116–119, 124
- VPN 268
- vSphere 11
- VSS 25
- vStorage API Array Integration (VAAI) 7

W

- Welcome panel 239, 243

X

- XCLI 7, 9, 28, 69, 76–79, 83–85, 90–91, 95, 99–103, 111, 123–124, 130, 134, 143–144, 151, 154, 169, 172–173, 178, 182, 184–185, 188, 191–192, 194–195, 197, 209–210, 212–213, 226, 235, 237, 245, 247, 269–270
- XCLI command 85, 173, 191, 209
 - event_list 191
 - example 99
- XCLI Session 9, 85, 99–100, 111, 123, 130, 166, 188, 235
- XCLI utility 100–101
- XIV 1–2, 4–6, 8–10, 47–49, 51–52, 59, 61–62, 64–72, 74–79, 150, 152, 203–210, 212–213, 215, 221–224, 237,

- 245, 247–255, 257–260, 267–268, 270, 275–276
- XIV GUI 89, 91, 100, 104, 114, 134, 136, 145–146, 181, 185–186, 188, 191, 222, 237, 247, 249
 - LDAP configuration settings 181, 185
 - LDAP server 188
 - Viewing events 191
- XIV Remote Support Center (XRSC) 77, 268–269
- XIV Role 155, 171–173, 179, 182–183, 186
 - mapping 179, 183
- XIV role
 - mapping 183
- XIV Secure Remote Support (XSRC) 65
- XIV Storage
 - Account 133, 174
 - hardware 84, 90, 93
 - Management 47, 76, 79
 - Management GUI 91, 93, 119, 121
 - Management software 84–86, 99
 - Management software compatibility 84
 - Manager 84, 94
 - Manager GUI 8
 - Manager installation file 86
 - Subsystem 222, 224
 - Subsystem TPC report 222
 - System 1–6, 9, 14–18, 20–26, 30–31, 35–36, 38, 40, 43–45, 47–49, 51, 53, 65–67, 69–71, 75–79, 84, 90–92, 94, 101, 113, 122, 124, 131–136, 139, 145, 147, 149, 169–170, 191, 195, 197–200, 204, 206, 209, 224, 226–227, 235–236, 245, 268
 - System architecture 18
 - System Graphical User Interface 197
 - System installation 70
 - System reliability 43
 - System software 7, 84
 - System time 235
 - System virtualization 113
 - Systems 113, 146, 186
- XIV Storage Manager 84–85, 89
- XIV Storage System 14, 83–84, 89–94, 99, 101, 103, 113, 115, 122, 203–205, 210, 214–215, 221, 224, 237, 245, 247, 249–250, 252, 254, 256, 258, 267, 270
 - administrator 91
 - architecture 5, 14, 16, 18, 20, 35–36, 44–45
 - common command execution syntax 101
 - communicate 76
 - configuration 113
 - data 199
 - design 5
 - distribution algorithm 39
 - family 48
 - Graphical User Interface 197
 - hardware 49, 67, 76
 - installation 66, 133
 - internal operating environment 17, 35
 - iSCSI configuration 75
 - logical architecture 20
 - logical hierarchy 20
 - main GUI window 125
 - management functionality 90
 - Management main window 92
 - Overview 1
 - point 128, 130
 - rack 65, 78
 - reliability 35–36
 - reserves capacity 23
 - software 5–6, 84
 - stripe 199
 - stripes data 198
 - time 199
 - track 199
 - use 101
 - verifie 133
 - virtualization 18, 113
 - volume 200
- XIV subsystem 200
- XIV System 26, 55, 77, 97, 99, 132–133, 146–147, 150–151, 153–154, 170, 175–176, 185, 187–189, 215, 219, 245, 253, 268–269, 284–288, 291–293, 295
- XIV system 145, 150–151, 154, 166, 170, 174, 178, 181–182
 - address range 253, 258–259
 - configuring individual probes 219
 - corresponding parameters 151
 - following information 222
 - fully qualified domain name 217
 - LDAP authenticated user logs 178
 - LDAP authentication 285
 - local repository 167
 - same password 145
 - secure communications 187, 293
 - software component 269
 - user logs 153
- XIV TOP 197, 231
- xiv_development 133, 135
- xiv_maintenance 135
- XIVGUI 7
- XIVTop 7–8
- XSRC 65



IBM XIV Storage System: Architecture, Implementation, and Usage

(0.5" spine)
0.475" <-> 0.873"
250 <-> 459 pages



IBM XIV Storage System Architecture, Implementation, and Usage



**Support provided for
2 TB SATA disk drives**

**Detailed system
configuration
instructions**

**LDAP Configuration
Wizard
Secure Remote
Support**

This IBM Redbooks publication describes the concepts, architecture, and implementation of the IBM XIV Storage System (2810-A14 and 2812-A14), which is designed to be a scalable enterprise storage system based upon a grid array of hardware components. It can attach to both Fibre Channel Protocol (FCP) and iSCSI capable hosts.

In the first few chapters of this book, we provide details about many of the unique and powerful concepts that form the basis of the XIV Storage System logical and physical architecture. We explain how the system was designed to eliminate direct dependencies between the hardware elements and the software that governs the system.

In subsequent chapters, we explain the planning and preparation tasks that are required to deploy the system in your environment. This explanation is followed by a step-by-step procedure of how to configure and administer the system. We provide illustrations of how to perform those tasks by using the intuitive, yet powerful XIV Storage Manager GUI or the Extended Command Line Interface (XCLI).

The book contains comprehensive information about how to integrate the XIV Storage System for authentication in an LDAP environment and outlines the requirements and summarizes the procedures for attaching the system to various host platforms.

We also discuss performance characteristics of the IBM XIV Storage System and present options available for alerting and monitoring, including an enhanced Secure Remote Support capability.

This book is intended for those individuals who want an understanding of the XIV Storage System, and also targets readers who need detailed advice about how to configure and use the system.

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, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7659-02

ISBN 0738435082